



HAL
open science

Testing the Satisfiability of Formulas in Separation Logic with Permissions

Nicolas Peltier

► **To cite this version:**

Nicolas Peltier. Testing the Satisfiability of Formulas in Separation Logic with Permissions. TABLEAUX 2023 32nd International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, Sep 2023, Prague, Czech Republic. hal-04163846

HAL Id: hal-04163846

<https://hal.science/hal-04163846v1>

Submitted on 17 Jul 2023

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.

Testing the Satisfiability of Formulas in Separation Logic with Permissions

Nicolas Peltier

Université Grenoble Alpes, LIG, CNRS, Inria, Grenoble INP, F-38000 Grenoble, France

Abstract. We investigate the satisfiability problem for a fragment of Separation Logic (SL) with inductively defined spatial predicates and permissions. We show that the problem is undecidable in general, but decidable under some restrictions on the rules defining the semantics of the spatial predicates. Furthermore, if the satisfiability of permission formulas can be tested in exponential time for the considered permission model then SL satisfiability is EXPTIME complete.

1 Introduction

Separation Logic [14, 22] (SL) is a dialect of bunched logic [18] that is widely used in verification for reasoning on programs manipulating pointer-based data structures. It constitutes the theoretical basis of several industrial scale automated static program analyzers [1, 7, 2]. SL formulas describe *heaps*, with atoms asserting that some location (i.e., a memory address) is allocated and refers to some tuple of locations (i.e., a record), combined with a special connective $*$, called *separating conjunction*, which is used to compose heaps. Custom data structures may be described in this setting by using spatial predicates, the semantics of which is defined using *inductive rules*, similar to those used for defining recursive structures in usual programming languages. Such rules allow one to describe heaps of unbounded size with some particular structure such as lists or trees. In this setting, existing work usually focuses on the fragment of SL called *symbolic heaps* (defined as separating conjunctions of SL atoms).

Usually, SL formulas are interpreted in the *standard heap model*, where heaps are defined as partial finite functions mapping locations to tuples of locations and where the separating conjunction $*$ is interpreted as the disjoint union of heaps. Both the satisfiability and entailment problems have been extensively investigated for this heap model. It was proven that the satisfiability problem is EXPTIME complete [6], whereas the entailment problem is undecidable in general, and 2-EXPTIME complete provided the inductive rules meet some syntactic conditions [13, 11, 12, 15] which are general enough to capture usual data structures used in programming. The combination of spatial reasoning with theory reasoning has also been thoroughly investigated, see for instance [20, 21, 19, 23, 16]).

However, richer models exist (see for instance [8]) accounting for additional features of dynamic memory. The automation of reasoning in these models received little attention. One such model that is of practical relevance is *separation logic with permissions* [3, 5], where allocated locations are associated with so called *permissions* used to model the ownership of a given heap region (e.g., a process may have `read` or `write` permission over some location). The heap composition operator that is used to define the interpretation of the separating conjunction is more complex in this framework than in the above case: non disjoint heaps can be combined if they agree on all the locations on which they are both defined and if the corresponding permissions can be combined (for instance it is natural to assume that `read` permissions can be freely combined but not `write` permissions). The framework is thus parameterized by some *permission model* describing which permissions are available and how they can be combined. In [10] algorithms are provided to decide the satisfiability and entailment problems for SL formulas (symbolic heaps) with permissions in the case of lists, i.e., when all allocated locations refer to a single location (i.e., to a record of size 1) and when there is only one spatial predicate $\text{lseg}_p(x, y)$ denoting a list segment from x to y , with permission p . The provided algorithms are generic w.r.t. the permission model, and it is proven that these problems are in NP and co-NP, respectively, assuming that some oracle exists for testing the satisfiability of permission formulas in the considered model.

In the present paper, we investigate the satisfiability problem for SL formulas with permission defined over arbitrary spatial predicates, with user-defined inductive rules. The goal is to allow for more genericity by tackling custom data structures (such as trees, cyclic lists, doubly linked lists etc.) with arbitrary permissions. The addition of permissions makes satisfiability testing much more difficult: we prove that the problem is undecidable in general, and we devise syntactic conditions on the inductive rules for which the problem is EXPTIME-complete. The restrictions are similar – but stronger – to those given in [13] to ensure the decidability of the entailment problem in the standard heap model. In particular, the inductive rules defining the predicate `lseg` mentioned above fulfill these restrictions¹, as well as other usual data structures such as cyclic list, trees etc. (however, doubly linked lists or trees with parent links are not captured). The considered inductive rules use a special connective \circ (different from $*$) that is interpreted as a disjoint union. As we shall see, this is both more natural for defining data structures (see also [5]) and required for deciding satisfiability. Due to lack of space, proofs can be found in the appendix.

2 Definitions

Syntax. We first briefly review some basic notations. If \mathbf{x} and \mathbf{y} are finite sequences, then we denote by $\mathbf{x}\mathbf{y}$ the concatenation of \mathbf{x} and \mathbf{y} . We denote by $|\mathbf{x}|$ the length of \mathbf{x} and by \mathbf{x}_i its i -th element (if $1 \leq i \leq |\mathbf{x}|$). If $E \subseteq \{1, \dots, |\mathbf{x}|\}$ then $\mathbf{x}|_E$ denotes the set $\{\mathbf{x}_i \mid i \in E\}$. With a slight abuse of notations, a finite

¹ provided the considered lists are not empty.

sequence \mathbf{x} is sometimes identified with the set $\{\mathbf{x}|_i \mid i = 1, \dots, |\mathbf{x}|\}$, for instance, we may write $x \in (\mathbf{u} \cup \mathbf{v}) \setminus \mathbf{w}$ to state that x occurs in \mathbf{u} or \mathbf{v} but not in \mathbf{w} .

We consider a multisorted framework, with two sorts \mathbf{l} (for locations) and \mathbf{p} (for permissions). Let $\mathcal{V}_{\mathbf{l}}$ and $\mathcal{V}_{\mathbf{p}}$ be two countably infinite disjoint sets of *variables* with $\mathcal{V} \stackrel{\text{def}}{=} \mathcal{V}_{\mathbf{l}} \cup \mathcal{V}_{\mathbf{p}}$, where $\mathcal{V}_{\mathbf{l}}$ and $\mathcal{V}_{\mathbf{p}}$ denote location variables and permission variables, respectively. The set of *permission terms* $\mathcal{T}_{\mathbf{p}}$ denotes the set of terms built inductively as usual on the set of variables $\mathcal{V}_{\mathbf{p}}$ and the binary function \oplus (written in infix notation). A *points-to atom* is an expression of the form $x \xrightarrow{p} (y_1, \dots, y_k)$ with $x, y_1, \dots, y_k \in \mathcal{V}_{\mathbf{l}}$ and $p \in \mathcal{T}_{\mathbf{p}}$. An *equational atom* is an expression of the form $x \simeq y$ or $x \not\simeq y$ with either $x, y \in \mathcal{V}_{\mathbf{l}}$ or $x, y \in \mathcal{T}_{\mathbf{p}}$.

We consider two disjoint sets of predicate symbols $\mathcal{P}_{\mathbf{p}}$ and \mathcal{P} . The set $\mathcal{P}_{\mathbf{p}}$ denotes *permission predicates*, where each predicate $\hat{P} \in \mathcal{P}_{\mathbf{p}}$ is associated with a unique arity $\#(\hat{P})$. A *permission atom* is an expression of the form $\hat{P}(p_1, \dots, p_n)$, $\hat{P} \in \mathcal{P}_{\mathbf{p}}$, $n = \#(\hat{P})$ and $p_1, \dots, p_n \in \mathcal{T}_{\mathbf{p}}$. \mathcal{P} is a finite set of *spatial predicate symbols*. Each symbol $P \in \mathcal{P}$ is associated with a *spatial arity* $\#_{\mathbf{l}}(P) \in \mathbb{N}$ and with an *arity* $\#(P) \in \mathbb{N}$, with $\#(P) > \#_{\mathbf{l}}(P) > 0$ ($\#_{\mathbf{l}}(P)$ and $\#(P) - \#_{\mathbf{l}}(P)$ denote the number of arguments of P that are of sort \mathbf{l} and \mathbf{p} , respectively). A *predicate atom* is an expression of the form $P(x_1, \dots, x_n, p_1, \dots, p_m)$, with $n = \#_{\mathbf{l}}(P)$, $n + m = \#(P)$, $x_1, \dots, x_n \in \mathcal{V}_{\mathbf{l}}$ and $p_1, \dots, p_m \in \mathcal{T}_{\mathbf{p}}$. A *spatial atom* is either a points-to atom or a predicate atom.

The set of *formulas* is built inductively as usual on the logical constants \mathbf{emp} , and \perp and on the set of spatial, equational and permission atoms, using the special connectives $*$ and \circ and existential quantification on variables of sort \mathbf{l} only (existential quantification over variables of type \mathbf{p} is not allowed). The connective $*$ is usually called *separating conjunction*, and we call \circ the *disjoint conjunction* (it is intended to capture the disjoint union of heaps²). Formulas are taken up to associativity and commutativity of the symbols $*$ and \circ , up to the commutativity of $\simeq, \not\simeq$ and up to prenex form. We denote by $|\phi|$ the size of ϕ . For technical convenience, we assume that the symbols \circ and $*$ have weight of 1 and 2, respectively, and that all atoms have size 1. For conciseness, a formula $\exists x_1 \dots \exists x_n \phi$ will often be written $\exists \mathbf{x} \phi$, with $\mathbf{x} = (x_1, \dots, x_n)$. A *permission formula* is a formula containing no spatial atoms and no equational atom of the form $x \simeq y$ or $x \not\simeq y$ with $x, y \in \mathcal{V}_{\mathbf{l}}$ (note that \mathbf{emp} is a permission formula). A formula is *spatial* if all the atoms occurring in it are spatial. A *pure formula* is a formula that contains no spatial atom (it is not necessarily a permission formula, as it may contain equations or disequations between locations) A *symbolic heap* is a formula containing no occurrence of \circ , and a \circ -*formula* is a formula containing no occurrence of $*$.

A variable x is *free* in a formula ϕ if it occurs in ϕ outside of the scope of any quantifier binding x . The set of variables (freely) occurring in a term (or formula) ϕ is denoted by $fv(\phi)$. A *substitution* is a function mapping every variable in $\mathcal{V}_{\mathbf{l}}$ to a variable in $\mathcal{V}_{\mathbf{l}}$ and every variable in $\mathcal{V}_{\mathbf{p}}$ to a term in $\mathcal{T}_{\mathbf{p}}$. The

² The connective \circ is called *strong separating conjunction* in [5] and written $*$ (whereas $*$ is written \circledast). Our notations are mostly consistent with those in [10].

domain of a substitution σ (denoted by $\text{dom}(\sigma)$) is the set of variables x such that $\sigma(x) \neq x$. A substitution of domain $\{x_1, \dots, x_n\}$ with $\sigma(x_i) = t_i$ is denoted by $\{x_i \leftarrow t_i \mid i = 1, \dots, n\}$, or $\{\mathbf{x} \leftarrow \mathbf{t}\}$, with $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{t} = (t_1, \dots, t_n)$. For all formulas or terms ϕ , we denote by $\phi\sigma$ the formula or term obtained from ϕ by replacing every free occurrence of a variable x by $\sigma(x)$.

Semantics. Permissions are interpreted in some permission model:

Definition 1 (Adapted from [10]). A permission model \mathfrak{P} is a triple

$$(\mathcal{P}_{\mathfrak{P}}, \oplus_{\mathfrak{P}}, (\hat{P}_{\mathfrak{P}})_{\hat{P} \in \mathcal{P}_{\mathfrak{P}}})$$

where $\mathcal{P}_{\mathfrak{P}}$ is a non empty set, called the set of permissions, $\oplus_{\mathfrak{P}} : \mathcal{P}_{\mathfrak{P}}^2 \rightarrow \mathcal{P}_{\mathfrak{P}}$ is a binary partial function that is commutative, associative and cancellative, and $\hat{P}_{\mathfrak{P}} \subseteq \mathcal{P}_{\mathfrak{P}}^{\#(\hat{P})}$, for all $\hat{P} \in \mathcal{P}_{\mathfrak{P}}$. If $\pi, \pi' \in \mathcal{P}_{\mathfrak{P}}$, we write $\pi \leq_{\mathfrak{P}} \pi'$ if $\pi = \pi' \vee (\exists \pi'' \in \mathcal{P}_{\mathfrak{P}} \pi' = \pi \oplus_{\mathfrak{P}} \pi'')$.

In what follows, \mathfrak{P} always denotes a permission model. If $\pi \in \mathcal{P}_{\mathfrak{P}}$ and $n \in \mathbb{N}$, we denote by π^n the permission $\pi \oplus_{\mathfrak{P}} \dots \oplus_{\mathfrak{P}} \pi$ (n times), note that π^n is not necessarily defined and implicitly depends on the considered permission model, which will always be clear from the context. In contrast to [10], we do not assume that a maximal “total” permission $1_{\mathfrak{P}}$ exists, we allow instead for arbitrary predicates over permissions (the total permission can be encoded as a unary predicate symbol T , with $T_{\mathfrak{P}} = \{1_{\mathfrak{P}}\}$).

Example 2. Assume that $\mathcal{P}_{\mathfrak{p}} = \emptyset$. A simple example of permission model is $\mathfrak{w} = (\{\mathbf{read}, \mathbf{write}\}, \oplus_{\mathfrak{w}}, \emptyset)$, with $\mathbf{read} \oplus_{\mathfrak{w}} \mathbf{read} = \mathbf{read}$ and $\mathbf{write} \oplus_{\mathfrak{w}} \pi$ is undefined for all $\pi \in \{\mathbf{read}, \mathbf{write}\}$. Another example (from [4]) is $\mathfrak{f} = (]0, 1], \oplus_{\mathfrak{f}}, \emptyset)$ where $]0, 1]$ denotes the interval of rational numbers, with $\pi \oplus_{\mathfrak{f}} \pi' = \pi + \pi'$ if $\pi + \pi' \leq 1$ and $\pi \oplus_{\mathfrak{f}} \pi'$ is undefined otherwise (\mathfrak{f} stands for *fractional*).

Let \mathcal{L} be a countably infinite set of *locations*. A *store* (for a given permission model \mathfrak{P}) is a total mapping associating every variable in \mathcal{V}_1 to an element of \mathcal{L} and every variable in $\mathcal{V}_{\mathfrak{p}}$ to an element of $\mathcal{P}_{\mathfrak{P}}$. A store can be extended into a partial mapping from $\mathcal{T}_{\mathfrak{p}}$ to $\mathcal{P}_{\mathfrak{P}}$ inductively defined as follows: $\mathfrak{s}(p_1 \oplus p_2) \stackrel{\text{def}}{=} \mathfrak{s}(p_1) \oplus_{\mathfrak{P}} \mathfrak{s}(p_2)$. Note that the obtained mapping is partial since $\mathfrak{s}(p_1) \oplus_{\mathfrak{P}} \mathfrak{s}(p_2)$ is not always defined. If x_1, \dots, x_n are pairwise distinct variables in \mathcal{V}_1 and $\ell_1, \dots, \ell_n \in \mathcal{L}$, we denote by $\mathfrak{s}\{x_i \leftarrow \ell_i \mid i = 1, \dots, n\}$ the store \mathfrak{s}' coinciding with \mathfrak{s} on every variable not occurring in $\{x_1, \dots, x_n\}$ and such that $\mathfrak{s}'(x_i) = \ell_i$ for all $i = 1, \dots, n$.

A *heap* (for a given permission model \mathfrak{P}) is a partial finite function from \mathcal{L} to $\mathcal{L}^* \times \mathcal{P}_{\mathfrak{P}}$. The domain of a heap \mathfrak{h} is denoted by $\text{dom}(\mathfrak{h})$, and we denote by $|\mathfrak{h}|$ the finite cardinality of $\text{dom}(\mathfrak{h})$. A heap of domain ℓ_1, \dots, ℓ_n such that $\mathfrak{h}(\ell_i) = (\ell_1^i, \dots, \ell_{k_i}^i, \pi_i)$ (for all $i \in \{1, \dots, n\}$) will be denoted as a set $\{(\ell_i, \ell_1^i, \dots, \ell_{k_i}^i, \pi_i) \mid i = 1, \dots, n\}$. For every heap \mathfrak{h} we denote by $\text{loc}(\mathfrak{h})$ the set $\{\ell_i \mid \ell_0 \in \text{dom}(\mathfrak{h}), \mathfrak{h}(\ell_0) = (\ell_1, \dots, \ell_k, \pi), 0 \leq i \leq k\}$. A heap may be viewed as a directed (labeled) graph: the locations in $\text{loc}(\mathfrak{h})$ are the vertices of the graph

and there is a edge from ℓ to ℓ' if $\mathfrak{h}(\ell) = (\ell_1, \dots, \ell_n, \pi)$ and $\ell' = \ell_i$ for some $i \in \{1, \dots, n\}$.

A *subheap* of \mathfrak{h} is any heap \mathfrak{h}' such that $\text{dom}(\mathfrak{h}') \subseteq \text{dom}(\mathfrak{h})$ and $\mathfrak{h}'(\ell) = \mathfrak{h}(\ell)$ for all $\ell \in \text{dom}(\mathfrak{h}')$. A *p-weakening* of \mathfrak{h} (w.r.t. some permission model \mathfrak{P}) is any heap \mathfrak{h}' such that $\text{dom}(\mathfrak{h}') = \text{dom}(\mathfrak{h})$ and for all $\ell \in \text{dom}(\mathfrak{h})$, if $\mathfrak{h}(\ell) = (\ell_1, \dots, \ell_n, \pi)$ then $\mathfrak{h}'(\ell) = (\ell_1, \dots, \ell_n, \pi')$ with $\pi' \leq_{\mathfrak{P}} \pi$. We write $\mathfrak{h}' \leq_1 \mathfrak{h}$ (resp. $\mathfrak{h}' \leq_p \mathfrak{h}$) if \mathfrak{h}' is a subheap (resp. a p-weakening) of \mathfrak{h} . The relation \leq denotes the composition of \leq_1 and \leq_p . We write $\mathfrak{h} \sim \mathfrak{h}'$ if \mathfrak{h} and \mathfrak{h}' only differ by the permissions, i.e., $\text{dom}(\mathfrak{h}) = \text{dom}(\mathfrak{h}')$ and for all $\ell \in \text{dom}(\mathfrak{h})$, if $\mathfrak{h}'(\ell) = (\ell_1, \dots, \ell_n, \pi')$ then there exists π such that $\mathfrak{h}(\ell) = (\ell_1, \dots, \ell_n, \pi)$.

