



Decision Problems in a Logic for Reasoning About Reconfigurable Distributed Systems

Marius Bozga^(✉) , Lucas Bueri , and Radu Iosif

Univ. Grenoble Alpes, CNRS, Grenoble INP, VERIMAG, 38000 Saint-Martin-d'Hères, France
marius.bozga@univ-grenoble-alpes.fr

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 [37] and Separation Logic [39]. We study the complexity of the satisfiability and entailment problems for the configuration logic under consideration. Additionally, we consider the robustness property of degree boundedness (is every component involved in a bounded number of interactions?), 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 [35]), the distributed systems community has recently put massive effort in designing algorithms for *reconfigurable systems*, whose network topologies change at runtime [23]. However, dynamic reconfiguration in the form of software or network upgrades has been recognized as one of the most important sources of cloud service outage [25].

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 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 [37], or Separation Logic [39]. 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 a robustness property [27], namely *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. For space reasons, the proofs of the technical results are given in [5].

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 a component $[x]@q$, an interaction $\langle x.out, z.in \rangle$ and a separate $\text{chain}_{h',t'}(z, y)$. Figure 1(b) depicts the unfolding of the inductive definition of the token ring, with the




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

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 logic introduced above 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.

The proof in Fig. 1(c) uses *local axioms* specifying, for each basic operation, only those components and interactions required to avoid faulting, with a *frame rule* $\{\phi\} P \{\psi\} \Rightarrow \{\phi * [F]\} P \{\psi * F\}$; for readability, the frame formulae (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 [3] or REO [2]. Many such ADLs have extensions that describe programmed reconfiguration, e.g., [19,30], classified according to the underlying formalism used to define their operational semantics: *process algebras* [13,33], *graph rewriting* [32,41,44], *chemical reactions* [43] (see the surveys [7,11]). Unfortunately, only few ADLs support formal verification, mainly in the flavour of runtime verification [10,17,20,31] or finite-state model checking [14].

Parameterized verification of unbounded networks of distributed processes uses mostly hard-coded coordinating architectures (see [4] for a survey). A first attempt at specifying architectures by logic is the *interaction logic* of Konnov et al. [29], 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 [34]. 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* [26,32,40] 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 [36], initially with a non-interfering (race-free) semantics [8] and later combining ideas of assume-and-rely-guarantee [28,38] with local reasoning [22,42] and abstract notions of framing [15,16,21]. 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) [39]. 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 [18]. 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 (Sect. 4) 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 including zero. For a set A , we define $A^1 \stackrel{\text{def}}{=} A$, $A^{i+1} \stackrel{\text{def}}{=} A^i \times A$, for all $i \geq 1$, 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 $\mathbb{B} = (\mathcal{P}, \mathcal{Q}, \rightarrow)$, where $\mathcal{P}$ is a finite set of *ports*, $\mathcal{Q}$ is a finite set of *states* and $\rightarrow \subseteq \mathcal{Q} \times \mathcal{P} \times \mathcal{Q}$ is a transition relation. Intuitively, each transition $q \xrightarrow{p} q'$ of the behavior 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 $\mathbb{B} = (\mathcal{P}, \mathcal{Q}, \rightarrow)$ are 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.
- $\rho : \mathbb{C} \rightarrow \mathcal{Q}$ 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 $q \in \mathcal{Q}$.

The last condition requires that there is an infinite pool of components in each state $q \in \mathcal{Q}$; since $\mathbb{C}$ is infinite and $\mathcal{Q}$ is finite, this condition is feasible. For example, the configurations of the token ring from Fig. 1(a) are $(\{c_1, \dots, c_n\}, \{(c_i, out, c_{(i \bmod n)+1}, in) \mid$

$i \in [1, n]\}, p)$, where $p : \mathbb{C} \rightarrow \{H, T\}$ 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. Note that the components c_i are not necessary part of the configuration. 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.