Example 3. Consider the permission model \mathfrak{f} defined in Example 2 with $\mathcal{L} = \mathbb{N}$. Then

$$\begin{aligned} \mathfrak{h}_0 &= \{(0, 0, 1, 0.1), (1, 0, 0, 0.2)\}, & \mathfrak{h}_1 &= \{(0, 0, 1, 0.1)\}, \\ \mathfrak{h}_2 &= \{(0, 0, 1, 0.1), (1, 0, 0, 0.1)\} & \mathfrak{h}_3 &= \{(1, 0, 0, 0.1)\} \end{aligned}$$

are heaps, and we have, e.g., $\mathfrak{h}_0(0) = (0, 1, 0.1)$ (meaning that the location 0 is allocated and refers to (0, 1), with permission 0.1), $\mathfrak{h}_1 \leq_1 \mathfrak{h}_0$, $\mathfrak{h}_2 \leq_p \mathfrak{h}_0$, $\mathfrak{h}_3 \leq_1 \mathfrak{h}_2$, and $\mathfrak{h}_3 \leq \mathfrak{h}_0$. Moreover, $\mathfrak{h}_0 \sim \mathfrak{h}_2$.

Heaps can be composed using the following partial operator. If $\mathfrak{h}_1, \mathfrak{h}_2$ are heaps, then $\mathfrak{h}_1 \sqcup \mathfrak{h}_2$ is defined iff for all $\ell \in \text{dom}(\mathfrak{h}_1) \cap \text{dom}(\mathfrak{h}_2)$, we have $\mathfrak{h}_i(\ell) = (\ell_1^i, \dots, \ell_{k_i}^i, \pi_i)$ (for all $i = 1, 2$) where $k_1 = k_2$, $\ell_j^1 = \ell_j^2$ for all $j \in \{1, \dots, k_1\}$ and $\pi_1 \oplus_{\mathfrak{P}} \pi_2$ is defined. Then $\mathfrak{h}_1 \sqcup \mathfrak{h}_2$ is defined as follows: if $\ell \in \text{dom}(\mathfrak{h}_i) \setminus \text{dom}(\mathfrak{h}_j)$ with $(i, j) \in \{(1, 2), (2, 1)\}$ then $(\mathfrak{h}_1 \sqcup \mathfrak{h}_2)(\ell) \stackrel{\text{def}}{=} \mathfrak{h}_i(\ell)$, and if $\ell \in \text{dom}(\mathfrak{h}_1) \cap \text{dom}(\mathfrak{h}_2)$ then $(\mathfrak{h}_1 \sqcup \mathfrak{h}_2)(\ell) \stackrel{\text{def}}{=} (\ell_1^1, \dots, \ell_{k_1}^1, \pi_1 \oplus_{\mathfrak{P}} \pi_2)$.

Example 4. Consider the permission model \mathfrak{f} defined in Example 2, with $\mathcal{L} = \mathbb{N}$ and the following heaps:

$$\begin{aligned} \mathfrak{h}_0 &= \{(0, 0, 0.5), (1, 0, 0.6)\} & \mathfrak{h}_1 &= \{(0, 0, 0.5), (1, 0, 0.2), (2, 0.1)\} \\ \mathfrak{h}_2 &= \{(0, 0, 0.5), (1, 0, 0.6)\} & \mathfrak{h}_3 &= \{(0, 0, 0.1), (1, 0.1)\} \end{aligned}$$

Then $\mathfrak{h}_0 \sqcup \mathfrak{h}_1$ is defined, and we have: $\mathfrak{h}_0 \sqcup \mathfrak{h}_1 = \{(0, 0, 1), (1, 0, 0.8), (2, 0.1)\}$. However, neither $\mathfrak{h}_0 \sqcup \mathfrak{h}_2$ nor $\mathfrak{h}_0 \sqcup \mathfrak{h}_3$ is defined (in the former case the permissions of location 1 cannot be combined (as $0.6 + 0.6 > 1$) and in the latter case the location 1 is associated with distinct tuples, (0) and (1), respectively).

A *structure* (for a given permission model \mathfrak{P}) is a pair $(\mathfrak{s}, \mathfrak{h})$ where \mathfrak{s} is a store and \mathfrak{h} is a heap for \mathfrak{P} . It is *injective* if \mathfrak{s} is injective. A location ℓ is *allocated* in a structure $(\mathfrak{s}, \mathfrak{h})$ or in a heap \mathfrak{h} if $\ell \in \text{dom}(\mathfrak{h})$, and a variable x is *allocated* in $(\mathfrak{s}, \mathfrak{h})$ if $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$.

The semantics of spatial predicate is defined by inductive rules. A *set of inductive definitions* (SID) is a set of *rules* of the form $P(x_1, \dots, x_n, y_1, \dots, y_m) \Leftarrow \phi$ where $n = \#_1(P)$, $n + m = \#(P)$, x_1, \dots, x_n are pairwise distinct variables in \mathcal{V}_1 , y_1, \dots, y_m are pairwise distinct variables in \mathcal{V}_p , and ϕ is a formula such that $\text{fv}(\phi) \subseteq \{x_1, \dots, x_n, y_1, \dots, y_m\}$. We write $P(z_1, \dots, z_n, p_1, \dots, p_m) \Leftarrow_{\mathcal{R}} \psi$ iff \mathcal{R} contains a rule $P(x_1, \dots, x_n, y_1, \dots, y_m) \Leftarrow \phi$ with $\psi = \phi\{x_i \leftarrow z_i, y_j \leftarrow p_j \mid i \in \{1, \dots, n\}, j \in \{1, \dots, m\}\}$.

Definition 5. (Semantics) For every permission model \mathfrak{P} and SID \mathcal{R} , the satisfiability relation $\models_{\mathcal{R}}^{\mathfrak{P}}$ is the smallest relation between structures (for \mathfrak{P}) and formulas such that:

1. $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \mathbf{emp}$ iff $\mathfrak{h} = \emptyset$.
2. $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} x \xrightarrow{p} (y_1, \dots, y_k)$ if $\mathfrak{s}(p)$ is defined and $\mathfrak{h} = \{(\mathfrak{s}(x), \mathfrak{s}(y_1), \dots, \mathfrak{s}(y_k), \mathfrak{s}(p))\}$. Note that this entails that $\text{dom}(\mathfrak{h}) = \{\mathfrak{s}(x)\}$.
3. $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} x \simeq y$ (resp. $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} x \not\simeq y$) if $\mathfrak{h} = \emptyset$, $\mathfrak{s}(x)$ and $\mathfrak{s}(y)$ are defined and $\mathfrak{s}(x) = \mathfrak{s}(y)$ (resp. $\mathfrak{s}(x) \neq \mathfrak{s}(y)$).
4. $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \hat{P}(p_1, \dots, p_n)$ with $\hat{P} \in \mathcal{P}_{\mathfrak{P}}$ if $\mathfrak{s}(p_i)$ is defined for all $i \in \{1, \dots, n\}$, $(\mathfrak{s}(p_1), \dots, \mathfrak{s}(p_n)) \in \hat{P}_{\mathfrak{P}}$ and $\mathfrak{h} = \emptyset$.
5. $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} P(x_1, \dots, x_n, \pi_1, \dots, \pi_m)$ with $P \in \mathcal{P}$ if there exists ϕ such that $P(x_1, \dots, x_n, \pi_1, \dots, \pi_m) \leftarrow_{\mathcal{R}} \phi$ and $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$.
6. $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_1 * \phi_2$ if there exist heaps $\mathfrak{h}_1, \mathfrak{h}_2$ such that $\mathfrak{h}_1 \sqcup \mathfrak{h}_2$ is defined, $\mathfrak{h} = \mathfrak{h}_1 \sqcup \mathfrak{h}_2$ and $(\mathfrak{s}, \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$ for all $i = 1, 2$.
7. $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_1 \circ \phi_2$ if there exists heaps $\mathfrak{h}_1, \mathfrak{h}_2$ such that $\text{dom}(\mathfrak{h}_1) \cap \text{dom}(\mathfrak{h}_2) = \emptyset$, $\mathfrak{h} = \mathfrak{h}_1 \sqcup \mathfrak{h}_2$ and $(\mathfrak{s}, \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$ for all $i = 1, 2$.
8. $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \exists x \phi$ if $(\mathfrak{s}\{x \leftarrow \ell\}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$ for some $\ell \in \mathcal{L}$.

A structure $(\mathfrak{s}, \mathfrak{h})$ such that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$ is an $(\mathcal{R}, \mathfrak{P})$ -model of ϕ . A formula admitting an $(\mathcal{R}, \mathfrak{P})$ -model is $(\mathcal{R}, \mathfrak{P})$ -satisfiable. Two formulas are sat-equivalent (w.r.t. $\mathcal{R}, \mathfrak{P}$) if they are both $(\mathcal{R}, \mathfrak{P})$ -satisfiable or both $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable.

Example 6. The formula $x \overset{u}{\mapsto} (y, z) \circ x \overset{u'}{\mapsto} (y', z')$ is $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable, as x cannot be allocated in disjoint parts of the heap. $x \overset{u}{\mapsto} (y) * x \overset{u'}{\mapsto} (y') * y \not\simeq y'$ is also $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable, as x cannot refer to two distinct records, but $x \overset{u}{\mapsto} (y, z) * x \overset{u'}{\mapsto} (y', z')$ admits the model (on the permission model \mathfrak{f}) $(\mathfrak{s}, \mathfrak{h})$ with $\mathfrak{s}(x) = 0$, $\mathfrak{s}(y) = \mathfrak{s}(y') = 1$, $\mathfrak{s}(z) = \mathfrak{s}(z') = 2$, $\mathfrak{s}(u) = 0.5$, $\mathfrak{s}(u') = 0.2$ and $\mathfrak{h} = \{(0, 1, 2, 0.7)\}$.

Note that there is no logical constant \top (true): no formula can be satisfied on all heaps. The constant \mathbf{emp} is similar to \top but it states that the heap is empty. For all formulas ϕ, ψ , we write $\phi \models_{\mathcal{R}}^{\mathfrak{P}} \psi$ iff the implication $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi \implies (\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \psi$ holds for all structures $(\mathfrak{s}, \mathfrak{h})$, and $\phi \equiv_{\mathcal{R}}^{\mathfrak{P}} \psi$ iff we have both $\phi \models_{\mathcal{R}}^{\mathfrak{P}} \psi$ and $\psi \models_{\mathcal{R}}^{\mathfrak{P}} \phi$. If ϕ contains no predicate symbols in \mathcal{P} , then the truth value of ϕ in $(\mathfrak{s}, \mathfrak{h})$ does not depend on \mathcal{R} . We thus may write $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} \phi$ instead of $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$. If, moreover, ϕ is pure, then $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} \phi$ holds only if \mathfrak{h} is empty. We will write $\mathfrak{s} \models^{\mathfrak{P}} \phi$ to state that $(\mathfrak{s}, \emptyset) \models^{\mathfrak{P}} \phi$. Finally, if ϕ contains only equalities between variables then its semantics does not depend on \mathcal{R} and \mathfrak{P} thus we write $\mathfrak{s} \models \phi$ to state that $(\mathfrak{s}, \emptyset) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$. Note that the semantics of $\phi_1 \circ \phi_2$ and $\phi_1 * \phi_2$ coincide if ϕ_1 or ϕ_2 is pure, and also coincide with that of the usual standard conjunction if both ϕ_1 and ϕ_2 are pure.

Shorthands. If $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{y} = (y_1, \dots, y_m)$ are sequences of variables in \mathcal{V}_1 then $\mathbf{x} \simeq \mathbf{y}$ denotes the formula \perp if $n \neq m$ and $(x_1 \simeq y_1) \circ \dots \circ (x_n \simeq y_n)$ otherwise. For every permission term p , we denote by $\text{def}(p)$ the atom $p \simeq p$. By definition, $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \text{def}(p)$ iff $\mathfrak{s}(p)$ is defined and $\mathfrak{h} = \emptyset$.

3 \mathfrak{h} -Regular Systems

We focus on SIDs of some particular form, defined below.

Definition 7. *A rule is \mathfrak{h} -regular if it is of the following form:*

$$P(\mathbf{x}, \mathbf{y}) \Leftarrow \exists u_1, \dots, u_n (x \xrightarrow{p} (v_1, \dots, v_k) \circ Q_1(u_1, \mathbf{y}_1) \dots \circ Q_n(u_n, \mathbf{y}_n) \circ \phi)$$

where $\{u_1, \dots, u_n\} \subseteq \{v_1, \dots, v_k\}$, \mathbf{y}_i is a vector of variables³, $Q_i \in \mathcal{P}$ and ϕ is pure. We assume by α -renaming that x, \mathbf{y} do not occur in $\{u_1, \dots, u_n\}$. A SID \mathcal{R} is \mathfrak{h} -regular if all the rules in \mathcal{R} are \mathfrak{h} -regular.

Note that the right-hand side formula contains only the disjoint separation connective \circ and not the usual separating conjunction $*$. As we will see (Theorem 33) this is crucial for the decidability of the satisfiability problem. However, as already observed in [5], this is also justified from a practical point of view. Assume for instance that we want to define the predicate lseg introduced in [10], denoting a list segment from x to y with some permission z . The following rules can be used⁴: $\text{lseg}(x, y, z) \Leftarrow x \xrightarrow{z} (y)$ $\text{lseg}(x, y, z) \Leftarrow \exists u (x \xrightarrow{z} (u) \circ \text{lseg}(u, y, z))$. A structure $(\mathfrak{s}, \mathfrak{h})$ satisfies $\text{lseg}(x, y, z)$ if $\mathfrak{h} = \{(\ell_i, \ell_{i+1}, \mathfrak{s}(z)) \mid i = 1, \dots, n\}$ with $n > 0$, $\mathfrak{s}(x) = \ell_1$, $\mathfrak{s}(y) = \ell_{n+1}$ and $\ell_i \neq \ell_j$ if $i \neq j$ and $i, j \in \{1, \dots, n\}$. This fits in with the definition in [10] (except that $n > 0$). In contrast, if one uses instead the connective $*$: $\text{lseg}(x, y, z) \Leftarrow \exists u (x \xrightarrow{z} (u) * \text{lseg}(u, y, z))$, then one could obtain models where the list “loops” on itself an arbitrary number of times, such as, for instance $(\mathfrak{s}, \{(\mathfrak{s}(x), \mathfrak{s}(x), p)\})$, with $\mathfrak{s}(y) = \mathfrak{s}(x)$ and $p = \mathfrak{s}(z)^n$, for any $n > 0$ such that $\mathfrak{s}(z)^n$ is defined. In the former definition, $\mathfrak{s}(y)$ possibly occurs in $\{\ell_1, \dots, \ell_n\}$, but each location can only be allocated once.

Intuitively, \mathfrak{h} -regular sets of inductive rules generate heaps with a regular structure (in the sense that it may be represented by a tree automaton [9]), enriched with some additional edges (referring to the nodes corresponding to the variables passed as parameters to the spatial predicates at some recursive calls). These additional edges may refer to locations corresponding to free variables (e.g. the root of the structure) but also to existential variables (for instance they may refer to the parent node in the tree). \mathfrak{h} -Regular SID are related to the PCE systems introduced in [13] (for **p**rogressing, **c**onnected and **e**stablished), extended to formulas with permissions, but our conditions are slightly stronger, because we require that every existential variable be allocated at the next recursive call. Note that structures with mixed permissions are allowed, for instance

³ i.e., compound permission terms are not allowed in predicate atoms.

⁴ As \mathfrak{h} -regular rules allocate exactly one location, we assume that the segment is non empty, the case of an empty segment must be considered apart.

the rules $P(x, z_1, z_2) \Leftarrow x \overset{z_1}{\mapsto} ()$ and $P(x, z_1, z_2) \Leftarrow \exists u (x \overset{z_1}{\mapsto} (u) \circ P(u, z_2, z_1))$ defines a list with permissions alternating between z_1 and z_2 . Rules with compound permission terms in points-to or permission atoms are allowed (such as $P(x, y_1, y_2) \Leftarrow x \overset{y_1 \oplus y_2}{\mapsto} () \circ \text{def}(y_1 \oplus y_1)$), but not those with compound permission terms in spatial predicate atoms⁵ (e.g., $P(x, y_1, y_2) \Leftarrow x \overset{y_1}{\mapsto} () \circ Q(x, y_1 \oplus y_2)$ is *not* \mathfrak{h} -regular).

For every quantifier-free formula ϕ , we denote by $\text{roots}(\phi)$ the set of variables x (called the *roots of ϕ*) inductively defined as follows: $\text{roots}(x \overset{P}{\mapsto} (y_1, \dots, y_k)) \stackrel{\text{def}}{=} \{x\}$, $\text{roots}(P(x, y_1, \dots, y_k)) \stackrel{\text{def}}{=} \{x\}$, $\text{roots}(\exists y \phi) = \text{roots}(\phi) \setminus \{y\}$, $\text{roots}(\phi) = \emptyset$ if ϕ is pure and $\text{roots}(\phi_1 * \phi_2) = \text{roots}(\phi_1 \circ \phi_2) = \text{roots}(\phi_1) \cup \text{roots}(\phi_2)$. By Definition 7, roots are always allocated:

Proposition 8. *Let \mathcal{R} be a \mathfrak{h} -regular SID. If $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$ and $x \in \text{roots}(\phi)$ then $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$. Consequently, every formula of the form $\phi_1 \circ \phi_2$ with $\text{roots}(\phi_1) \cap \text{roots}(\phi_2) \neq \emptyset$ is $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable.*

The conditions in Definition 7 are actually not sufficient to ensure that the satisfiability problem is decidable:

Theorem 9. *If there exist (not necessary distinct) permissions $\pi_1, \pi_2 \in \mathcal{P}_{\mathfrak{P}}$ such that $\pi_1 \oplus_{\mathfrak{P}} \pi_2$ is defined, then the $(\mathcal{R}, \mathfrak{P})$ -satisfiability problem is undecidable for \mathfrak{h} -regular SID \mathcal{R} .*

To ensure decidability, we need to further restrict the way existential variables are passed as parameters during recursive calls. This is the goal of the next definition.

Definition 10. *Assume that \mathcal{R} is \mathfrak{h} -regular. Given two spatial predicates P and Q , of arities n and m respectively, we write $P \bowtie_{\mathcal{R}} Q$ if $P(x, x_1, \dots, x_{n-1}) * Q(x, y_1, \dots, y_{m-1})$ is $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable⁶ (where $x_1, \dots, x_{n-1}, y_1, \dots, y_{m-1}$ denote pairwise distinct variables of the appropriate sorts). We denote by $\gamma_{\mathcal{R}}$ the function associating every predicate symbol P of spatial arity n to a subset of $\{2, \dots, n\}$ inductively defined as follows: for every rule $P(x_1, \dots, x_n, \mathbf{u}) \Leftarrow \exists y_1, \dots, y_m \phi$ in \mathcal{R} , for every predicate atom $Q(z_1, \dots, z_k, \mathbf{u}_k)$ in ϕ with $\#_1(Q) = k$ and for all $i \in \{2, \dots, k\}$:*

1. $z_i \in \{y_1, \dots, y_m\} \Rightarrow i \in \gamma_{\mathcal{R}}(Q)$.
2. $z_i \in \{x_j \mid j \in \gamma_{\mathcal{R}}(P)\} \Rightarrow i \in \gamma_{\mathcal{R}}(Q)$.