Below we define the composition of configurations, as the union of disjoint sets of components and interactions:

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

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 3. *The degree of a configuration $\gamma = (C, I, p)$ 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 $x, y, x_1, \dots \in \mathbb{V}$, $q \in \mathbb{Q}$ 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. These formulæ are also referred to as *atoms*. The connective $*$ is called the *separating conjunction*. We use the shorthand $[x]@q \stackrel{\text{def}}{=} [x] * x@q$. For instance, 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.

A formula is said to be *pure* if and only if it is a separating conjunction 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; we note $\text{fv}(\phi)$ 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 4. 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 $\gamma \models_\Delta^\vee \phi$ 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 (Definition 1):

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

If ϕ is a sentence, the satisfaction relation $\gamma \models_\Delta^\vee \phi$ 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 $\gamma \models^\vee \phi$. 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 $\gamma \models_\Delta^\vee \phi$ only if $\gamma \models_\Delta^\vee \psi$.

2.3 Separation Logic

Separation Logic (SL) [39] 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:

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

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

¹ We use the universe $\mathbb{C}$ here for simplicity, the definition works with any countably infinite set.

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

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 5. 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. $\text{Bnd}[\Delta, A]$: is the set $\{\delta(\gamma) \mid \gamma \models_{\Delta} \exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})\}$ finite?
3. $\text{Entl}[\Delta, A, B]$: does $A(x_1, \dots, x_{\#A}) \models_{\Delta} \exists x_{\#B+1} \dots \exists x_{\#A} . B(x_1, \dots, x_{\#B})$ hold?

The size of a formula ϕ is 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:

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

For a 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 Δ to $\text{arity}(\Delta) \leq k$ and $\text{intertime}(\Delta) \leq \ell$, 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 (Definition 5), the complexity for the case $k < \infty, \ell < \infty$ matches the one for the case $k < \infty, \ell = \infty$.

Satisfiability (1) and entailment (3) 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.

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 [18]. 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 1. The following SID describes star topologies with a central controller connected to an unbounded number of workers stations:

$$\begin{aligned} \text{Controller}(x) &\leftarrow [x] * \text{Worker}(x) \\ \text{Worker}(x) &\leftarrow \exists y . \langle x.out, y.in \rangle * [y] * \text{Worker}(x) \quad \text{Worker}(x) \leftarrow \text{emp} \quad \blacksquare \end{aligned}$$

3 Satisfiability

We show that the satisfiability problem (Definition 5, 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. In the following, the order of terms in (dis-)equalities is not important i.e., we consider $x = y$ (resp. $x \neq y$) and $y = x$ (resp. $y \neq x$) to be the same formula.

Definition 6. 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 Definition 6 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 $x \not\approx_\pi y$.

Base tuples constitute the abstract domain used by the algorithms for checking satisfiability (point 1 of Definition 5) and boundedness (point 2 of Definition 5), defined as follows:

Definition 7. A base tuple is a triple $\mathbf{t} = (C^\sharp, I^\sharp, \pi)$, 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.

Intuitively, a base tuple is an abstract representation of a configuration, where components (resp. interactions) are represented by variables (resp. tuples of variables). 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 atoms 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. We lift the composition, projection and substitution operations to sets of satisfiable base tuples, as usual.

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 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 formulae:

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

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:

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

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

Let $\Delta^\sharp$ 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. Since SatBase is finite, the least solution can be attained in a finite number of steps, using a Kleene iteration (see Fig. 2).

We state below the main result leading to an elementary recursive algorithm for the satisfiability problem (Theorem 1). The intuition is that, if $\mu\vec{X}.\Delta^\sharp(A)$ is not empty, then it contains only satisfiable base tuples, from which a model of $A(x_1, \dots, x_{\#A})$ can be built.