⁵ Otherwise the unfolding of spatial predicates could yield terms of arbitrary depth.

⁶ In practice, as this condition is hard to test, some stronger syntactic condition can be tested instead, for instance one can check that all the formulas ϕ and ϕ' such that $P(x, x_1, \dots, x_{n-1}) \Leftarrow_{\mathcal{R}} \phi$ and $Q(x, y_1, \dots, y_{m-1}) \Leftarrow_{\mathcal{R}} \phi'$ are of the form $\phi = (x \mapsto (\mathbf{u}) \circ \psi)$ and $\phi' = (x \mapsto (\mathbf{u}') \circ \psi')$ with $|\mathbf{u}| \neq |\mathbf{u}'|$ (this condition is used in Theorem 33 and for the EXPTIME-hardness proof in Theorem 32.). More generally, it is sufficient to test that the “shape” of the structures generated by P and Q , up to a certain fixed unfolding depth, are incompatible.

Let \mathcal{P}^* be a subset of \mathcal{P} , such that: (3) $P \in \mathcal{P}^* \implies \gamma_{\mathcal{R}}(P) = \emptyset$; and (4) $P \in \mathcal{P}^* \wedge Q \in \mathcal{P} \setminus \mathcal{P}^* \implies P \bowtie_{\mathcal{R}} Q$. A \mathfrak{h} -regular rule is \exists -restricted (w.r.t. \mathcal{R} and \mathcal{P}^*) if it satisfies the following condition (using the notations of Definition 7):

$$5. \forall i \in \{1, \dots, n\} \forall j \in \{1, \dots, n\} (u_i \in \mathbf{y}_j \implies Q_i \in \mathcal{P}^*).$$

A SID \mathcal{R} is \exists -restricted if all the rules in \mathcal{R} are \exists -restricted.

Conditions 1 and 2 in Definition 10 are meant to ensure that $\gamma_{\mathcal{R}}(P)$ denotes the indices of the parameters of P that may (but do not have to) be instantiated by some existential variable introduced during the unfolding of the inductive rules in \mathcal{R} (the other parameters may only be instantiated by variables occurring in the initial formula). Condition 1 corresponds to a base case, where an existential variable is passed as a parameter to a predicate symbol, and Condition 2 handles the inductive case, when the variable is carried through recursive calls⁷. Then, Condition 5 ensures that an existential variable may only be passed as a parameter to a predicate symbol if it is the root of a structure defined by an atom $Q_i(\mathbf{y}_i)$ containing no variables introduced by unfolding (by Condition 3).

Example 11. The rules of the predicate `lseg` are \exists -restricted (with $\mathcal{P}^* = \emptyset$). Indeed, they contain only one existential variable u , which occurs only as the first argument of a predicate. Hence Condition 5 in Definition 10 trivially holds. If \mathcal{R} contains no other rule then $\gamma_{\mathcal{R}}(\text{lseg}) = \emptyset$. Note that $\gamma_{\mathcal{R}}(\text{lseg})$ depends on the entire set \mathcal{R} . For instance, if \mathcal{R} contains a rule $P(x, y) \Leftarrow \exists u (x \xrightarrow{y} (u) \circ \text{lseg}(u, u, y))$ then the second argument of `lseg` may be instantiated by an existential variable hence $\gamma_{\mathcal{R}}(\text{lseg}) = \{2\}$, and the latter rule is not \exists -restricted. On the other hand, if $\mathcal{P}^* = \{Q\}$, then the rules $Q(x, y) \Leftarrow x \xrightarrow{y} ()$, $R(x, y) \Leftarrow \exists u, v (x \xrightarrow{y} (u, v) \circ \text{lseg}(u, v, y) \circ Q(v, y))$ are \exists -restricted, with $\mathcal{P}^* = \{Q\}$. Indeed, the variable u occurs only at the root of a predicate, and the variable v is the root of $Q(v, y)$. Note that $\text{lseg}(x, y, z) * Q(x, u)$ and $R(x, y) * Q(x, u)$ are $(\mathcal{R}, \mathfrak{B})$ -unsatisfiable, thus $\text{lseg} \bowtie_{\mathcal{R}} Q$ and $R \bowtie_{\mathcal{R}} Q$.

Intuitively, the structures generated by \exists -restricted rules are regular tree-shaped structures, enriched with two kinds of additional edges: (i) a *bounded* number of *arbitrary* edges (corresponding to free variables, which may be freely passed as arguments to any predicate, thus may be referred to in an arbitrary way); (ii) an *unbounded* number of other edges (corresponding to existential variables) which are only allowed to point to structures that contain no edge of type (ii). Condition 4 ensures that the structures containing only edges of type (i) do not overlap with those containing both kinds of edges. Note that the conditions of Definition 10 always hold if the existential variables occur only as roots (with $\mathcal{P}^* = \mathcal{P}$ or $\mathcal{P}^* = \emptyset$). In this case there is no edge of type (ii), i.e., the obtained structures are regular sets of trees with a bounded number of additional edges (for instance

⁷ For generality, one could assume that all the equalities occurring in the rules are propagated before $\gamma_{\mathcal{R}}$ is computed (so that existential variables are eliminated if they are equal to a free variable), but this is not essential for our purposes hence the corresponding formal definitions are omitted.

trees with pointers to the root, or cyclic lists). Note that doubly linked lists cannot be captured (as they contain an unbounded number of additional edges from every node to the previous one). In the following we devise an algorithm to test the $(\mathcal{R}, \mathfrak{F})$ -satisfiability of symbolic heaps when \mathcal{R} is \exists -restricted.

4 A Decision Procedure For Testing Satisfiability

Before entering into technical details we start with a general overview of the procedure for testing satisfiability (assuming the considered SID is \exists -restricted).

1. Starting with a formula of the form $\delta_1 * \dots * \delta_n$ where the δ_i 's are atoms, we first reduce every spatial atom δ_i into an equivalent disjunction of \circ -conjunctions $\delta_1^i \circ \dots \circ \delta_{m_i}^i$ such that the *only* free variables allocated by an atom δ_j^i are its roots $roots(\delta_j^i)$ (as δ_j^i is an atom, $card(roots(\delta_j^i)) \leq 1$). Due to the particular properties of the \mathfrak{h} -regular rules (more precisely, due to the fact that the rules satisfy the “establishment” property of [13], i.e., every existential variable is allocated), this entails that, for all structures $(\mathfrak{s}, \mathfrak{h}_{i,j})$ satisfying δ_j^i , the domains of $\mathfrak{h}_{i,j}$ and $\mathfrak{h}_{i',j'}$ are either equal (if δ_j^i and $\delta_{j'}^{i'}$ have the same roots) or disjoint (otherwise). Indeed, the establishment property ensures that the considered heaps have no “pending edges” (i.e., no location that is referred to but not allocated), other than those denoted by free variables. This step can be considered as the key part of the procedure. It requires to (automatically) enrich the language with additional predicates and rules, and the termination of the transformation crucially depends on the conditions on \exists -restricted rules. For instance, an atom $\mathbf{lseg}(x, x)$ occurring in a formula with free variables x, y could be written $(x \simeq y \circ \mathbf{lseg}(x, x)) \vee \mathbf{lseg}'(x, x, y) \vee (\mathbf{lseg}'(x, y, y) \circ \mathbf{lseg}'(y, x, x))$ where $\mathbf{lseg}'(u, v, w)$ denotes a list segment from u to v not allocating w . The previous decomposition depends on whether y is equal to x and whether y occurs in the list segment from x to x .
2. By distributivity, we get at this point $*$ -conjunctions of \circ -conjunctions of atoms. Taking advantage of the previous property, we then reduce these formulas into \circ -conjunctions of $*$ -conjunctions of atoms, by regrouping the atoms with the same roots, e.g., $(P(x, y) \circ Q(y, x)) * (P'(x, y) \circ Q'(y, x))$ may be written $(P(x, y) * P'(x, y)) \circ (Q(y, x) * Q'(y, x))$.
3. Next, we show that a $*$ -conjunction of atoms sharing the same root (such as $P(x, y) * P'(x, y)$ or $Q(y, x) * Q'(y, x)$) can be denoted by a single atom, the rules of which are obtained by “merging” the rules of the initial atoms.
4. At this point we get a \circ -conjunction of atoms. To ensure that the formula is satisfiable it suffices to test that all these atoms have a model and that all these models are compatible, w.r.t. the equality constraints, allocated locations and permission constraints. To this aim, we construct finite abstractions of the models of the considered atoms using a bottom-up fixpoint algorithm.

In the next subsections, each of these steps is explained in details.

4.1 Normalization

We first show that every formula can be transformed into an equivalent formula (that we call *normalized*) in which every allocated variable occurs as a root:

Definition 12. A formula ϕ is normalized if it is of the form $\exists \mathbf{x} \psi$ where ψ is quantifier-free and for all spatial atoms δ in ψ , for all $(\mathcal{R}, \mathfrak{P})$ -models $(\mathfrak{s}, \mathfrak{h})$ of δ and for all variables $y \in \text{fv}(\psi)$: $\mathfrak{s}(y) \in \text{dom}(\mathfrak{h}) \iff y \in \text{roots}(\psi)$.

For instance, $\text{lseg}(x, y)$ is not normalized, because y may be allocated (e.g., if $\mathfrak{s}(x) = \mathfrak{s}(y)$) and does not occur in $\text{roots}(\text{lseg}(x, y)) = \{x\}$. To enforce this condition, we introduce new predicate symbols (called *derived predicates*), the rules of which can be automatically computed from those of the predicates already occurring in this formula. We first define predicate symbols that ensure that some given variable is not allocated.

Definition 13. For all predicate atoms $P(\mathbf{x}, \mathbf{p})$ (where \mathbf{x} and \mathbf{p} are vectors of location variables and permission terms, respectively) and for all location variables v , we denote by $P(\mathbf{x}, \mathbf{p})[v]^-$ any atom of the form $Q(\mathbf{x}, v, \mathbf{p})$, where Q is a fresh predicate symbol, associated with the rules:

$$Q(\mathbf{y}, w, \mathbf{z}) \Leftarrow \exists \mathbf{u} (Q_1(\mathbf{y}_1, \mathbf{p}_1)[w]^- \circ \dots \circ Q_m(\mathbf{y}_m, \mathbf{p}_m)[w]^- \circ \phi \circ \mathbf{y}|_1 \neq w)$$

for all rules $P(\mathbf{y}, \mathbf{z}) \Leftarrow \exists \mathbf{u} (Q_1(\mathbf{y}_1, \mathbf{p}_1) \circ \dots \circ Q_m(\mathbf{y}_m, \mathbf{p}_m) \circ \phi)$ in \mathcal{R} (up to AC), where \mathbf{y}, \mathbf{y}_i are vectors of location variables, \mathbf{z}, \mathbf{p}_i are vectors of permission variables, and ϕ contains no predicate atom.

For instance $\text{lseg}(x, y, z)[u]^-$ is a predicate atom $Q(x, y, u, z)$ defined by the following rules: $\{Q(x, y, u, z) \Leftarrow \exists x'(x \xrightarrow{z} (x') \circ Q(x', y, u, z) \circ x \neq u), Q(x, y, u, z) \Leftarrow x \xrightarrow{z} (y) \circ x \neq u\}$. It denotes a list segment from x to y not allocating u . The following result is straightforward to prove:

Proposition 14. For every \exists -restricted SID \mathcal{R} , the set \mathcal{R} enriched with the rules associated with the predicate Q corresponding to $P(\mathbf{x}, \mathbf{p})[v]^-$ in Definition 13 is \exists -restricted, with $\gamma_{\mathcal{R}}(Q) = \gamma_{\mathcal{R}}(P)$ and $Q \in \mathcal{P}^* \iff P \in \mathcal{P}^*$.

Intuitively the structures that satisfy $P(\mathbf{x}, \mathbf{p})[v]^-$ are exactly those that satisfy $P(\mathbf{x}, \mathbf{p})$ and do not allocate v :

Lemma 15. For all \mathfrak{h} -regular SID \mathcal{R} , $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} P(\mathbf{x}, \mathbf{p})[v]^-$ iff $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} P(\mathbf{x}, \mathbf{p})$ and $\mathfrak{s}(v) \notin \text{dom}(\mathfrak{h})$.

The operator $\delta \mapsto \delta[x]^-$ can be applied recursively, e.g., one can consider atoms of the form $\delta[x]^- [y]^-$, etc. For all predicate atoms δ , we denote by $\text{unalloc}(\delta)$ the set of variables inductively defined as follows: $\text{unalloc}(\delta[x]^-) \stackrel{\text{def}}{=} \{x\} \cup \text{unalloc}(\delta)$, and $\text{unalloc}(\delta) \stackrel{\text{def}}{=} \emptyset$ if δ is not of the form $\delta'[x]^-$. The following proposition is an immediate consequence of Lemma 15:

Proposition 16. If $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \delta$ then $\mathfrak{s}(x) \notin \text{dom}(\mathfrak{h})$, for all $x \in \text{unalloc}(\delta)$.

Next, we define predicate symbols allowing one to remove some part of a structure. Intuitively, the expression $(\phi \dashv\bullet \psi)$ will hold exactly in the structures that satisfy ψ when a disjoint structure satisfying ϕ is added. For instance given the rules $\mathbf{tree}(x, y) \Leftarrow \exists x_1, x_2 x \overset{y}{\dashv} (x_1, x_2) \circ \mathbf{tree}(x_1, y) \circ \mathbf{tree}(x_2, y)$ and $\mathbf{tree}(x, y) \Leftarrow x \overset{y}{\dashv} ()$, $\mathbf{tree}(z, y)$ and $\mathbf{tree}(x, y)$ denote binary trees with roots z and x , respectively, and $\mathbf{tree}(z, y) \dashv\bullet \mathbf{tree}(x, y)$ denotes a tree of root x with a “hole” at z (the structures satisfying $\mathbf{tree}(z, y) \dashv\bullet \mathbf{tree}(x, y)$ are obtained from models of $\mathbf{tree}(x, y)$ by removing the part of the heap that corresponds to $\mathbf{tree}(z, y)$). The formula $\phi \dashv\bullet \psi$ is similar to the *strong magic wand* introduced in [17] and to the *context predicates* in [12] and also close in spirit to the separating implication of SL although the semantics are slightly different.

Definition 17. For all finite sequences of predicate atoms $P_i(\mathbf{x}_i, \mathbf{p}_i)$ (with $i = 0, \dots, n$), where \mathbf{x}_i and \mathbf{p}_i are vectors of location variables and permission terms, respectively, we denote by $(P_1(\mathbf{x}_1, \mathbf{p}_1) \circ \dots \circ P_n(\mathbf{x}_n, \mathbf{p}_n)) \dashv\bullet P_0(\mathbf{x}_0, \mathbf{p}_0)$ any atom $P(\mathbf{x}, \mathbf{p})$ with $\mathbf{x} = \mathbf{x}_0 \dots \mathbf{x}_n$, $\mathbf{p} = \mathbf{p}_0 \dots \mathbf{p}_n$, and such that $P = P_0$ if $n = 0$ and otherwise P is a fresh symbol associated with rules of the form

$$P(\mathbf{y}, \mathbf{z}) \Leftarrow \exists \mathbf{w} (\psi_1 \circ \dots \circ \psi_m \circ \phi)$$

for all rules

$$P_0(\mathbf{y}_0, \mathbf{z}_0) \Leftarrow \exists \mathbf{w} (Q_1(\mathbf{u}_1, \mathbf{q}_1) \circ \dots \circ Q_m(\mathbf{u}_m, \mathbf{q}_m) \circ \phi)$$

in \mathcal{R} and for all decompositions $\alpha_1 \circ \dots \circ \alpha_m = P_1(\mathbf{y}_1, \mathbf{z}_1) \circ \dots \circ P_n(\mathbf{y}_n, \mathbf{z}_n)$ (up to AC, where the α_i 's may be empty), where:

- \mathbf{y}_i and \mathbf{z}_i are sequences of pairwise distinct location and permission variables, respectively, with $|\mathbf{y}_i| = |\mathbf{x}_i|$ and $|\mathbf{z}_i| = |\mathbf{p}_i|$;
- $\mathbf{y} = \mathbf{y}_0 \dots \mathbf{y}_n$, $\mathbf{z} = \mathbf{z}_1 \dots \mathbf{z}_n$;
- ψ_i is of one of the following forms:
 - either $\alpha_i \dashv\bullet Q_i(\mathbf{u}_i, \mathbf{q}_i)$;
 - or $\mathbf{y}_j \simeq \mathbf{u}_i \circ \mathbf{z}_j \simeq \mathbf{q}_i$, if $\alpha_i = P_j(\mathbf{y}_j, \mathbf{z}_j)$ and $P_j = Q_i$.