Lemma 1. *$\text{Sat}[\Delta, A]$ has a positive answer if and only if $\mu\vec{X}.\Delta^\sharp(A) \neq \emptyset$.*

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 Definition 7). 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.*

The upper bounds are consequences of the fact that the size of a satisfiable base tuple is bounded by a simple exponential in the $\min(\text{arity}(\Delta), \text{intsize}(\Delta))$, hence the number of such tuples is doubly exponential in $\min(\text{arity}(\Delta), \text{intsize}(\Delta))$. The lower bounds are by a polynomial reduction from the satisfiability problem for SL [9].

Example 2. 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:

$$A_i(x_1, \dots, x_n) \leftarrow \bigstar_{j=0}^{n-i} A_{i+1}(x_1, \dots, x_{i-1}, [x_i, \dots, x_n]^j), \quad \text{for all } i \in [1, n-1]$$

$$A_n(x_1, \dots, x_n) \leftarrow \langle x_1.p, \dots, x_n.p \rangle \quad A_n(x_1, \dots, x_n) \leftarrow \text{emp}$$

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 Degree Boundedness

The boundedness problem (Definition 5, point 2) asks for the existence of a bound on the degree (Definition 3) of the models of a sentence $\exists x_1 \dots \exists x_{\#A} . A(x_1, \dots, x_{\#A})$. Intuitively, the $\text{Bnd}[\Delta, A]$ problem has a negative answer if and only if there are increasingly large unfoldings (i.e., 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. We formalize the notion of unfolding below:

Definition 8. Given a predicate A and a sequence $(r_1, i_1), \dots, (r_n, i_n) \in (\Delta \times \mathbb{N})^+$, where $r_1 : A(x_1, \dots, x_{\#A}) \leftarrow \phi \in \Delta$, 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 9. 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 $(r, i) \in \Delta \times \mathbb{N}$:

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

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). A path $(A_1, t_1) \rightsquigarrow^{(r_1, i_1)} (A_2, t_2) \rightsquigarrow^{(r_2, i_2)} \dots \rightsquigarrow^{(r_n, i_n)} (A_n, t_n)$ in $\mathcal{G}(\Delta)$ induces a unique

input: a SID Δ **output:** $\mathcal{G}(\Delta) = (V, E)$

- 1: initially $V := \emptyset, E := \emptyset$
- 2: **for** $A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m \cdot \phi \in \Delta$, 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** $r : A(x_1, \dots, x_{\#A}) \leftarrow \exists y_1 \dots \exists y_m \cdot \phi * \bigstar_{\ell=1}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell) \in \Delta$ **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: $E := E \cup \{(\langle A, t \rangle, (r, \ell), (B_\ell, t_\ell)) \mid t \in X, \ell \in [1, h]\}$

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

unfolding $A_1(x_1, \dots, x_{\#A_1}) \xrightarrow{(r_1, i_1) \dots (r_n, i_n)}_\Delta \phi$ (Definition 8). 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.

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 $\exists x_1 \dots \exists x_{\#A} \cdot A(x_1, \dots, x_{\#A})$ 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 $B_0(x_1, \dots, x_{\#B_0}) \leftarrow \exists y_1 \dots \exists y_m \cdot \phi * \bigstar_{\ell=2}^h B_\ell(z_1^\ell, \dots, z_{\#B_\ell}^\ell) \in \Delta$, 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:

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

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

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

There exists a family of models (with respect to Δ) of $\exists x_1 \dots \exists x_{\#A} \cdot A(x_1, \dots, x_{\#A})$ of unbounded degree if and only if these are models of $\exists x_1 \dots \exists x_{\#A+1} \cdot A'(x_1, \dots, x_{\#A+1})$ (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 2. *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 $t \in \text{SatBase}$.*

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

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 $k = \infty, \ell = \infty$.*

5 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 (Sect. 1.1).

We show that the entailment problem is undecidable, in general (Thm. 3), 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 10. *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 11. *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 Sect. 1.1, together with rules (4) and (5) is progressing, connected and e-restricted.

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 3. *$\text{Entl}[\Delta, A, B]$ 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).*

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 Proposition 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 [18]. We recall the definition of SL, interpreted over heaps $h : \mathbb{C} \rightarrow_{\text{fin}} \mathbb{C}^{\mathfrak{R}}$, introduced in Sect. 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 (Definition 10).

Definition 12. 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 $\phi = \exists t_1 \dots \exists t_m . x_1 \mapsto (y_1, \dots, y_{\mathfrak{R}}) * \psi$, where ψ contains only predicate and equality atoms,
2. connected if and only if $z_1 \in \{x_i \mid i \in \lambda_{\bar{\Delta}}(\bar{A})\} \cup \{y_1, \dots, y_{\mathfrak{R}}\}$, 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 (Definition 11); 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 Definition 11). A tight upper bound on the complexity of the entailment problem between SL formulae, interpreted by progressing, connected and e-restricted SIDs, is given below:

Theorem 4 ([18]). *The SL entailment problem is in $2^{2^{\text{poly}(\text{width}(\bar{\Delta}) \cdot \log \text{size}(\bar{\Delta}))}}$, 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 a undirected *Gaifman graph*, in which every tuple from a relation (hyperedge) becomes a clique [24]. In a similar vein, we encode a configuration, of degree at most $\mathfrak{B}$, by a heap of degree $\mathfrak{K}$ (Definition 13), such that $\mathfrak{K}$ is defined using the following integer function:

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

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 $\mathbb{B} = (\mathcal{P}, Q, \rightarrow)$ (Sect. 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 $\text{Tuples}_I^j(c) \stackrel{\text{def}}{=} \{\langle c_1, \dots, c_n \rangle \mid (c_1, p_1, \dots, c_n, p_n) \in I, \tau_j = \langle p_1, \dots, p_n \rangle, c \in \{c_1, \dots, c_n\}\}$ be the tuples of components from an interaction of type τ_j from I , that contain a given component c .

Definition 13. Given a configuration $\gamma = (C, I, \rho)$, such that $\delta(\gamma) \leq \mathfrak{B}$, a Gaifman heap for γ is a heap $h : \mathbb{C} \rightarrow_{\text{fin}} \mathbb{C}^{\mathfrak{K}}$, where $\mathfrak{K} \stackrel{\text{def}}{=} \text{pos}(0, M + 1, N)$, $\text{dom}(h) = \text{nodes}(\gamma)$ and, for all $c_0 \in \text{dom}(h)$, such that $h(c_0) = \langle c_1, \dots, c_{\mathfrak{K}} \rangle$, the following hold:

1. $c_1 = c_0$ if and only if $c_0 \in C$,
2. for all $j \in [1, M]$, $\text{Tuples}_I^j(c) = \{\mathbf{c}_1, \dots, \mathbf{c}_s\}$ if and only if there exist integers $0 \leq k_1 < \dots < k_s < \mathfrak{B}$, such that $\langle h(c_0) \rangle_{\text{inter}(k_i, j)} = \mathbf{c}_i$, 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 Definition 13), 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 Definition 13), whereas the last N entries are used to represent the state of the component (condition 3 of Definition 13). 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 3. Figure 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 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 $\eta(x) \in \mathbb{V}^{\mathfrak{K}}$,




Fig. 4. Gaifman Heap for a Chain Configuration

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 $\bar{A}_\ell$ of arity $\#\bar{A}_\ell = (\#A + 1) \cdot \#A$, where $\iota : [1, \#A] \times [1, M] \rightarrow 2^{[0, \#B-1]}$ 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 $\bar{A}_\ell(x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(x_{\#A}))$ (point 2 of Definition 13). Then $\bar{\Delta}$ contains rules of the form:

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

for which Δ has a *stem rule* $A(x_1, \dots, x_{\#(A)}) \leftarrow \exists y_1 \dots \exists y_m \cdot \psi * \pi * \bigstar_{\ell=1}^h B^\ell(z_1^\ell, \dots, z_{\#(B^\ell)}^\ell)$, 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 14. 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 $Y_{i,j} \subseteq [0, \#B-1]$, such that:

- $\|Y_{i,j}\| = \|I_{\psi, \pi}^j(x_i)\|$, where $I_{\psi, \pi}^j(x_i)$ 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 $z_s \approx_\pi x_i$, 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 Definition 13, 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 Definition 13). In particular, the definition of $\text{InterAtoms}_\psi(x_i)$ uses the fact that the rule is well-formed.