For instance $\mathbf{tree}(z, y) \dashv\bullet \mathbf{tree}(x, y)$ denotes an atom $P(x, z, y, y)$ with the rules:

$$\begin{aligned} P(x, z, y_1, y_2) &\Leftarrow \exists x_1, x_2 (x \overset{y_1}{\dashv} (x_1, x_2) \circ P(x_1, z, y_1, y_2) \circ \mathbf{tree}(x_2, z, y_1)) \\ P(x, z, y_1, y_2) &\Leftarrow \exists x_1, x_2 (x \overset{y_1}{\dashv} (x_1, x_2) \circ \mathbf{tree}(x_1, z, y_1) \circ P(x_2, z, y_1, y_2)) \\ P(x, z, y_1, y_2) &\Leftarrow \exists x_1, x_2 (x \overset{y_1}{\dashv} (x_1, x_2) \circ x_1 \simeq z \circ y_1 \simeq y_2 \circ \mathbf{tree}(x_2, z, y_1)) \\ P(x, z, y_1, y_2) &\Leftarrow \exists x_1, x_2 (x \overset{y_1}{\dashv} (x_1, x_2) \circ \mathbf{tree}(x_1, z, y_1) \circ x_2 \simeq z \circ y_1 \simeq y_2) \end{aligned}$$

For readability, all the expressions of the form $\mathbf{emp} \dashv\bullet \mathbf{tree}(x_2, z, y_1)$ have been replaced by $\mathbf{tree}(x_2, z, y_1)$. Note that the rules are not \mathfrak{h} -regular, as x_1 and x_2 do not occur as roots in every rule, but they can easily be transformed into \mathfrak{h} -regular rules by replacing x_1 and x_2 by z in the third and fourth rule, respectively (using the equations $x_1 \simeq z$ and $x_2 \simeq z$). The definition can be applied recursively (i.e., P_0, \dots, P_n may be derived predicates). The next proposition is an immediate consequence of Definition 17:

Proposition 18. *Let \mathcal{R} be a \mathfrak{h} -regular SID. The rules associated with any predicate P corresponding to an expression $\alpha \dashv\bullet \delta$ (Definition 17) are \mathfrak{h} -regular, up to the following equivalence: $\exists x (x \simeq y \circ \phi) \equiv_{\mathcal{R}}^{\exists} \phi\{x \leftarrow y\}$. Moreover, the rules are also \exists -restricted, with $\gamma_{\mathcal{R}}(P) = \gamma_{\mathcal{R}}(P_0)$ and $P \in \mathcal{P}^* \iff P_0 \in \mathcal{P}^*$. Finally if $\alpha = \text{emp}$ then $(\alpha \dashv\bullet \delta) = \delta$.*

Note that, however, the implication $P \in \mathcal{P}^* \wedge Q \in \mathcal{P} \setminus \mathcal{P}^* \implies P \bowtie_{\mathcal{R}} Q$ (Condition 4 in Definition 10) does *not* necessarily hold for derived predicates P, Q . The following lemma states a form of modus ponens, relating the connective \circ with $\dashv\bullet$:

Lemma 19. *If \mathcal{R} is \mathfrak{h} -regular then $P(\mathbf{x}, \mathbf{p}) \circ ((P(\mathbf{x}, \mathbf{p}) \circ \alpha) \dashv\bullet Q(\mathbf{y}, \mathbf{q})) \models_{\mathcal{R}}^{\exists} \alpha \dashv\bullet Q(\mathbf{y}, \mathbf{q})$.*

The next lemma states that every predicate atom allocating x can be written as a \circ -formula in which x occurs as a root.

Lemma 20. *Assume that \mathcal{R} is \exists -restricted. Let \mathbf{y}, \mathbf{p} be vectors of location variables and permission terms, respectively. If $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} Q(\mathbf{y}, \mathbf{p})$, $\mathfrak{s}(x) \neq \mathfrak{s}(\mathbf{y}|_1)$ and $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$, then there exist atoms of the form $P(x, \mathbf{z}, \mathbf{q})$, $P_i(x_i, \mathbf{y}_i, \mathbf{q}_i)$ (with $i \in \{1, \dots, n\}$), where $\mathbf{z} \subseteq \mathbf{y} \cup \{x_1, \dots, x_n\}$, $\mathbf{y}_i \subseteq \{\mathbf{y}|_j \mid j \notin \gamma_{\mathcal{R}}(Q)\}$, $\mathbf{q} \subseteq \mathbf{p}$ and $\mathbf{q}_i \subseteq \mathbf{p}$, such that: $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \exists x_1, \dots, x_n (\beta \circ (\beta \dashv\bullet Q(\mathbf{y}, \mathbf{p})))$, with $\beta = P(x, \mathbf{z}, \mathbf{q}) \circ \bigcirc_{i=1}^m P_i(x_i, \mathbf{y}_i, \mathbf{q}_i)$. Moreover, $P_i \in \mathcal{P}^*$, $\{x_1, \dots, x_n\} \subseteq (x, \mathbf{z})|_{\gamma_{\mathcal{R}}(P)}$ and $y \in \mathbf{y} \cap \mathbf{z} \wedge y \notin \{\mathbf{y}|_j \mid j \notin \gamma_{\mathcal{R}}(Q)\} \implies y \in (x, \mathbf{z})|_{\gamma_{\mathcal{R}}(P)}$.*

Intuitively, since x is allocated and the rules are \mathfrak{h} -regular, then necessarily some predicate atom of the form $P(x, \mathbf{z}, \mathbf{q})$ must be called at some point during the unfolding of the rules. Using $\dashv\bullet$, this predicate can be removed from the call tree of $Q(\mathbf{y}, \mathbf{p})$ and lifted at the root level in the formula. The atom $P(x, \mathbf{z}, \mathbf{q})$ may contain variables not occurring in $Q(\mathbf{y}, \mathbf{p})$ corresponding to existential variables introduced by unfolding. As the rules are \exists -restricted, all such variables x_i must themselves appear as the root of some predicate atom $P_i(x_i, \mathbf{y}_i, \mathbf{q}_i)$ which contains (beside x_i) only variables occurring in $Q(\mathbf{y}, \mathbf{p})$ (since $\gamma_{\mathcal{R}}(P_i) = \emptyset$, due to Condition 5 in Definition 10). Again, these atoms can be moved at the root level. See Appendix E for details.

Definition 21. *For all atoms $Q(\mathbf{y}, \mathbf{p})$ we denote by $\delta[x]^+$ the set of formulas of the form $\exists x_1, \dots, x_n (\beta \circ (\beta \dashv\bullet Q(\mathbf{y}, \mathbf{p})))$ as defined in Lemma 20. We also denote by $\delta[x]^=$ the formula: $\delta \circ (x \simeq \mathbf{y}|_1)$.*

For every model of δ , $\delta[x]^-$ holds if x is not allocated in δ , $\delta[x]^=$ holds if x is equal to the root of δ and $\delta[x]^+$ holds if x is allocated but is not the root of δ . The following result follows immediately from Lemmata 19 and 20:

Lemma 22. *Assume that \mathcal{R} is \exists -restricted. Let $x \in \mathcal{V}_1$. For every predicate atom δ such that $x \notin \text{roots}(\delta)$, and for all structures $(\mathfrak{s}, \mathfrak{h})$: $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \delta$ iff there exists $\psi \in \{\delta[x]^-, \delta[x]^=\} \cup \delta[x]^+$ such that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \psi$.*

For instance the atom $\mathbf{lseg}(x, y, z)$ holds iff one of the formulas $\mathbf{lseg}(x, y, z) \circ x \simeq y$, $\mathbf{lseg}(x, y, z)[y]^-$ or $\mathbf{lseg}(y, y, z) \circ (\mathbf{lseg}(y, y, z) \dashv\bullet \mathbf{lseg}(x, y, z))$ holds. The second formula corresponds to the case where y is not allocated, and the first and third ones correspond to the case where there is a loop on y . By applying repeatedly Lemma 22 on every variable x and atom δ we eventually obtain a disjunction of normalized formulas:

Lemma 23. *Let \mathcal{R} be a \exists -restricted SID. There exists an algorithm transforming any symbolic heap ϕ containing no points-to atom into a set of normalized formulas Ψ such that for all structures $(\mathfrak{s}, \mathfrak{h})$: $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$ iff there exists $\psi \in \Psi$ such that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \psi$. Furthermore, every formula in Ψ is a (quantified) separating conjunction of \circ -formulas.*

4.2 Commuting Separating and Disjoint Connections

The next step consists in showing that – under some particular conditions enforced by the previous transformation – the operator $*$ can be pushed innermost in the formula (below the operator \circ). To this aim, we exploit an essential property of \mathfrak{h} -regular SIDs, namely that all the locations that occur in the heap of some model of a formula ϕ but are not allocated correspond to a variable in $fv(\phi)$. We shall denote by $cut(L, L', \mathfrak{h})$ the set of locations reachable from L in \mathfrak{h} , from a path not crossing L' :

Definition 24. *Let \mathfrak{h} be a heap, let $L, L' \subseteq \mathcal{L}$. We denote by $cut(L, L', \mathfrak{h})$ the set of locations inductively defined as follows: $L \subseteq cut(L, L', \mathfrak{h})$, and if $\ell' \in cut(L, L', \mathfrak{h})$, $\mathfrak{h}(\ell') = (\ell_1, \dots, \ell_k, \pi)$, $i \in \{1, \dots, k\}$ and $\ell_i \notin L'$ then $\ell_i \in cut(L, L', \mathfrak{h})$.*

The following lemma characterizes the domain of the part of the heap satisfying some formula ϕ :

Lemma 25. *Let \mathcal{R} be a \mathfrak{h} -regular SID and let ϕ be a \circ -formula containing no quantifier. Let \mathfrak{s} be a store and let $\mathfrak{h}, \mathfrak{h}'$ be heaps, with $\mathfrak{h}' \leq \mathfrak{h}$. Let V be a set of variables, with $fv(\phi) \subseteq V \cup roots(\phi)$ and $\mathfrak{s}(V) \cap dom(\mathfrak{h}') = \emptyset$. If $(\mathfrak{s}, \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} \phi$ then $dom(\mathfrak{h}') = cut(\mathfrak{s}(roots(\phi)), \mathfrak{s}(V), \mathfrak{h})$.*

The commutation property, pushing $*$ below \circ , is given by Lemma 26:

Lemma 26. *Let \mathcal{R} be a \mathfrak{h} -regular SID. Let $V \subseteq \mathcal{V}_1$ and let ϕ be a normalized formula, of the form $\phi = \phi' \circ (\mathfrak{*}_{i=1}^n (\phi_i \circ \psi_i) * \psi')$, where, for all $i \in \{1, \dots, n\}$, $roots(\phi_i) = V$ and $(roots(\psi_i) \cup roots(\psi')) \cap V = \emptyset$. Then ϕ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable iff $(\phi' \circ \mathfrak{*}_{i=1}^n \phi_i) \circ ((\mathfrak{*}_{i=1}^n \psi_i) * \psi')$ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable.*

Roughly speaking, as $roots(\phi_i) = V$ and ϕ_i is normalized, it is possible to prove, using the characterization given in Lemma 25, that the parts of the heap that correspond to the formulas ϕ_i have all the same domain. This entails that the heaps corresponding to the formulas ψ_i and ψ' are disjoint, which permits to prove that $\mathfrak{*}_{i=1}^n (\phi_i \circ \psi_i)$ can be written $(\mathfrak{*}_{i=1}^n \phi_i) \circ (\mathfrak{*}_{i=1}^n \psi_i)$, yielding the result. The detailed proof is given in Appendix H.

4.3 Merging of Spatial Predicates

We show that, under some particular conditions, it is possible to replace the separating conjunction of two spatial atoms having the same root by a single spatial atom. The rules defining this atom are obtained by combining the rules of the two initial atoms. More precisely, consider any \mathfrak{h} -regular SID \mathcal{R} and two spatial atoms $P(x, \mathbf{y}, \mathbf{p})$ and $P'(x, \mathbf{y}', \mathbf{p}')$ sharing the same root x , where \mathbf{y}, \mathbf{y}' are vectors of location variables and \mathbf{p} and \mathbf{p}' are vectors of permission terms. We denote by $P(x, \mathbf{y}, \mathbf{p}) \nabla P'(x, \mathbf{y}', \mathbf{p}')$ any atom $Q(x, \mathbf{y}, \mathbf{y}', \mathbf{p}, \mathbf{p}')$ where Q is associated with rules of the form:

$$Q(v, \mathbf{w}, \mathbf{w}', \mathbf{z}, \mathbf{z}') \Leftarrow \exists u_1, \dots, u_n \quad v \stackrel{q}{\mapsto} (v_1, \dots, v_k) \\ \circ \bigcirc_{i=1}^n (Q_i(u_i, \mathbf{y}_i, \mathbf{q}_i) \nabla Q'_i(u_i, \mathbf{y}'_i, \mathbf{q}'_i)) \circ \phi \circ \phi' \circ \psi$$

with $q \stackrel{\text{def}}{=} p \oplus p'$, for all pairs of rules of the following forms in \mathcal{R} (with the same numbers k and n , and up to α -renaming, so that the rules share the same existential variables):

$$P(v, \mathbf{w}, \mathbf{z}) \Leftarrow \exists u_1, \dots, u_n \quad v \stackrel{p}{\mapsto} (v_1, \dots, v_k) \circ \bigcirc_{i=1}^n Q_i(u_i, \mathbf{y}_i, \mathbf{q}_i) \circ \phi \\ P'(v, \mathbf{w}', \mathbf{z}') \Leftarrow \exists u_1, \dots, u_n \quad v \stackrel{p'}{\mapsto} (v'_1, \dots, v'_k) \circ \bigcirc_{i=1}^n Q'_i(u_i, \mathbf{y}'_i, \mathbf{q}'_i) \circ \phi'$$

where $\psi = \bigcirc_{i=1}^k (v_i \simeq v'_i)$. Note that all the produced rules are \mathfrak{h} -regular⁸.

Lemma 27. *Let \mathcal{R} be a \mathfrak{h} -regular SID. Let $x \in \mathcal{V}_1$ and let $(\mathfrak{s}, \mathfrak{h})$ be a structure such that $\mathfrak{s}(y) \notin \text{dom}(\mathfrak{h})$ holds for all variables y such that $\mathfrak{s}(x) \neq \mathfrak{s}(y)$. Then $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} P(x, \mathbf{y}, \mathbf{p}) \nabla P'(x, \mathbf{y}', \mathbf{p}') \iff (\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} P(x, \mathbf{y}, \mathbf{p}) * P'(x, \mathbf{y}', \mathbf{p}')$.*

The result crucially depends on the fact that the parts of the heap that correspond to $P(x, \mathbf{y}, \mathbf{p})$ and $P'(x, \mathbf{y}', \mathbf{p}')$ respectively must share the same domain, since otherwise, as \mathcal{R} is \mathfrak{h} -regular, a free variable would be allocated, contradicting the hypothesis. This ensures that the heap can be generated by the above rules (see Appendix I for details).

4.4 Heap Abstractions and Main Result

As we shall see later, the previous transformations can be used to transform any symbolic heap into a \circ -formula (while preserving satisfiability). The final step is to devise an algorithm to test the satisfiability of \circ -formulas. As it is done in [6] for standard heap models, the algorithm works by constructing relevant abstractions of the models of the predicate atoms. It suffices to keep track of the truth value of the equational atoms, of the allocated variables and of the permission atoms satisfied by the structure.

⁸ However \exists -restrictedness is not necessarily preserved.

Definition 28. A heap abstraction is a tuple $\mathbf{a} = (V_{\mathbf{a}}, \sim_{\mathbf{a}}, A_{\mathbf{a}}, \rho_{\mathbf{a}})$ where $V_{\mathbf{a}}$ is a finite set of variables, $\sim_{\mathbf{a}}$ is an equivalence relation on the variables of sort $\mathbf{1}$ occurring in $V_{\mathbf{a}}$, $A_{\mathbf{a}}$ is a subset of $V_{\mathbf{a}} \cap \mathcal{V}_1$, closed under $\sim_{\mathbf{a}}$ (i.e., for all $x, y \in \mathcal{V}_1$: $x \in A_{\mathbf{a}} \wedge x \sim_{\mathbf{a}} y \implies y \in A_{\mathbf{a}}$), and $\rho_{\mathbf{a}}$ is a permission formula (with variables in $V_{\mathbf{a}}$).

Definition 29. Let $(\mathfrak{s}, \mathfrak{h})$ be a structure and let $\mathbf{a} = (V_{\mathbf{a}}, \sim_{\mathbf{a}}, A_{\mathbf{a}}, \rho_{\mathbf{a}})$ be a heap abstraction. We write $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} \mathbf{a}$ if all the following conditions are satisfied: (i) For all variables $x, y \in V_{\mathbf{a}} \cap \mathcal{V}_1$: $x \sim_{\mathbf{a}} y \iff \mathfrak{s}(x) = \mathfrak{s}(y)$; (ii) for all $x \in V_{\mathbf{a}} \cap \mathcal{V}_1$, $x \in A_{\mathbf{a}} \iff \mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$; and (iii) $\mathfrak{s} \models^{\mathfrak{P}} \rho_{\mathbf{a}}$. A heap abstraction is \mathfrak{P} -satisfiable if there exists a structure $(\mathfrak{s}, \mathfrak{h})$ such that $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} \mathbf{a}$.

Proposition 30. A heap abstraction \mathbf{a} is \mathfrak{P} -satisfiable iff $\rho_{\mathbf{a}}$ is \mathfrak{P} -satisfiable.

For all \circ -formulas ϕ , we define a set of heap abstractions $\mathfrak{A}(\phi)$ by mutual induction as follows. The sets $\mathfrak{A}(\phi)$ are the least sets of heap abstractions satisfying the following properties, for all finite sets of variables⁹ $V \supseteq \text{fv}(\phi)$ and for all equivalence relations \sim on $V \cap \mathcal{V}_1$: (i) if $\phi = x \stackrel{P}{\mapsto} (y_1, \dots, y_n)$ then $(V, \sim, \{y \mid y \sim x\}, \text{def}(P)) \in \mathfrak{A}(\phi)$. (ii) if $\phi = x \simeq y$ (resp. $x \not\sim y$) with $x, y \in \mathcal{V}_1$ and $x \sim y$ (resp. $x \not\sim y$) then $(V, \sim, \emptyset, \mathbf{emp}) \in \mathfrak{A}(\phi)$; (iii) if ϕ is a permission formula then $(V, \sim, \emptyset, \phi) \in \mathfrak{A}(\phi)$; (iv) if $\phi = \exists x \psi$, $(V, \sim, A, \rho) \in \mathfrak{A}(\psi)$ then $(V \setminus \{x\}, \sim', A \setminus \{x\}, \rho) \in \mathfrak{A}(\phi)$, where \sim' denotes the restriction of \sim to the variables distinct from x , i.e., $\sim' \stackrel{\text{def}}{=} \{(u, v) \mid u \sim v \wedge u, v \neq x\}$ (note that x cannot occur in ρ , since quantification over permission variables is not allowed); (v) if $\phi = \phi_1 \circ \phi_2$, $(V, \sim, A_i, \rho_i) \in \mathfrak{A}(\phi_i)$ (for all $i = 1, 2$) with $A_1 \cap A_2 = \emptyset$, then $(V, \sim, A_1 \cup A_2, \rho_1 \circ \rho_2) \in \mathfrak{A}(\phi)$; (vi) if $\phi = P(\mathbf{x}, \mathbf{p})$ and $\phi \leftarrow_{\mathcal{R}} \xi$ then $\mathfrak{A}(\xi) \subseteq \mathfrak{A}(\phi)$.

Lemma 31. A \circ -formula ϕ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable iff at least one of the abstractions in $\mathfrak{A}(\phi)$ is \mathfrak{P} -satisfiable.

Putting things together we get the following result:

Theorem 32. If \mathfrak{P} -satisfiability is decidable for permission formulas, then there exists an algorithm that, for every \exists -restricted SID, decides whether a given formula ϕ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable. If, moreover, \mathfrak{P} -satisfiability is in EXPTIME, then $(\mathcal{R}, \mathfrak{P})$ -satisfiability is also in EXPTIME (for \exists -restricted SID). Finally, for every permission model \mathfrak{P} , $(\mathcal{R}, \mathfrak{P})$ -satisfiability is EXPTIME-hard (for \exists -restricted SID).

5 Using Separating Conjunctions Inside Rules

To end the paper, we wish to point out that the satisfiability problem is undecidable from \exists -restricted SID if the disjoint separation \circ is replaced by the standard

⁹ For technical convenience we do not impose any bound on the cardinality of V , hence the set $\mathfrak{A}(\phi)$ is infinite. This simplifies the theoretical definition of the abstraction for disjoint conjunctions. In practice only variables occurring in the initial formula or in the rules need to be considered.

separating connective $*$ in the inductive definitions (see Definition 7). We think that the result is of some theoretical interest, although, as explained above, rules using \circ are actually more convenient for describing data structures. The notions of $*\mathfrak{h}$ -regular and $*\exists$ -restricted SID are defined exactly as \mathfrak{h} -regular SID and \exists -restricted SID (Definitions 7 and 10) except that the symbol \circ is replaced by $*$ everywhere (for conciseness the formal definitions are omitted).

Theorem 33. *Let \mathfrak{P} be any permission model and assume that for every $n \in \mathbb{N}$, there exists $\pi \in \mathcal{P}_{\mathfrak{P}}$ such that π^n is defined. The $(\mathcal{R}, \mathfrak{P})$ -satisfiability problem is undecidable for $*\exists$ -restricted SID.*

6 Conclusion and Future Work

An algorithm was devised to test the satisfiability of symbolic heaps in Separation Logic with inductively defined predicates and permissions, under some (syntactic) conditions on the inductive rules giving the semantics of the spatial predicates. The algorithm runs in exponential time, provided the satisfiability of permission formulas is in EXPTIME. In addition, we showed that some natural relaxings of these conditions make the problem undecidable (under some minimal assumptions on the permission model). The next step is to investigate the entailment problem for the considered fragment. The techniques devised in the present paper for transforming symbolic heaps into disjoint conjunctions of atoms should serve as a basis for this purpose, but the extension is not straightforward. Another (much easier) extension that could be of practical relevance is to consider formulas with labels (in the sense of [5]) which allow one to express additional equality conditions on some parts of the structures. In our context, labels would simply yield additional conditions on the decomposition generated during the normalization step: two formulas sharing the same label should be decomposed into formulas with the same set of roots. It could also be interesting to relax some of the conditions on the rules, for instance to allow for existential variables not occurring as roots in the rules. This is required to encode data structures with forward pointers, such as skip lists. It is also unclear whether Condition 4 in Definition 10 is required for decidability. Finally, the decision algorithm could probably be extended to handle arbitrary combinations of disjoint and separating conjunctions.

Acknowledgments. This work has been partially funded by the the French National Research Agency (ANR-21-CE48-0011)

References

1. J. Berdine, C. Calcagno, and P. W. O’Hearn. Smallfoot: Modular automatic assertion checking with separation logic. In F. S. de Boer, M. M. Bonsangue, S. Graf, and W. P. de Roever, editors, *Formal Methods for Components and Objects, 4th International Symposium, FMCO 2005, Amsterdam, The Netherlands, November 1-4, 2005, Revised Lectures*, volume 4111 of *Lecture Notes in Computer Science*, pages 115–137. Springer, 2005.

2. J. Berdine, B. Cook, and S. Ishtiaq. Slayer: Memory safety for systems-level code. In G. G. and Shaz Qadeer, editor, *Computer Aided Verification - 23rd International Conference, CAV 2011, Snowbird, UT, USA, July 14-20, 2011. Proceedings*, volume 6806 of *LNCS*, pages 178–183. Springer, 2011.
3. R. Bornat, C. Calcagno, P. W. O’Hearn, and M. J. Parkinson. Permission accounting in separation logic. In J. Palsberg and M. Abadi, editors, *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, pages 259–270. ACM, 2005.
4. J. Boyland. Fractional permissions. In D. Clarke, J. Noble, and T. Wrigstad, editors, *Aliasing in Object-Oriented Programming. Types, Analysis and Verification*, volume 7850 of *Lecture Notes in Computer Science*, pages 270–288. Springer, 2013.
5. J. Brotherston, D. Costa, A. Hobor, and J. Wickerson. Reasoning over permissions regions in concurrent separation logic. In S. K. Lahiri and C. Wang, editors, *Computer Aided Verification - 32nd International Conference, CAV 2020, Los Angeles, CA, USA, July 21-24, 2020, Proceedings, Part II*, volume 12225 of *Lecture Notes in Computer Science*, pages 203–224. Springer.
6. J. Brotherston, C. Fuhs, J. A. N. Pérez, and N. Gorogiannis. A decision procedure for satisfiability in separation logic with inductive predicates. In T. A. Henzinger and D. Miller, editors, *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS ’14, Vienna, Austria, July 14 - 18, 2014*, pages 25:1–25:10. ACM, 2014.
7. C. Calcagno and D. Distefano. Infer: An automatic program verifier for memory safety of C programs. In M. G. Bobaru, K. Havelund, G. J. Holzmann, and R. Joshi, editors, *NASA Formal Methods - Third International Symposium, NFM 2011, Pasadena, CA, USA, April 18-20, 2011. Proceedings*, volume 6617 of *Lecture Notes in Computer Science*, pages 459–465. Springer, 2011.
8. 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, Wroclaw, Poland, Proceedings*, pages 366–378. IEEE Computer Society, 2007.
9. H. Comon, M. Dauchet, R. Gilleron, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. Tree automata techniques and applications. Available on: <http://www.grappa.univ-lille3.fr/tata>, 1997.
10. S. Demri, É. Lozes, and D. Lugiez. On symbolic heaps modulo permission theories. In S. V. Lokam and R. Ramanujam, editors, *37th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2017, December 11-15, 2017, Kanpur, India*, volume 93 of *LIPICs*, pages 25:1–25:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.
11. M. Echenim, R. Iosif, and N. Peltier. Entailment checking in separation logic with inductive definitions is 2-exptime hard. In E. Albert and L. Kovács, editors, *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.
12. M. Echenim, R. Iosif, and N. Peltier. Decidable entailments in separation logic with inductive definitions: Beyond establishment. In *CSL 2021: 29th International Conference on Computer Science Logic*, EPiC Series in Computing. EasyChair, 2021.
13. R. Iosif, A. Rogalewicz, and J. Simacek. The tree width of separation logic with recursive definitions. In *Proc. of CADE-24*, volume 7898 of *LNCS*, 2013.

14. S. S. Ishtiaq and P. W. O'Hearn. Bi as an assertion language for mutable data structures. In *ACM SIGPLAN Notices*, volume 36, pages 14–26, 2001.
15. J. Katelaan and F. Zuleger. Beyond symbolic heaps: Deciding separation logic with inductive definitions. In E. Albert and L. Kovács, editors, *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 390–408. EasyChair, 2020.
16. Q. L. Le. Compositional satisfiability solving in separation logic. In F. Henglein, S. Shoham, and Y. Vizek, editors, *Verification, Model Checking, and Abstract Interpretation - 22nd International Conference, VMCAI 2021, Copenhagen, Denmark, January 17-19, 2021, Proceedings*, volume 12597 of *Lecture Notes in Computer Science*, pages 578–602. Springer, 2021.
17. K. Nakazawa, M. Tatsuta, D. Kimura, and M. Yamamura. Cyclic Theorem Prover for Separation Logic by Magic Wand. In *ADSL 18 (First Workshop on Automated Deduction for Separation Logics)*, July 2018. Oxford, United Kingdom.
18. P. W. O'Hearn and D. J. Pym. The logic of bunched implications. *Bull. Symb. Log.*, 5(2):215–244, 1999.
19. J. A. N. Pérez and A. Rybalchenko. Separation logic modulo theories. In C. Shan, editor, *Programming Languages and Systems - 11th Asian Symposium, APLAS 2013, Melbourne, VIC, Australia, December 9-11, 2013. Proceedings*, volume 8301 of *LNCS*, pages 90–106. Springer, 2013.
20. R. Piskac, T. Wies, and D. Zufferey. Automating separation logic using SMT. In N. Sharygina and H. Veith, editors, *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, volume 8044 of *LNCS*, pages 773–789. Springer, 2013.
21. X. Qiu, P. Garg, A. Stefanescu, and P. Madhusudan. Natural proofs for structure, data, and separation. In H. Boehm and C. Flanagan, editors, *ACM SIGPLAN PLDI '13*, pages 231–242. ACM, 2013.
22. J. Reynolds. Separation Logic: A Logic for Shared Mutable Data Structures. In *Proc. of LICS'02*, 2002.
23. Z. Xu, T. Chen, and Z. Wu. Satisfiability of compositional separation logic with tree predicates and data constraints. In L. de Moura, editor, *CADE 26*, volume 10395 of *LNCS*, pages 509–527. Springer, 2017.

A Proof of Proposition 8

The proof is by induction on the satisfiability relation. ϕ cannot be purely-spatial as $\text{roots}(\phi)$ would be empty, contradicting the hypothesis of the proposition. If ϕ is a points-to atom then by definition of $\text{roots}(\phi)$ it must be of the form $\phi = x \overset{P}{\mapsto} (y_1, \dots, y_k)$, so that $\text{dom}(\mathfrak{h}) = \{\mathfrak{s}(x)\}$ by definition of the semantics of points-to atoms. If ϕ is a spatial predicate atom then by definition it must be of the form $P(x, y_1, \dots, y_k)$, and by definition of the semantics of predicate atoms we have $\phi \leftarrow_{\mathcal{R}} \psi$ with $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \psi$. Since \mathcal{R} is \mathfrak{h} -regular, ψ is of the form $\exists \mathbf{u}(x \overset{P}{\mapsto} (y'_1, \dots, y'_k) \circ \psi')$, with $x \notin \mathbf{u}$, so that $x \in \text{roots}(\psi)$, and we get $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$ by the induction hypothesis. If $\phi = \exists y \psi$ then by definition $x \in \text{roots}(\psi)$ and $x \neq y$, moreover, there exists a location ℓ such that $\mathfrak{s}\{y \leftarrow \ell\} \models_{\mathcal{R}}^{\exists} \psi$. By the induction hypothesis we get $\mathfrak{s}\{y \leftarrow \ell\}(x) \in \text{dom}(\mathfrak{h})$, thus $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$ as $x \neq y$. If $\phi = \phi_1 \circ \phi_2$ or $\phi = \phi_1 * \phi_2$ then there exists $i \in \{1, 2\}$ such that $x \in \text{roots}(\phi_i)$. Moreover, there exist heaps \mathfrak{h}_i (for $i = 1, 2$) with $\mathfrak{h} = \mathfrak{h}_1 \sqcup \mathfrak{h}_2$ and $(\mathfrak{s}, \mathfrak{h}_i) \models_{\mathcal{R}}^{\exists} \phi_i$. By the induction hypothesis we get $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h}_i)$, thus $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$, as $\text{dom}(\mathfrak{h}) = \text{dom}(\mathfrak{h}_1) \cup \text{dom}(\mathfrak{h}_2)$.

B Proof of Theorem 9

The proof is by reduction from the Post Correspondence Problem (PCP), that is well-known to be undecidable. Consider a natural number n and two sequences of words μ_1, \dots, μ_n and ν_1, \dots, ν_n . A *potential witness* is a word ω such that for all $\lambda \in \{\mu, \nu\}$, there exists a sequence $I^\lambda(1), \dots, I^\lambda(\kappa^\lambda)$ with $\omega = \lambda_{I^\lambda(1)} \dots \lambda_{I^\lambda(\kappa^\lambda)}$ and $\kappa^\lambda > 0$. It is a *witness* if, moreover, $\kappa^\mu = \kappa^\nu = \kappa$ and $I^\mu(i) = I^\nu(i)$ for all $i \in \{1, \dots, \kappa\}$. The PCP consists in determining whether a witness exists for two given sequences of words. We assume, to ease the encoding, that the solution sequence $I^\lambda(1), \dots, I^\lambda(\kappa^\lambda)$ ends with a dummy index 0, with $\mu_0 = \nu_0 = \#$ (this entails that we must have $\kappa^\lambda > 1$). For every potential witness w , for every $\lambda \in \{\mu, \nu\}$ and for every $i \in \{1, \dots, \kappa^\lambda\}$, we denote by $\zeta^\lambda(i)$ the position of the word $\lambda_{I^\lambda(i)}$ inside w , formally defined as follows: $\zeta^\lambda(i) \stackrel{\text{def}}{=} |\lambda_{I^\lambda(1)} \dots \lambda_{I^\lambda(i-1)}| + 1$. For every $j \in \{1, \dots, |w|\}$, we also denote by $\eta^\lambda(j)$ the element $i \in \{1, \dots, \kappa^\lambda\}$ such that $\omega|_j$ occurs in $\lambda_{I^\lambda(i)}$, formally defined as follows: $\eta^\lambda(j) = i$ if $\zeta^\lambda(i) \leq j < \zeta^\lambda(i+1)$.

For simplicity, we assume that all the natural numbers in $\{1, \dots, n\}$ are taken as variables of sort **1**. Under this assumption, a potential witness ω may be represented by a structure $(\mathfrak{s}, \mathfrak{h})$ defined as follows: $\mathfrak{h} = \mathfrak{h}^\mu \sqcup \mathfrak{h}^\nu \sqcup \mathfrak{h}'$, with $\mathfrak{h}' \stackrel{\text{def}}{=} \{(\ell'_j, \ell^\mu_{\eta^\mu(j)}, \ell^\nu_{\eta^\nu(j)}, \ell'_{j+1}, \pi_1) \mid 1 \leq j \leq |\omega|\}$ and for all $\lambda \in \{\mu, \nu\}$: $\mathfrak{h}^\lambda \stackrel{\text{def}}{=} \{(\ell^\lambda_{i+1}, \mathfrak{s}(I^\lambda(i)), \ell^\lambda_i, \pi_1) \mid 1 \leq i \leq \kappa^\lambda\}$. Intuitively, every heap \mathfrak{h}^λ encodes the sequence $I^\lambda(\kappa^\lambda), \dots, I^\lambda(1)$ as a linked list in which the i -th element refers to (the image of) $I^\lambda(i)$ (note that the list is reversed, the reason will become clear later). Then \mathfrak{h}' encodes the word ω as a list where the j -th element of the list refers to the cells $\ell^\mu_{\eta^\mu(j)}$ and $\ell^\nu_{\eta^\nu(j)}$ in the lists \mathfrak{h}^μ and \mathfrak{h}^ν that correspond to the words in which $\omega|_j$ occurs.

Let \mathbf{v} be a sequence of variables containing all numbers of $\{1, \dots, n\}$, as well as a special variable `nil` (marking the end of the lists), and variables u_μ and u_ν denoting the first cell in the list \mathfrak{h}^λ and \mathfrak{h}^ν . It is straightforward to check that the following rules generate structures of the above form (for all $i_\mu, i_\nu, i'_\mu, i'_\nu \in \{1, \dots, n\}$ and for all $j_\mu, j_\nu \in \{1, \dots, M+1\}$, where M denotes the maximal length of the word in the sequences μ or ν).

Intuitively, P allocates the first element of the potential witness, by guessing the first element i in the solution sequence¹⁰, then call $P_{2,2,i,i}$. The predicate $P_{j_\mu, j_\nu, i_\mu, i_\nu}$ allocates the next elements, starting from a character occurring simultaneously at position j_μ in the word i_μ in the sequence μ and at position j_ν in the word i_ν in the sequence ν (thus initially $i_\mu = i_\nu = i$ and $j_\mu = j_\nu = 2$, as the first character of the first word has already been allocated by the predicate P). Conditions are added on the rules to ensure that the two characters are identical in both sequences. Each time one reaches the start of a new word $I^\lambda(i)$ in the witness, $Q(y, y', i, z)$ is called to allocate the elements of the list \mathfrak{h}^λ , where y' is a pointer to the *previous* element in the list $I^\lambda(1), \dots, I^\lambda(\kappa^\lambda)$ (thus initially $y' = \text{nil}$). If the two words end on the same character then this means that we have reached the end of the potential witness¹¹, and we only have to allocate the locations corresponding to the dummy index 0 and the dummy character $\#$. Note that the locations corresponding to 0 are associated with the variables u_μ and u_ν . The list must be constructed in reverse order to ensure that the obtained rules are \mathfrak{h} -regular (if the lists were constructed in the usual order, then the pointer to the next cell could not be allocated before one gets to the next word, hence the corresponding rule would not be \mathfrak{h} -regular, as one would need to introduce an existential variable without immediately allocating it).

$$\begin{aligned}
P(x, \mathbf{v}, z) \Leftarrow \exists x', y_\mu, y_\nu \quad & x \xrightarrow{z} (y_\mu, y_\nu, x') \\
& \circ Q(y_\mu, \text{nil}, i, z) \circ Q(y_\nu, \text{nil}, i, z) \circ P_{2,2,i,i}(x', y_\mu, y_\nu, \mathbf{v}, z) \\
& \text{if } \mu_i|_1 = \nu_i|_1 \quad (1)
\end{aligned}$$

$$Q(y, y', i, z) \Leftarrow y \xrightarrow{z} (i, y') \quad (2)$$

$$\begin{aligned}
P_{j_\mu, j_\nu, i_\mu, i_\nu}(x, y_\mu, y_\nu, \mathbf{v}, z) \Leftarrow \exists x' \quad & x \rightarrow (y_\mu, y_\nu, x', z) \\
& \circ P_{j_\mu+1, j_\nu+1, i_\mu, i_\nu}(x', y_\mu, y_\nu, \mathbf{v}, z) \\
& \text{if } j_\mu \leq |\mu_{i_\mu}|, j_\nu \leq |\nu_{i_\nu}| \text{ and } \mu_{i_\mu}|_{j_\mu} = \nu_{i_\nu}|_{j_\nu} \quad (3)
\end{aligned}$$

$$\begin{aligned}
P_{j_\mu, j_\nu, i_\mu, i_\nu}(x, y_\mu, y_\nu, \mathbf{v}, z) \Leftarrow \exists x, y'_\mu \quad & x \rightarrow (y'_\mu, y_\nu, x', z) \\
& \circ Q(y'_\mu, i'_\mu, y_\mu, z) \circ P_{2, j_\nu+1, i'_\mu, i_\nu}(x', y'_\mu, y_\nu, \mathbf{v}, z) \\
& \text{if } j_\mu = |\mu_{i_\mu}| + 1, j_\nu \leq |\nu_{i_\nu}| \text{ and } \mu_{i'_\mu}|_1 = \nu_{i_\nu}|_{j_\nu} \quad (4)
\end{aligned}$$

¹⁰ For simplicity we only encode potential witnesses where the two sequences of indices start with the same index, i.e., $I^\mu(1) = I^\nu(1)$. It is clear that this is not restrictive as we eventually want to encode witnesses.