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 $\bar{A}_i(x_1, \dots, x_{\#A}, \eta(x_1), \dots, \eta(x_{\#A})) \Vdash_{\bar{\Delta}} \exists x_{\#B+1} \dots \exists x_{\#B} \exists \eta(x_{\#B+1}) \dots \exists \eta(x_{\#B}) \cdot \bar{B}_i(x_1, \dots, x_{\#B}, \eta(x_1), \dots, \eta(x_{\#B}))$, in combination with the upper bound provided by Theorem 4, 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 (Proposition 1), which impacts the size of $\bar{\Delta}$ and time needed to generate it from Δ .

Theorem 5. *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 $4\text{EXP} \cap 2\text{EXP-hard}$.*

6 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 [39], 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. Ahrens, E., Bozga, M., Iosif, R., Katoen, J.: Local reasoning about parameterized reconfigurable distributed systems. CoRR, abs/2107.05253 (2021)
2. Arbab, F.: Reo: a channel-based coordination model for component composition. Math. Struct. Comput. Sci. **14**(3), 329–366 (2004)
3. Basu, A., Bozga, M., Sifakis, J.: Modeling heterogeneous real-time components in BIP. In: Fourth IEEE International Conference on Software Engineering and Formal Methods (SEFM 2006), pp. 3–12. IEEE Computer Society (2006)
4. Bloem, R., et al.: Decidability of Parameterized Verification. Synthesis Lectures on Distributed Computing Theory. Morgan & Claypool Publishers (2015)
5. Bozga, M., Bueri, L., Iosif, R.: Decision problems in a logic for reasoning about reconfigurable distributed systems. CoRR, abs/2202.09637 (2022)
6. Bozga, M., Iosif, R., Sifakis, J.: Verification of component-based systems with recursive architectures. CoRR, abs/2112.08292 (2021)
7. Bradbury, J., Cordy, J., Dingel, J., Wermelinger, M.: A survey of self-management in dynamic software architecture specifications. In: Proceedings of the 1st ACM SIGSOFT workshop on Self-managed systems, pp. 28–33. ACM (2004)
8. Brookes, S., O’Hearn, P.W.: Concurrent separation logic. ACM SIGLOG News **3**(3), 47–65 (2016)