¹¹ We assume, w.l.o.g., that only witnesses of minimal length are considered.

$$\begin{aligned}
P_{j_\mu, j_\nu, i_\mu, i_\nu}(x, y_\mu, y_\nu, \mathbf{v}, z) &\Leftarrow \exists x, y'_\nu \quad x \rightarrow (y_\mu, y'_\nu, x', z) \\
&\quad \circ Q(y'_\nu, i'_\nu, y_\nu, z) \circ P_{j_\mu+1, 2, i_\mu, i'_\nu}(x', y_\mu, y'_\nu, \mathbf{v}, z) \\
&\quad \text{if } j_\mu \leq |\mu_{i_\mu}|, j_\nu = |\nu_{i_\nu}| + 1 \text{ and } \mu_{i_\mu}|_{j_\mu} = \nu_{i_\nu}|_1 \quad (5)
\end{aligned}$$

$$\begin{aligned}
P_{j_\mu, j_\nu, i_\mu, i_\nu}(x, y_\mu, y_\nu, \mathbf{v}, z) &\Leftarrow \exists u', u'' \quad x \rightarrow (u', u'', \mathbf{nil}, z) \circ Q(u', y_\mu, 0, z) \\
&\quad \circ Q(u'', y_\nu, 0, z) \circ u' \simeq u_\mu \circ u'' \simeq u_\nu \text{ if } j_\mu = |\mu_{i_\mu}| + 1 \text{ and } j_\nu = |\nu_{i_\nu}| + 1 \quad (6)
\end{aligned}$$

To ensure that the considered potential witness is indeed a witness, it only remains to check that the sequences Q^μ and Q^ν are identical. To this aim, we introduce a predicate that generates structures of the form $\mathfrak{h}'' \sqcup \mathfrak{h}^\mu \sqcup \mathfrak{h}^\nu$ with $\mathfrak{h}'' \stackrel{\text{def}}{=} \{(\ell''_i, \ell^\mu_i, \ell^\nu_i, \ell''_{i+1}, \pi_2) \mid i \in \{1, \dots, \kappa\}\}$, $\mathfrak{h}^{\lambda'} \stackrel{\text{def}}{=} \{(\ell^\lambda_{i+1}, \mathfrak{s}(I^\lambda(i)), \ell^\lambda_i, \pi_2) \mid 1 \leq i \leq \kappa\}$ ($\mathfrak{h}^{\lambda'}$ is identical to \mathfrak{h}^λ , except for the permissions) and $\mathfrak{s}(I^\lambda) = \mathfrak{s}(I^\nu)$, for all $i \in \{1, \dots, \kappa\}$.

$$\begin{aligned}
R(x, y_\mu, y_\nu, \mathbf{v}, z) &\Leftarrow \exists x', y'_\mu, y'_\nu \quad x \xrightarrow{z} (y'_\mu, y'_\nu, x') \circ R(x', y'_\mu, y'_\nu, z) \\
&\quad \circ Q(y'_\mu, y_\mu, i, z) \circ Q(y'_\nu, y_\nu, i, z) \quad (7)
\end{aligned}$$

$$\begin{aligned}
R(x, y_\mu, y_\nu, \mathbf{v}, z) &\Leftarrow \exists u', u'' \quad x \xrightarrow{z} (u', u'', \mathbf{nil}) \circ Q(u', y_\mu, i, z) \circ Q(u'', y_\nu, i, z) \\
&\quad \circ u' \simeq u_\mu \circ u'' \simeq u_\nu \quad (8)
\end{aligned}$$

It is clear that the formula $P(x_1, \mathbf{v}, z_1) \circ R(x_2, \mathbf{nil}, \mathbf{nil}, \mathbf{v}, z_2)$ is $(\mathcal{R}, \mathfrak{F})$ -satisfiable iff the considered instance of the PCP admits a solution.

C Proof of Lemma 15

The proof is by induction on $|\mathfrak{h}|$.

- Assume that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{F}} P(\mathbf{x}, \mathbf{p})[v]^-$. By definition of the semantics of predicate atoms, there exists a formula ψ such that $P(\mathbf{x}, \mathbf{p})[v]^- \Leftarrow_{\mathcal{R}} \psi$ and $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{F}} \psi$. By definition of the rules defining $P(\mathbf{x}, \mathbf{p})[v]^-$ and of the relation $\Leftarrow_{\mathcal{R}}$, this entails that \mathcal{R} contains a rule $P(\mathbf{y}, \mathbf{z}) \Leftarrow \exists \mathbf{u} (\bigcirc_{i=1}^m Q_i(\mathbf{y}_i, \mathbf{p}_i) \circ \phi)$ and we have:

$$\psi = \exists \mathbf{u} (\bigcirc_{i=1}^m Q_i(\mathbf{y}_i, \mathbf{p}_i)[w]^- \circ \phi \circ \mathbf{y}|_1 \neq w) \sigma$$

with $\sigma = \{\mathbf{y} \leftarrow \mathbf{x}, w \leftarrow v, \mathbf{z} \leftarrow \mathbf{p}\}$. Again by definition of the semantics, there exist a sequence of locations ℓ and disjoint heaps $\mathfrak{h}_0, \dots, \mathfrak{h}_m$ such that $\mathfrak{h} = \mathfrak{h}_0 \sqcup \dots \sqcup \mathfrak{h}_m$, $(\mathfrak{s}', \mathfrak{h}_0) \models_{\mathcal{R}}^{\mathfrak{F}} \phi \sigma$, $\mathfrak{s}'(\mathbf{y}|_1 \sigma) \neq \mathfrak{s}'(w \sigma)$ and $(\mathfrak{s}', \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{F}} Q_i(\mathbf{y}_i, \mathbf{p})[w]^- \sigma = Q_i(\mathbf{y}_i \sigma, \mathbf{p} \sigma)[v]^-$ (for all $i \in \{1, \dots, m\}$), with $\mathfrak{s}' = \mathfrak{s}\{\mathbf{u} \leftarrow \ell\}$. Since \mathcal{R} is \mathfrak{h} -regular, ϕ contains exactly one points-to atom with left-hand side $\mathbf{y}|_1$, hence $\text{dom}(\mathfrak{h}_0) = \{\mathfrak{s}'(\mathbf{y}|_1 \sigma)\} = \{\mathfrak{s}(\mathbf{x}|_1)\}$. Thus $\mathfrak{h}_0 \neq \emptyset$ and $|\mathfrak{h}_i| < |\mathfrak{h}|$ (for all $i \in \{1, \dots, m\}$). By the induction hypothesis this entails that $(\mathfrak{s}', \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{F}} Q_i(\mathbf{y}_i \sigma, \mathbf{p} \sigma)$ and $\mathfrak{s}(v) = \mathfrak{s}'(v) \notin \text{dom}(\mathfrak{h}_i)$. As $\mathfrak{s}'(\mathbf{y}|_1 \sigma) \neq \mathfrak{s}'(w \sigma)$, we have $\mathfrak{s}(\mathbf{x}|_1) \neq \mathfrak{s}(v)$ i.e. $\mathfrak{s}(v) \notin \text{dom}(\mathfrak{h}_0)$, hence $\mathfrak{s}(v) \notin \text{dom}(\mathfrak{h}) = \bigcup_{i=0}^m \text{dom}(\mathfrak{h}_i)$. Furthermore, $(\mathfrak{s}', \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{F}} \bigcirc_{i=1}^m Q_i(\mathbf{y}_i \sigma, \mathbf{p} \sigma) \circ \phi \sigma = (\bigcirc_{i=1}^m Q_i(\mathbf{y}_i, \mathbf{p}) \circ \phi) \sigma$ and $P(\mathbf{x}, \mathbf{p}) \Leftarrow_{\mathcal{R}} \exists \mathbf{u} (\bigcirc_{i=1}^m Q_i(\mathbf{y}_i, \mathbf{p}) \circ \phi) \sigma$, consequently $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{F}} P(\mathbf{x}, \mathbf{p})$.

- Conversely, assume that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} P(\mathbf{x}, \mathbf{p})$ and that $\mathfrak{s}(v) \notin \text{dom}(\mathfrak{h})$. By definition of the semantics, this entails that \mathcal{R} contains a rule $P(\mathbf{y}, \mathbf{z}) \Leftarrow \psi$ with $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \psi\sigma$ and $\sigma = \{\mathbf{y} \leftarrow \mathbf{x}, \mathbf{z} \leftarrow \mathbf{p}\}$. As \mathcal{R} is \mathfrak{h} -regular, ψ is necessarily of the form $\exists \mathbf{u} (\bigcirc_{i=1}^m Q_i(\mathbf{y}_i, \mathbf{p}_i) \circ \phi)$, where ϕ contains no predicate atom and exactly one points-to atom with left-hand side $y|_1$. Moreover, there exist a vector of locations ℓ and disjoint heaps $\mathfrak{h}_0, \dots, \mathfrak{h}_m$ such that $(\mathfrak{s}', \mathfrak{h}_0) \models_{\mathcal{R}}^{\mathfrak{P}} \phi\sigma$, $(\mathfrak{s}', \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} Q_i(\mathbf{y}_i, \mathbf{p}_i)\sigma$ and $\mathfrak{s}' = \mathfrak{s}\{\mathbf{u} \leftarrow \ell\}$. Since $\mathfrak{s}'(v) = \mathfrak{s}(v) \notin \text{dom}(\mathfrak{h})$, we get $\mathfrak{s}'(v) \notin \text{dom}(\mathfrak{h})$ thus $\mathfrak{s}'(v) \notin \text{dom}(\mathfrak{h}_i)$. Moreover, $\text{dom}(\mathfrak{h}_0) = \{\mathfrak{s}'(\mathbf{y}|_1\sigma)\}$, hence $\mathfrak{s}'(v) \neq \mathfrak{s}'(\mathbf{y}|_1\sigma)$ and $|\mathfrak{h}_i| < |\mathfrak{h}|$ (for all $i \in \{1, \dots, m\}$). By the induction hypothesis, this entails that $(\mathfrak{s}', \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} Q_i(\mathbf{y}_i\sigma, \mathbf{p}\sigma)[v]^-$. By definition of the rules defining $P(\mathbf{x}, \mathbf{p})[v]^-$, $P(\mathbf{x}, \mathbf{p})[v]^- \Leftarrow_{\mathcal{R}} \exists \mathbf{u} (\bigcirc_{i=1}^m Q_i(\mathbf{y}_i, \mathbf{p})[w]^- \circ \phi \circ y|_1 \neq w)\sigma'$, with $\sigma' = \sigma\{w \leftarrow v\}$. We get $(\mathfrak{s}', \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} Q_i(\mathbf{y}_i, \mathbf{p})[w]^- \sigma'$ (for all $i \in \{1, \dots, m\}$), $(\mathfrak{s}', \mathfrak{h}_0) \models_{\mathcal{R}}^{\mathfrak{P}} \phi\sigma = \phi\sigma'$ and $\mathfrak{s}'(y|_1\sigma) \neq \mathfrak{s}'(w\sigma') = \mathfrak{s}'(v)$. Thus $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} P(\mathbf{x}, \mathbf{p})[v]^-$.

D Proof of Lemma 19

Assume that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} P(\mathbf{x}, \mathbf{p}) \circ ((P(\mathbf{x}, \mathbf{p}) \circ \alpha) \multimap Q(\mathbf{y}, \mathbf{q}))$. By definition, there exist disjoint heaps $\mathfrak{h}_1, \mathfrak{h}_2$ such that $\mathfrak{h} = \mathfrak{h}_1 \sqcup \mathfrak{h}_2$, $(\mathfrak{s}, \mathfrak{h}_1) \models_{\mathcal{R}}^{\mathfrak{P}} P(\mathbf{x}, \mathbf{p})$ and $(\mathfrak{s}, \mathfrak{h}_2) \models_{\mathcal{R}}^{\mathfrak{P}} (P(\mathbf{x}, \mathbf{p}) \circ \alpha) \multimap Q(\mathbf{y}, \mathbf{q})$. We show, by induction on $|\mathfrak{h}_2|$, that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \alpha \multimap Q(\mathbf{y}, \mathbf{q})$. By definition of the rules defining $(P(\mathbf{x}, \mathbf{p}) \circ \alpha) \multimap Q(\mathbf{y}, \mathbf{q})$, necessarily $(\mathfrak{s}, \mathfrak{h}_2) \models_{\mathcal{R}}^{\mathfrak{P}} \exists \mathbf{w} (\psi_1 \circ \dots \circ \psi_m \circ \phi)$, with $P(\mathbf{x}, \mathbf{p}) \Leftarrow_{\mathcal{R}} \bigcirc_{i=1}^m Q_i(\mathbf{u}_i, \mathbf{p}_i) \circ \phi$, $\alpha_1 \circ \dots \circ \alpha_m$ is a decomposition of $P(\mathbf{x}, \mathbf{p}) \circ \alpha$ (up to AC and neutrality of emp), and for every $i \in \{1, \dots, m\}$, either $\psi_i = \alpha_i \multimap Q_i(\mathbf{u}_i, \mathbf{q}_i)$ or $\psi_i = (\mathbf{y}_j \simeq \mathbf{u}_i \circ \mathbf{p}_j \simeq \mathbf{q}_i)$ with $\alpha_i = P_j(\mathbf{y}_j, \mathbf{p}_j)$ and $P_j = Q_i$. Thus there exist locations ℓ and disjoint heaps $\mathfrak{h}_2^0, \dots, \mathfrak{h}_2^m$ such that $(\mathfrak{s}', \mathfrak{h}_2^0) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$, $(\mathfrak{s}', \mathfrak{h}_2^i) \models_{\mathcal{R}}^{\mathfrak{P}} \psi_i$ (for all $i \in \{1, \dots, m\}$) and $\mathfrak{s}' = \mathfrak{s}\{\mathbf{w} \leftarrow \ell\}$. By definition, there exists $i \in \{1, \dots, m\}$ such that $P(\mathbf{x}, \mathbf{p})$ occurs in α_i , we assume by symmetry that $i = 1$, with $\alpha_1 = P(\mathbf{x}, \mathbf{p}) \circ \alpha'_1$. As \mathcal{R} is \mathfrak{h} -regular, ϕ contains exactly one points-to atom, thus $\mathfrak{h}_0 \neq \emptyset$ and $|\mathfrak{h}_i| < |\mathfrak{h}|$ (for all $i \in \{1, \dots, m\}$). We distinguish two cases.

- Assume that ψ_1 is of the first form above. We have $(\mathfrak{s}', \mathfrak{h}_1 \sqcup \mathfrak{h}_2^1) \models_{\mathcal{R}}^{\mathfrak{P}} P(\mathbf{x}, \mathbf{p}) \circ ((P(\mathbf{x}, \mathbf{p}) \circ \alpha'_1) \multimap Q_1(\mathbf{y}_1, \mathbf{q}_1))$. By the induction hypothesis, we get $(\mathfrak{s}', \mathfrak{h}_1 \sqcup \mathfrak{h}_2^1) \models_{\mathcal{R}}^{\mathfrak{P}} \alpha'_1 \multimap Q_1(\mathbf{y}_1, \mathbf{q}_1)$. This entails that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \exists \mathbf{w} (\alpha'_1 \multimap Q_1(\mathbf{y}_1, \mathbf{q}_1) \circ \bigcirc_{i=2}^m \psi_m \circ \phi)$, hence $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \alpha \multimap Q(\mathbf{y}, \mathbf{q}_1)$, by definition of the rules defining $\alpha \multimap Q(\mathbf{y}, \mathbf{q}_1)$ (since $\alpha = \alpha'_1 \circ \bigcirc_{i=2}^m \alpha_i$).
- Assume that ψ_1 is of the second form above. Then $\mathfrak{s}'(\mathbf{x}) = \mathfrak{s}'(\mathbf{y}_1)$, $\mathfrak{s}'(\mathbf{p}) = \mathfrak{s}'(\mathbf{q}_1)$, and $Q_1 = P$ (with $\alpha' = \text{emp}$), so that $(\mathfrak{s}', \mathfrak{h}_1) \models_{\mathcal{R}}^{\mathfrak{P}} Q_1(\mathbf{y}_1, \mathbf{q}) = \text{emp} \multimap Q_1(\mathbf{y}_1, \mathbf{p})$. Therefore $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \exists \mathbf{w} (\text{emp} \multimap Q_1(\mathbf{y}_1, \mathbf{q}_1) \circ \bigcirc_{i=2}^m \psi_m \circ \phi)$, with $\alpha = \text{emp} \circ \bigcirc_{i=2}^m \alpha_i$. This entails the result, by definition of the rules defining $\alpha \multimap Q(\mathbf{y}, \mathbf{q}_1)$.