9. Brotherston, J., Fuhs, C., Pérez, J.A.N., Gorogiannis, N.: A decision procedure for satisfiability in separation logic with inductive predicates. In: CSL-LICS, pp. 25:1–25:10. ACM (2014)
10. Bucchiarone, A., Galeotti, J.P.: Dynamic software architectures verification using dynalloy. *Electron. Commun. Eur. Assoc. Softw. Sci. Technol.* **10** (2008). <https://doi.org/10.14279/tuj.eceasst.10.145>
11. Butting, A., Heim, R., Kautz, O., Ringert, J.O., Rumpe, B., Wortmann, A.: A classification of dynamic reconfiguration in component and connector architecture description. In: *Proceedings of MODELS 2017 Satellite Event: Workshops (ModComp)*. CEUR Workshop Proceedings, vol. 2019, pp. 10–16. CEUR-WS.org (2017)
12. Calcagno, C., O’Hearn, P.W., Yang, H.: Local action and abstract separation logic. In: 22nd IEEE Symposium on Logic in Computer Science (LICS 2007), 10–12 July 2007, Wroclaw, Poland, Proceedings, pp. 366–378. IEEE Computer Society (2007)
13. Cavalcante, E., Batista, T.V., Oquendo, F.: Supporting dynamic software architectures: from architectural description to implementation. In: Bass, L., Lago, P., Kruchten, P. (eds.) 12th Working IEEE/IFIP Conference on Software Architecture, WICSA 2015, pp. 31–40. IEEE Computer Society (2015)
14. Clarke, D.: A basic logic for reasoning about connector reconfiguration. *Fundam. Inf.* **82**(4), 361–390 (2008)
15. Dinsdale-Young, T., Birkedal, L., Gardner, P., Parkinson, M., Yang, H.: Views: compositional reasoning for concurrent programs. *SIGPLAN Not.* **48**(1), 287–300 (2013)
16. Dinsdale-Young, T., Dodds, M., Gardner, P., Parkinson, M.J., Vafeiadis, V.: Concurrent abstract predicates. In: D’Hondt, T. (ed.) ECOOP 2010. LNCS, vol. 6183, pp. 504–528. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-14107-2_24
17. Dormoy, J., Kouchnarenko, O., Lanoix, A.: Using temporal logic for dynamic reconfigurations of components. In: Barbosa, L.S., Lumpe, M. (eds.) FACS 2010. LNCS, vol. 6921, pp. 200–217. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-27269-1_12
18. Echenim, M., Iosif, R., Peltier, N.: Unifying decidable entailments in separation logic with inductive definitions. In: Platzer, A., Sutcliffe, G. (eds.) CADE 2021. LNCS (LNAI), vol. 12699, pp. 183–199. Springer, Cham (2021). https://doi.org/10.1007/978-3-030-79876-5_11
19. El-Ballouli, R., Bensalem, S., Bozga, M., Sifakis, J.: Programming dynamic reconfigurable systems. *Int. J. Softw. Tools Technol. Transf.* **23**, 701–719 (2021)
20. El-Hokayem, A., Bozga, M., Sifakis, J.: A temporal configuration logic for dynamic reconfigurable systems. In: Hung, C., Hong, J., Bechini, A., Song, E. (eds.) SAC 2021: The 36th ACM/SIGAPP Symposium on Applied Computing, Virtual Event, Republic of Korea, 22–26 March 2021, pp. 1419–1428. ACM (2021)
21. Farka, F., Nanevski, A., Banerjee, A., Delbianco, G.A., Fábregas, I.: On algebraic abstractions for concurrent separation logics. *Proc. ACM Program. Lang.* **5**(POPL), 1–32 (2021)
22. Feng, X., Ferreira, R., Shao, Z.: On the relationship between concurrent separation logic and assume-guarantee reasoning. In: De Nicola, R. (ed.) ESOP 2007. LNCS, vol. 4421, pp. 173–188. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-71316-6_13
23. Foerster, K., Schmid, S.: Survey of reconfigurable data center networks: enablers, algorithms, complexity. *SIGACT News* **50**(2), 62–79 (2019)
24. Gaifman, H.: On local and non-local properties. *Stud. Log. Found. Math.* **107**, 105–135 (1982)
25. Gunawi, H.S., et al.: Why does the cloud stop computing? Lessons from hundreds of service outages. In: *Proceedings of the Seventh ACM Symposium on Cloud Computing, SoCC 2016*, pp. 1–16. Association for Computing Machinery, New York (2016)
26. Hirsch, D., Inverardi, P., Montanari, U.: Graph grammars and constraint solving for software architecture styles. In: *Proceedings of the Third International Workshop on Software Architecture, ISAW 1998*, pp. 69–72. Association for Computing Machinery, New York (1998)