E Proof of Lemma 20

The proof is by induction on $|\mathfrak{h}|$. Assume that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} Q(\mathbf{y}, \mathbf{p})$, $\mathfrak{s}(x) \neq \mathbf{y}|_1$ and $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$. By definition of the semantics, there exists a formula ψ such that $Q(\mathbf{y}, \mathbf{p}) \Leftarrow_{\mathcal{R}} \psi$ and $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \psi$, and, since \mathcal{R} is \mathfrak{h} -regular, ψ is of the form $\exists \mathbf{w} (\bigcirc_{i=1}^n Q_i(\mathbf{u}_i, \mathbf{p}_i) \circ \phi)$, where ϕ contains exactly one points-to atom (which left-hand side is necessarily $\mathbf{y}|_1$) and $\mathbf{p}_i \subseteq \mathbf{p}$ (for all $i = 1, \dots, n$). Therefore there exist a sequence of locations ℓ and disjoint heaps $\mathfrak{h}_0, \dots, \mathfrak{h}_n$ such that $\mathfrak{h} = \mathfrak{h}_0 \sqcup \dots \sqcup \mathfrak{h}_n$, $(\mathfrak{s}', \mathfrak{h}_0) \models_{\mathcal{R}}^{\exists} \phi$ and $(\mathfrak{s}', \mathfrak{h}_i) \models_{\mathcal{R}}^{\exists} Q_i(\mathbf{u}_i, \mathbf{p}_i)$, for all $i \in \{1, \dots, n\}$, with $\mathfrak{s}' = \mathfrak{s}\{\mathbf{w} \leftarrow \ell\}$. Furthermore, we have $\text{dom}(\mathfrak{h}_0) = \{\mathfrak{s}'(\mathbf{y}|_1)\} \neq \emptyset$, and thus $|\mathfrak{h}_i| < |\mathfrak{h}|$, for all $i \in \{1, \dots, n\}$. Since $\mathfrak{s}(x) \neq \mathfrak{s}(\mathbf{y}|_1)$ we have $\mathfrak{s}(x) \notin \text{dom}(\mathfrak{h}_0)$, and, as $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h})$, this entails that $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h}_i)$, for some $i \in \{1, \dots, n\}$. Assume by symmetry that $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h}_1)$. We distinguish two cases:

- If $\mathfrak{s}'(\mathbf{u}_1|_1) \neq \mathfrak{s}'(x)$, then, by the induction hypothesis, we deduce that there exist atoms $P(x, \mathbf{z}, \mathbf{q})$ and $P_i(x, \mathbf{y}_i, \mathbf{q}_i)$ (for $i \in \{1, \dots, m\}$) such that

$$(\mathfrak{s}', \mathfrak{h}_1) \models_{\mathcal{R}}^{\exists} \exists \mathbf{x} (\beta \circ (\beta \multimap Q_1(\mathbf{u}_1, \mathbf{p}_1)))$$

with $\mathbf{x} = (x_1, \dots, x_n)$, $\beta = P(x, \mathbf{z}, \mathbf{q}) \circ \bigcirc_{i=1}^m P_i(x_i, \mathbf{y}_i, \mathbf{q}_i)$, $\mathbf{z} \subseteq \mathbf{u}_1 \cup \mathbf{x}$, $\mathbf{q} \cup \mathbf{q}_i \subseteq \mathbf{p}_1 \subseteq \mathbf{p}$, $\mathbf{y}_i \subseteq \{\mathbf{u}_1|_j \mid j \notin \gamma_{\mathcal{R}}(Q_1)\}$ (for all $i \in \{1, \dots, m\}$), $P_i \in \mathcal{P}^*$, $\mathbf{x} \subseteq (x, \mathbf{z})|_{\gamma_{\mathcal{R}}(P)}$ and $\mathbf{y} \in \mathbf{u}_1 \cap \mathbf{z} \wedge \mathbf{y} \notin \{\mathbf{u}_1|_j \mid j \notin \gamma_{\mathcal{R}}(Q_1)\} \implies \mathbf{y} \in (x, \mathbf{z})|_{\gamma_{\mathcal{R}}(P)}$. By α -renaming, we assume that $\mathbf{x} \cap \mathbf{y} = \emptyset$. Thus there exist a store \mathfrak{s}'' coinciding with \mathfrak{s}' on all variables not occurring in \mathbf{x} , and disjoint heaps $\mathfrak{h}'_1, \mathfrak{h}''_1$ such that $\mathfrak{h}_1 = \mathfrak{h}'_1 \cup \mathfrak{h}''_1$, $(\mathfrak{s}'', \mathfrak{h}'_1) \models_{\mathcal{R}}^{\exists} \beta$ and $(\mathfrak{s}'', \mathfrak{h}''_1) \models_{\mathcal{R}}^{\exists} (\beta \multimap Q_1(\mathbf{u}_1, \mathbf{p}_1))$.

We show that $\mathbf{y} \in \mathbf{y} \cap \mathbf{z} \wedge \mathbf{y} \notin \{\mathbf{y}|_j \mid j \notin \gamma_{\mathcal{R}}(Q)\} \implies \mathbf{y} \in (x, \mathbf{z})|_{\gamma_{\mathcal{R}}(P)}$. Assume that $\mathbf{z} \in \mathbf{y} \cap \mathbf{z}$, with $\mathbf{z} \notin \{\mathbf{y}|_j \mid j \notin \gamma_{\mathcal{R}}(Q)\}$. As $\mathbf{z} \subseteq \mathbf{x} \cup \mathbf{u}_1$, necessarily $\mathbf{z} = \mathbf{u}_1|_{j'}$ for some $j' \in \{1, \dots, \#_1(Q_1)\}$. Assume that $j' \notin \gamma_{\mathcal{R}}(Q_1)$. By Condition 2 in Definition 10 we deduce that there exists $j \notin \gamma_{\mathcal{R}}(Q)$ such that $\mathbf{z} = \mathbf{y}|_j$, which contradicts the above assumption. Otherwise (i.e., if $\mathbf{z} = \mathbf{u}_1|_{j'} \implies j' \in \gamma_{\mathcal{R}}(Q_1)$, for all j') then $\mathbf{y} \notin \{\mathbf{u}_1|_j \mid j \notin \gamma_{\mathcal{R}}(Q_1)\}$ and the proof follows from the induction hypothesis above.

We show that $\mathbf{y}_i \subseteq \{\mathbf{y}|_{j'} \mid j' \notin \gamma_{\mathcal{R}}(Q)\}$, for all $i \in \{1, \dots, m\}$. Let $v \in \mathbf{y}_i$. Since $\mathbf{y}_i \subseteq \{\mathbf{u}_1|_j \mid j \notin \gamma_{\mathcal{R}}(Q_1)\}$, we get $v = \mathbf{u}_1|_j$, for some $j \notin \gamma_{\mathcal{R}}(Q_1)$. By definition of $\gamma_{\mathcal{R}}(Q_1)$ (see Conditions 1 and 2 in Definition 10), this entails that $v = \mathbf{y}|_{j'}$, for some $j' \notin \gamma_{\mathcal{R}}(Q)$.

We show that $\mathbf{z} \subseteq \mathbf{y} \cup \mathbf{x} \cup \{\mathbf{u}_i|_1 \mid i = 2, \dots, n\}$. Let V be the set of variables occurring in \mathbf{z} , but not in $\mathbf{y} \cup \mathbf{x}$. Consider any variable $v \in V$. Necessarily, $v \in \mathbf{u}_1$ (as $\mathbf{z} \subseteq \mathbf{u}_1 \cup \mathbf{x}$), and $v \in \mathbf{w}$ (as $\mathbf{u}_1 \subseteq \mathbf{y} \cup \mathbf{w}$, by definition of the unfolding). Note that, by Condition 1 in Definition 10, this entails that $v \in \mathbf{u}_1$ with $v = \mathbf{u}_1|_j \implies j \in \gamma_{\mathcal{R}}(Q_1)$, so that $v \in (x, \mathbf{z})|_{\gamma_{\mathcal{R}}(P)}$. As \mathcal{R} is \mathfrak{h} -regular and $v \in \mathbf{w}$, necessarily there exists $k \in \{2, \dots, m\}$ such that $v = \mathbf{u}_k|_1$. We assume, by symmetry, that the set of indices k such that $\mathbf{u}_k|_1 \in V$ is of the form $\{2, \dots, n'\}$ (with possibly $n' = 1$, as V may be empty), so that $\mathbf{z} \subseteq \mathbf{y} \cup \mathbf{x} \cup \{\mathbf{u}_i|_1 \mid i = 2, \dots, n'\}$.

Since \mathcal{R} is \exists -restricted (see Condition 5 in Definition 10) we get $\gamma_{\mathcal{R}}(Q_k) = \emptyset$, for all $k \in \{2, \dots, n'\}$, so that $\{\mathbf{u}_k|_j \mid j \neq 1\} \subseteq \{\mathbf{y}|_j \mid j \notin \gamma_{\mathcal{R}}(Q)\}$.

Let $\mathfrak{h}' \stackrel{\text{def}}{=} \mathfrak{h}_0 \sqcup \mathfrak{h}_1 \sqcup \mathfrak{h}_{n'+1} \sqcup \dots \sqcup \mathfrak{h}_n$ and $\mathfrak{h}'' \stackrel{\text{def}}{=} \mathfrak{h}_2 \sqcup \dots \sqcup \mathfrak{h}_{n'}$. Observe that \mathfrak{h}' and \mathfrak{h}'' are disjoint and that $\mathfrak{h} = \mathfrak{h}' \sqcup \mathfrak{h}''$. We get: $(\mathfrak{s}'', \mathfrak{h}') \models_{\mathcal{R}}^{\exists} (\beta \dashv\bullet Q(\mathbf{y}, \mathbf{p})) \circ \bigcirc_{i=n'+1}^n Q_i(\mathbf{y}_i, \mathbf{q}_i) \circ \phi$. By definition of the rules defining $\beta \dashv\bullet Q(\mathbf{y}, \mathbf{p})$, this entails that $(\mathfrak{s}'', \mathfrak{h}') \models_{\mathcal{R}}^{\exists} \beta' \dashv\bullet Q(\mathbf{y}, \mathbf{p})$ with $\beta' \stackrel{\text{def}}{=} \beta \circ \bigcirc_{i=2}^{n'} Q_i(\mathbf{u}_i, \mathbf{p}_i)$. Moreover, $(\mathfrak{s}'', \mathfrak{h}'') \models_{\mathcal{R}}^{\exists} P(x, \mathbf{z}, \mathbf{q}) \circ \bigcirc_{i=1}^m P_i(x_i, \mathbf{y}_i, \mathbf{q}_i) \circ \bigcirc_{i=2}^{n'} Q_i(\mathbf{u}_i, \mathbf{q}_i)$. Thus:

$$(\mathfrak{s}'', \mathfrak{h}) \models_{\mathcal{R}}^{\exists} P(x, \mathbf{z}, \mathbf{q}) \circ \bigcirc_{i=1}^m P_i(x_i, \mathbf{y}_i, \mathbf{q}_i) \circ \bigcirc_{i=2}^{n'} Q_i(\mathbf{u}_i, \mathbf{q}_i) \circ (\beta' \dashv\bullet Q(\mathbf{y}, \mathbf{p}))$$

with $\beta' = P(x, \mathbf{z}, \mathbf{q}) \circ \bigcirc_{i=1}^m P_i(x_i, \mathbf{y}_i, \mathbf{q}_i) \circ \bigcirc_{i=2}^{n'} Q_i(\mathbf{u}_i, \mathbf{p}_i)$. Consequently,

$$(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \exists \mathbf{v} (P(x, \mathbf{z}, \mathbf{q}) \circ \bigcirc_{i=1}^m P_i(x_i, \mathbf{y}_i, \mathbf{q}_i) \circ \bigcirc_{i=2}^{n'} Q_i(\mathbf{u}_i, \mathbf{q}_i) \circ (\beta' \dashv\bullet Q(\mathbf{y}, \mathbf{p})))$$

with $\mathbf{v} = \mathbf{x} \cup V = \mathbf{x} \cup \{\mathbf{u}_i|_1 \mid i = 2, \dots, n'\}$, which completes the proof, as we have proven that $\mathbf{z} \subseteq \mathbf{y} \cup \mathbf{x} \cup \{\mathbf{u}_i|_1 \mid i = 2, \dots, n'\}$, $\mathbf{y}_i \cup \{\mathbf{u}_k|_j \mid j \neq 1\} \subseteq \{\mathbf{y}|_{j'} \mid j' \notin \gamma_{\mathcal{R}}(Q)\}$, $\mathbf{q} \cup \mathbf{q}_i \cup \mathbf{p}_i \subseteq \mathbf{p}$, $P_i \in \mathcal{P}^*$ and $\mathbf{x} \cup V \cup (\mathbf{y}|_{\gamma_{\mathcal{R}}(Q)} \cap \mathbf{z}) \subseteq (x, \mathbf{z})|_{\gamma_{\mathcal{R}}(P)}$.

– If $\mathfrak{s}'(\mathbf{u}_1|_1) = \mathfrak{s}'(x)$, i.e., $\mathfrak{s}'(\mathbf{u}_1|_1) = \mathfrak{s}(x)$, then we let $P = Q_1$, $\mathbf{p} = \mathbf{p}_1$, and $\mathbf{u}_1|_1 \cdot \mathbf{z} = \mathbf{u}_1$. As in the previous case, the variables \mathbf{v} occurring in \mathbf{z} but not in \mathbf{y} must occur in \mathbf{w} , thus must be the root of some atom $Q_i(\mathbf{u}_i, \mathbf{p}_i)$, for some $i \in \{2, \dots, n'\}$ with $n' \in \{1, \dots, n\}$, $Q_i \in \mathcal{P}^*$ and $\{\mathbf{u}_i|_j \mid j \neq 1\} \subseteq \{\mathbf{y}|_j \mid j \notin \gamma_{\mathcal{R}}(Q)\}$. By definition of $\dashv\bullet$, we have $(\mathfrak{s}'', \mathfrak{h}) \models_{\mathcal{R}}^{\exists} P(x, \mathbf{z}, \mathbf{q}) \circ \bigcirc_{i=2}^{n'} Q_i(\mathbf{u}_i, \mathbf{q}_i) \circ (\beta \dashv\bullet Q(\mathbf{y}, \mathbf{p}))$ with $\beta = P(x, \mathbf{z}, \mathbf{q}) \circ \bigcirc_{i=2}^{n'} Q_i(\mathbf{u}_i, \mathbf{p}_i)$, $\mathbf{z} \subseteq \mathbf{y} \cup \{\mathbf{u}_i|_1 \mid i = 2, \dots, n'\}$, $\mathbf{u}_i \subseteq \{\mathbf{y}|_{j'} \mid j' \notin \gamma_{\mathcal{R}}(Q)\}$ and $\mathbf{q} \cup \mathbf{p}_i \subseteq \mathbf{p}$, so that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \exists \mathbf{v} (P(x, \mathbf{z}, \mathbf{q}) \circ \bigcirc_{i=2}^{n'} Q_i(\mathbf{u}_i, \mathbf{q}_i) \circ (\beta \dashv\bullet Q(\mathbf{y}, \mathbf{p})))$. By definition of $\gamma_{\mathcal{R}}$, we have $\mathbf{v} \subseteq \mathbf{u}_1|_{\gamma_{\mathcal{R}}(P)}$ and $y \in \mathbf{y} \cap \mathbf{z} \wedge y \notin \{\mathbf{y}|_j \mid j \notin \gamma_{\mathcal{R}}(Q)\} \implies y \in (x, \mathbf{z})|_{\gamma_{\mathcal{R}}(P)}$. Again, the proof is completed.

F Proof of Lemma 23

We need the following:

Proposition 34. *Let δ be a predicate atom and let ϕ be a formula. If $x \in \text{roots}(\phi)$ then $\delta \circ \phi \equiv_{\mathcal{R}}^{\exists} \delta[x]^- \circ \phi$.*

Proof. The result is an immediate consequence of Lemma 15 and Proposition 8 (using the fact that x cannot be allocated in both parts of the heaps that correspond to $\text{roots}(\phi)$ and ϕ , respectively).

We apply the following (non deterministic) transformation, for all variables $x \in \text{fv}(\phi)$. We replace every atom δ occurring in ϕ such that $x \notin \text{roots}(\delta) \cup \text{unalloc}(\delta)$ by some (indeterministically chosen) formula in $\{\delta[x]^=, \delta[x]^-\} \cup \delta[x]^+$, yielding a set of formulas $\Phi(x)$. In the case where δ is replaced by $\delta[x]^=$, since by definition $\delta[x]^= = \delta \circ (x \simeq y)$ with $\{y\} = \text{roots}(\delta)$, the variable y may be replaced by x in the entire formula, so that we get an atom with root x . By Lemma 22, for all structures $(\mathfrak{s}, \mathfrak{h})$, $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \phi$ iff $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\exists} \psi$, for some formula $\psi \in \Phi(x)$, thus satisfiability is preserved. Observe that, for all formulas $\psi \in \Phi$ and for all atoms δ' in ψ , either x occurs at the root of some predicate atom

occurring in the same \circ -formula as δ' in ψ , or $x \in \text{unalloc}(\delta')$. The former case occurs when δ' occurs in $\delta[x]^-$ (after the replacement of y by x) or in some formula in $\delta[x]^+$, as, by definition, x occurs as a root in all formulas in $\delta[x]^+$, and all such formulas are \circ -formulas. The latter case occurs when δ' is of the form $\delta[x]^-$, as $x \in \text{unalloc}(\delta[x]^-)$, by definition of $\delta[x]^-$. Since the same variable x cannot be allocated in distinct parts of the heaps, this entails, by Lemma 15, that any atom δ' in ϕ with $x \notin \text{roots}(\delta) \cup \text{unalloc}(\delta)$ can be replaced by $\delta[x]^-$ without affecting the semantics of the formula. We then get a formula ϕ' such that $x \in \text{roots}(\delta') \cup \text{unalloc}(\delta')$ holds for all atom δ' in ϕ' , so that the implication $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \delta' \wedge \mathfrak{s}(x) \in \text{dom}(\mathfrak{h}) \implies x \in \text{roots}(\delta')$ holds, by Propositions 8 and 34.

This process is applied on every variable x , which eventually yield a set of normalized formulas Ψ . Note that new variables may be introduced into the formula, as $\delta[x]^+$ may contain variables (namely the variables x_1, \dots, x_n in Definition 21) not occurring in δ . The above transformation must be applied also on such variables, which, in principle, could lead to non termination. We prove that the algorithm terminates, by showing that no new variables are added when x does not occur in the initial formula ϕ (more precisely, we prove that the obtained formula is always $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable in this case, hence can be removed from the set of formulas at hand). By Lemma 20, as x is necessarily added by some replacement $\delta \rightarrow \delta[y]^+$, we observe the variable x necessarily occurs as the root of some atom $Q_i(x, \mathbf{y})$, with $Q_i \in \mathcal{P}^*$. We denote by ξ_1 the \circ -formula containing $Q_i(x, \mathbf{y})$. Assume that the transformation above is in turn applied on the variable x , yielding a formula ψ . By definition of $\delta[x]^+$, this entails that x occurs as the root of some atom $P(x, \mathbf{z})$. Moreover, if new variables are added, then, by Lemma 20, we must have $\gamma_{\mathcal{R}}(P) \neq \emptyset$, thus $P \in \mathcal{P} \setminus \mathcal{P}^*$ (by Condition 3 in Definition 10). We denote by ξ_2 the \circ -formula of ψ containing $P(x, \mathbf{z})$. The case where $\xi_1 = \xi_2$ can be dismissed, as x would occur twice as a root in the same \circ -formula, which would then be unsatisfiable. Therefore, for every model $(\mathfrak{s}, \mathfrak{h})$ of ψ , there exist heaps $\mathfrak{h}_1, \mathfrak{h}_2$ such that $(\mathfrak{s}, \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \xi_i$ and $\mathfrak{h}_1 \sqcup \mathfrak{h}_2 \leq \mathfrak{h}$. Observe that each time an atom of the form $\alpha \multimap P'(x, \mathbf{u})$ is introduced in the formula, the \circ -formula α is simultaneously added in the same \circ -formula, and we have (by Lemma 19) $\alpha \circ (\alpha \multimap P'(x, \mathbf{u})) \models_{\mathcal{R}}^{\mathfrak{P}} P'(x, \mathbf{u})$. Moreover, for every atom $P'(x, \mathbf{u})$ and variable y , we have (by Lemma 15) $P'(x, \mathbf{u})[y]^- \models_{\mathcal{R}}^{\mathfrak{P}} P'(x, \mathbf{u})$. By an easy induction on the structure of derived predicates, we deduce that for every atom $Q(x, \mathbf{u}')$ occurring in some \circ -formula ξ in ψ , and every model $(\mathfrak{s}, \mathfrak{h})$ of ξ , there exist an atom $Q'(x, \mathbf{u})$ and a heap \mathfrak{h}' such that $(\mathfrak{s}, \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} Q'(x, \mathbf{u})$ and $\mathfrak{h}' \leq \mathfrak{h}$, Q is either equal to Q' or derived from Q' , and Q' is not a derived predicate, i.e., occurs in the *initial* SID. Moreover, by Propositions 14 and 18, we have $Q' \in \mathcal{P}^* \iff Q \in \mathcal{P}^*$. Consequently, for all heaps $\mathfrak{h}_1, \mathfrak{h}_2$ such that $(\mathfrak{s}, \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \xi_i$ there exist heaps $\mathfrak{h}'_1, \mathfrak{h}'_2$ such that $(\mathfrak{s}, \mathfrak{h}'_1) \models_{\mathcal{R}}^{\mathfrak{P}} Q'_i(x_i, \mathbf{y}')$, $(\mathfrak{s}, \mathfrak{h}'_2) \models_{\mathcal{R}}^{\mathfrak{P}} P'(x_i, \mathbf{y}'')$, $\mathfrak{h}'_i \leq_1 \mathfrak{h}_i$ with $Q'_i \in \mathcal{P}^*$, $P' \in \mathcal{P} \setminus \mathcal{P}^*$ and Q'_i, P' are not derived predicates. However, by Condition 4 in Definition 10, $Q'_i(x_i, \mathbf{y}') * P'(x_i, \mathbf{y}'')$ is $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable. This entails that $\xi_1 * \xi_2$ (hence ψ) is also $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable.