27. Jansen, C., Katelaan, J., Matheja, C., Noll, T., Zuleger, F.: Unified reasoning about robustness properties of symbolic-heap separation logic. In: Yang, H. (ed.) ESOP 2017. LNCS, vol. 10201, pp. 611–638. Springer, Heidelberg (2017). https://doi.org/10.1007/978-3-662-54434-1_23
28. Jones, C.B.: Developing methods for computer programs including a notion of interference. Ph.D. thesis, University of Oxford, UK (1981)
29. Konnov, I.V., Kotecký, T., Wang, Q., Veith, H., Bliudze, S., Sifakis, J.: Parameterized systems in BIP: design and model checking. In: 27th International Conference on Concurrency Theory, CONCUR 2016, volume 59 of LIPIcs, pp. 30:1–30:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016)
30. Krause, C., Maraïkar, Z., Lazovik, A., Arbab, F.: Modeling dynamic reconfigurations in Reo using high-level replacement systems. *Sci. Comput. Program.* **76**, 23–36 (2011)
31. Lanoix, A., Dormoy, J., Kouchnarenko, O.: Combining proof and model-checking to validate reconfigurable architectures. *Electron. Notes Theor. Comput. Sci.* **279**(2), 43–57 (2011)
32. Le Metayer, D.: Describing software architecture styles using graph grammars. *IEEE Trans. Softw. Eng.* **24**(7), 521–533 (1998)
33. Magee, J., Kramer, J.: Dynamic structure in software architectures. In: ACM SIGSOFT Software Engineering Notes, vol. 21, no. 6, pp. 3–14. ACM (1996)
34. Mavridou, A., Baranov, E., Bliudze, S., Sifakis, J.: Configuration logics: modeling architecture styles. *J. Log. Algebr. Meth. Program.* **86**(1), 2–29 (2017)
35. Noormohammadpour, M., Raghavendra, C.S.: Datacenter traffic control: understanding techniques and tradeoffs. *IEEE Commun. Surv. Tutor.* **20**(2), 1492–1525 (2018)
36. O’Hearn, P.W.: Resources, concurrency, and local reasoning. *Theor. Comput. Sci.* **375**(1–3), 271–307 (2007)
37. O’Hearn, P.W., Pym, D.J.: The logic of bunched implications. *Bull. Symb. Log.* **5**(2), 215–244 (1999)
38. Owicki, S., Gries, D.: An axiomatic proof technique for parallel programs. In: Gries, D. (ed.) *Programming Methodology. Texts and Monographs in Computer Science*, pp. 130–152. Springer, New York (1978). https://doi.org/10.1007/978-1-4612-6315-9_12
39. Reynolds, J.C.: Separation logic: a logic for shared mutable data structures. In: *Proceedings of 17th IEEE Symposium on Logic in Computer Science (LICS 2002)*, 22–25 July 2002, Copenhagen, Denmark, pp. 55–74. IEEE Computer Society (2002)
40. Shtadler, Z., Grumberg, O.: Network grammars, communication behaviors and automatic verification. In: Sifakis, J. (ed.) CAV 1989. LNCS, vol. 407, pp. 151–165. Springer, Heidelberg (1990). https://doi.org/10.1007/3-540-52148-8_13
41. Taentzer, G., Goedicke, M., Meyer, T.: Dynamic change management by distributed graph transformation: towards configurable distributed systems. In: Ehrig, H., Engels, G., Krewowski, H.-J., Rozenberg, G. (eds.) TAGT 1998. LNCS, vol. 1764, pp. 179–193. Springer, Heidelberg (2000). https://doi.org/10.1007/978-3-540-46464-8_13
42. Vafeiadis, V., Parkinson, M.: A marriage of rely/guarantee and separation logic. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007. LNCS, vol. 4703, pp. 256–271. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-74407-8_18
43. Wermelinger, M.: Towards a chemical model for software architecture reconfiguration. *IEE Proc.-Softw.* **145**(5), 130–136 (1998)
44. Wermelinger, M., Fiadeiro, J.L.: A graph transformation approach to software architecture reconfiguration. *Sci. Comput. Program.* **44**(2), 133–155 (2002)

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

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