G Proof of Lemma 25

We use the following result, that is an easy consequence of Definition 24:

Proposition 35. $cut(L_1 \cup L_2, L', \mathfrak{h}) = cut(L_1, L' \cup L'_1, \mathfrak{h}) \cup cut(L_2, L' \cup L'_2, \mathfrak{h})$,
if $L_1 \cap L_2 = \emptyset$ and $L'_i \subseteq L_{3-i}$.

The proof of Lemma 25 is by induction on the satisfiability relation:

- If ϕ is pure then $roots(\phi) = dom(\mathfrak{h}') = \emptyset$ and the proof is immediate.
- If ϕ is a points-to atom $x \xrightarrow{P} (y_1, \dots, y_k)$ then by definition of the semantics $\mathfrak{h}' = \{\mathfrak{s}(x), \mathfrak{s}(y_1), \dots, \mathfrak{s}(y_n), \mathfrak{s}(p)\}$, hence, by definition of \leq , \mathfrak{h} is necessarily of the form $\{\mathfrak{s}(x), \mathfrak{s}(y_1), \dots, \mathfrak{s}(y_n), \pi\}$, where either $\pi = \mathfrak{s}(p)$ or $\pi = (\mathfrak{s}(p) \oplus_{\mathfrak{P}} \pi')$, for some $\pi' \in \mathcal{P}_{\mathfrak{P}}$. In both cases, we get $dom(\mathfrak{h}') = \{\mathfrak{s}(x)\}$, moreover, by Definition 24, $\{\mathfrak{s}(x)\} = cut(\mathfrak{s}(roots(\phi)), \mathfrak{s}(V), \mathfrak{h})$ since $\{x\} = roots(\phi)$ and for all $i \in \{1, \dots, n\}$ such that $\mathfrak{s}(y_i) \neq \mathfrak{s}(x)$, we have $y_i \in fv(\phi) \setminus \{x\}$, hence $y_i \in V$.
- If ϕ is of the form $\phi_1 \circ \phi_2$, then there exist heaps $\mathfrak{h}'_1, \mathfrak{h}'_2$ such that $dom(\mathfrak{h}'_1) \cap dom(\mathfrak{h}'_2) = \emptyset$, $\mathfrak{h}'_1 \sqcup \mathfrak{h}'_2 = \mathfrak{h}'$ and $(\mathfrak{s}, \mathfrak{h}'_i) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$, for all $i = 1, 2$. Let $V_i = V \cup roots(\phi_{3-i})$. Note that $\mathfrak{s}(V_i) \cap dom(\mathfrak{h}'_i) = \emptyset$ (indeed $dom(\mathfrak{h}'_i) \subseteq dom(\mathfrak{h}')$, hence $\mathfrak{s}(V) \cap dom(\mathfrak{h}'_i) = \emptyset$, and by Proposition 8 $\mathfrak{s}(roots(\phi_{3-i})) \subseteq dom(\mathfrak{h}'_{3-i})$, thus $\mathfrak{s}(roots(\phi_{3-i})) \cap dom(\mathfrak{h}'_i) = \emptyset$, as $dom(\mathfrak{h}'_1) \cap dom(\mathfrak{h}'_2) = \emptyset$). Furthermore, $fv(\phi_i) \subseteq fv(\phi) \subseteq V \cup roots(\phi) = V \cup roots(\phi_i) \cup roots(\phi_{3-i}) = V_i \cup roots(\phi_i)$. By definition of \leq , we have $\mathfrak{h}'_i \leq \mathfrak{h}$ for all $i = 1, 2$, thus, by the induction hypothesis, $dom(\mathfrak{h}'_i) = cut(\mathfrak{s}(roots(\phi_i)), \mathfrak{s}(V_i), \mathfrak{h})$. Consequently $dom(\mathfrak{h}') = dom(\mathfrak{h}'_1) \cup dom(\mathfrak{h}'_2) = cut(\mathfrak{s}(roots(\phi_1)), \mathfrak{s}(V_1), \mathfrak{h}) \cup cut(\mathfrak{s}(roots(\phi_2)), \mathfrak{s}(V_2), \mathfrak{h})$, and by Proposition 35 (applied with $L_i = L'_{3-i} = roots(\mathfrak{h}'_i)$ and $L' = V$), $dom(\mathfrak{h}') = cut(\mathfrak{s}(roots(\phi_1)) \cup \mathfrak{s}(roots(\phi_2)), \mathfrak{s}(V), \mathfrak{h})$. Since we have $\mathfrak{s}(roots(\phi_1)) \cup \mathfrak{s}(roots(\phi_2)) = \mathfrak{s}(roots(\phi))$, we get the result.
- If ϕ is a predicate atom $P(x, \mathbf{y}, z)$, then $roots(\phi) = \{x\}$ and $\phi \leftarrow_{\mathcal{R}} \phi'$, with $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi'$. Since \mathcal{R} is \mathfrak{h} -regular, ϕ' is of the form $\exists u_1, \dots, u_n (x \xrightarrow{P} (v_1, \dots, v_k) \circ \psi \circ \psi')$, where ψ' is pure, $roots(\psi) = \{u_1, \dots, u_n\}$ and $\{u_1, \dots, u_n\} \subseteq \{v_1, \dots, v_k\}$. Consequently, there exists a store \mathfrak{s}' coinciding with \mathfrak{s} on all variables not occurrence in $\{u_1, \dots, u_n\}$ such that $(\mathfrak{s}', \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} x \xrightarrow{P} (v_1, \dots, v_k) \circ \psi \circ \psi'$. Since ψ' is pure, we also have $(\mathfrak{s}', \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} x \xrightarrow{P} (v_1, \dots, v_k) \circ \psi$. As $fv(x \xrightarrow{P} (v_1, \dots, v_k) \circ \psi) \subseteq fv(\phi) \cup \{u_1, \dots, u_n\} \subseteq V \cup roots(x \xrightarrow{P} (v_1, \dots, v_k) \circ \psi)$, we get by the induction hypothesis:

$$\begin{aligned} dom(\mathfrak{h}') &= cut(\mathfrak{s}'(roots(x \xrightarrow{P} (v_1, \dots, v_k) \circ \psi)), \mathfrak{s}'(V), \mathfrak{h}) \\ \text{i.e. } dom(\mathfrak{h}') &= cut(\{\mathfrak{s}(x)\} \cup \mathfrak{s}'(\{u_1, \dots, u_n\}), \mathfrak{s}'(V), \mathfrak{h}) \end{aligned}$$

We assume by α -renaming that $\{u_1, \dots, u_n\} \cap V = \emptyset$ so that $\mathfrak{s}'(V) = \mathfrak{s}(V)$. By definition of the semantics we must have $\mathfrak{h}'(\mathfrak{s}(x)) = (\mathfrak{s}(v_1), \dots, \mathfrak{s}(v_k), \mathfrak{s}(p))$. Since $\mathfrak{h}' \leq \mathfrak{h}$, we get $\mathfrak{h}(\mathfrak{s}(x)) = (\mathfrak{s}'(v_1), \dots, \mathfrak{s}'(v_k), \pi)$ for some $\pi \in \mathcal{P}_{\mathfrak{P}}$. By Proposition 8, necessarily $\{\mathfrak{s}'(u_1), \dots, \mathfrak{s}'(u_n)\} \subseteq dom(\mathfrak{h}')$, since $roots(\psi) = \{u_1, \dots, u_n\}$. As $\mathfrak{s}(V) \cap dom(\mathfrak{h}') = \emptyset$, we get $\{\mathfrak{s}'(u_1), \dots, \mathfrak{s}'(u_n)\} \cap \mathfrak{s}(V) = \emptyset$. Since $\mathfrak{h}(\mathfrak{s}(x)) = (\mathfrak{s}'(v_1), \dots, \mathfrak{s}'(v_k), \pi)$ and $\{u_1, \dots, u_n\} \subseteq \{v_1, \dots, v_k\}$, this

entails, by Definition 24, that $\{\mathfrak{s}'(u_1), \dots, \mathfrak{s}'(u_n)\} \subseteq \text{cut}(\mathfrak{s}(\{x\}), \mathfrak{s}(V), \mathfrak{h})$. Thus $\text{cut}(\{\mathfrak{s}(x)\} \cup \mathfrak{s}'(\{x, u_1, \dots, u_n\}), \mathfrak{s}(V), \mathfrak{h}) = \text{cut}(\mathfrak{s}(\{x\}), \mathfrak{s}(V), \mathfrak{h})$, hence $\text{dom}(\mathfrak{h}') = \text{cut}(\mathfrak{s}(\text{roots}(\phi)), \mathfrak{s}(V), \mathfrak{h})$.

H Proof of Lemma 26

Let λ be a bijective mapping from \mathcal{L} to \mathcal{L} . For any tuple $(\ell_1, \dots, \ell_n, \pi)$ with $\ell_i \in \mathcal{L}$ and $\pi \in \mathcal{P}_{\mathfrak{P}}$, $\lambda((\ell_1, \dots, \ell_n, \pi))$ denotes the tuple $(\lambda(\ell_1), \dots, \lambda(\ell_n), \pi)$. For every heap \mathfrak{h} we denote by $\lambda(\mathfrak{h})$ the heap defined as follows: $\text{dom}(\lambda(\mathfrak{h})) \stackrel{\text{def}}{=} \lambda(\text{dom}(\mathfrak{h}))$ and for all $\ell \in \text{dom}(\mathfrak{h})$, $\lambda(\mathfrak{h})(\ell) \stackrel{\text{def}}{=} \lambda(\mathfrak{h}(\ell))$. The following proposition, showing that the truth value of a formula in a structure does not depend on the name of locations, can be established by an immediate induction on formulas:

Proposition 36. *Let λ be a bijective mapping from \mathcal{L} to \mathcal{L} . If $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$ then $(\lambda \circ \mathfrak{s}, \lambda(\mathfrak{h})) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$.*

Now, assume that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$. Then there exist heaps $\mathfrak{h}_i, \hat{\mathfrak{h}}_i, \mathfrak{h}'$ and $\hat{\mathfrak{h}}'$ (for $i \in \{1, \dots, n\}$) such that $\mathfrak{h} = \mathfrak{h}' \sqcup \bigsqcup_{i=1}^n (\mathfrak{h}_i \sqcup \hat{\mathfrak{h}}_i) \sqcup \hat{\mathfrak{h}}'$, $(\mathfrak{s}, \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} \phi'$, $(\mathfrak{s}, \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$, $(\mathfrak{s}, \hat{\mathfrak{h}}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \psi_i$, $(\mathfrak{s}, \hat{\mathfrak{h}}') \models_{\mathcal{R}}^{\mathfrak{P}} \psi'$, $\text{dom}(\mathfrak{h}_i) \cap \text{dom}(\hat{\mathfrak{h}}_i) = \emptyset$ (for all $i \in \{1, \dots, n\}$) and $\text{dom}(\mathfrak{h}') \cap \text{dom}(\bigsqcup_{i=1}^n (\mathfrak{h}_i \sqcup \hat{\mathfrak{h}}_i) \sqcup \hat{\mathfrak{h}}') = \emptyset$. Let $V' = \text{fv}(\phi) \setminus V$. By the hypothesis of the lemma $\text{roots}(\phi_i) = V$ (for all $i \in \{1, \dots, n\}$), thus $\text{fv}(\phi_i) \subseteq V' \cup \text{roots}(\phi_i)$. Since ϕ is normalized and $(\mathfrak{s}, \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$, we have $\forall x \in \text{fv}(\phi_i) : \mathfrak{s}(x) \in \text{dom}(\mathfrak{h}_i) \implies x \in \text{roots}(\phi_i)$, hence $\mathfrak{s}(V') \cap \text{dom}(\mathfrak{h}_i) = \emptyset$. Moreover, $\mathfrak{h}_i \leq \mathfrak{h}$, and by Lemma 25 (applied on ϕ_i , $\mathfrak{s}(V')$ and \mathfrak{h}_i), we deduce that $\text{dom}(\mathfrak{h}_i) = \text{cut}(\mathfrak{s}(\text{roots}(\phi_i)), \mathfrak{s}(V'), \mathfrak{h}) = \text{cut}(\mathfrak{s}(V), \mathfrak{s}(V'), \mathfrak{h})$. Therefore, $\text{dom}(\mathfrak{h}_i) = \text{dom}(\mathfrak{h}_j)$, for all $i, j \in \{1, \dots, n\}$, which entails that $\text{dom}(\mathfrak{h}_i) \cap \text{dom}(\hat{\mathfrak{h}}_j) = \emptyset$ (as $\text{dom}(\mathfrak{h}_j) \cap \text{dom}(\hat{\mathfrak{h}}_j) = \emptyset$). Let L be any infinite subset of \mathcal{L} containing no locations occurring in $\mathfrak{s}(\text{fv}(\phi))$ or $\text{loc}(\mathfrak{h})$ (such a set always exists as $\text{fv}(\phi)$ and \mathfrak{h} are both finite and \mathcal{L} is infinite). Let λ be any injective mapping from \mathcal{L} to $L \cup \mathfrak{s}(\text{fv}(\phi))$ such that $\lambda(\mathfrak{s}(x)) = \mathfrak{s}(x)$ for all $x \in \text{fv}(\phi)$. By Proposition 36, we have $(\lambda \circ \mathfrak{s}, \lambda(\mathfrak{h}_i)) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$, so that $(\mathfrak{s}, \bigsqcup_{i=1}^n \lambda(\mathfrak{h}_i)) \models_{\mathcal{R}}^{\mathfrak{P}} \star_{i=1}^n \phi_i$ (as $\lambda \circ \mathfrak{s}$ coincides with \mathfrak{s} on all variables in $\text{fv}(\phi)$ and $\bigsqcup_{i=1}^n \lambda(\mathfrak{h}_i)$ is defined, as $\bigsqcup_{i=1}^n \mathfrak{h}_i$ is defined).

We show that $\bigsqcup_{i=1}^n \lambda(\mathfrak{h}_i)$ and \mathfrak{h}' are disjoint. Assume, for the sake of contradiction, that $\ell \in \text{dom}(\lambda(\mathfrak{h}_i)) \cap \text{dom}(\mathfrak{h}')$. This entails that $\ell \in (L \cup \mathfrak{s}(\text{fv}(\phi))) \cap \text{dom}(\mathfrak{h}') \subseteq (L \cup \mathfrak{s}(\text{fv}(\phi))) \cap \text{loc}(\mathfrak{h})$. Since by definition of L , $L \cap \text{loc}(\mathfrak{h}) = \emptyset$ we deduce that $\ell = \mathfrak{s}(x)$, for some $x \in \text{fv}(\phi)$. By definition of λ , this entails that $\lambda(\ell) = \ell$, i.e., $\lambda^{-1}(\ell) = \ell$, so that $\ell \in \text{dom}(\mathfrak{h}_i)$. This entails that $\ell \in \text{dom}(\mathfrak{h}_i) \cap \text{dom}(\mathfrak{h}')$, which contradicts the fact that \mathfrak{h}_i and \mathfrak{h}' are disjoint.

Therefore, $\mathfrak{h}' \sqcup \bigsqcup_{i=1}^n \lambda(\mathfrak{h}_i)$ is defined and we get $(\mathfrak{s}, \mathfrak{h}' \sqcup \bigsqcup_{i=1}^n \lambda(\mathfrak{h}_i)) \models_{\mathcal{R}}^{\mathfrak{P}} \phi' \circ \star_{i=1}^n \phi_i$. As $(\mathfrak{s}, \hat{\mathfrak{h}}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \psi_i$ (for all $i \in \{1, \dots, n\}$) and $(\mathfrak{s}, \hat{\mathfrak{h}}') \models_{\mathcal{R}}^{\mathfrak{P}} \psi'$, we also have $(\mathfrak{s}, (\bigsqcup_{i=1}^n \hat{\mathfrak{h}}_i) \sqcup \hat{\mathfrak{h}}') \models_{\mathcal{R}}^{\mathfrak{P}} \star_{i=1}^n \psi_i * \psi'$, thus it only remains to prove that $\mathfrak{h}' \sqcup \bigsqcup_{i=1}^n \lambda(\mathfrak{h}_i)$ and $(\bigsqcup_{i=1}^n \hat{\mathfrak{h}}_i) \sqcup \hat{\mathfrak{h}}'$ are disjoint to prove that $(\phi' \circ \star_{i=1}^n \phi_i) \circ ((\star_{i=1}^n \psi_i) * \psi')$ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable. We know that $\text{dom}(\mathfrak{h}') \cap (\bigcup_{i=1}^n \text{dom}(\mathfrak{h}_i) \cup \text{dom}(\hat{\mathfrak{h}}')) = \emptyset$,

thus it suffices to show that $dom(\lambda(\mathfrak{h}_i)) \cap dom(\hat{\mathfrak{h}}_j) = \emptyset$ (for all $i, j \in \{1, \dots, n\}$) and that $dom(\lambda(\mathfrak{h}_i)) \cap dom(\hat{\mathfrak{h}}') = \emptyset$ (for all $i \in \{1, \dots, n\}$).

- Let $\ell \in dom(\lambda(\mathfrak{h}_i)) \cap dom(\hat{\mathfrak{h}}_j)$. By definition $\ell \in L \cup \mathfrak{s}(fv(\phi))$ and $\ell \in loc(\mathfrak{h})$, thus $\ell \in \mathfrak{s}(fv(\phi))$. By definition of λ , this entails that $\lambda^{-1}(\ell) = \ell$, so that $\ell \in dom(\mathfrak{h}_i)$, which contradicts the fact that \mathfrak{h}_i and $\hat{\mathfrak{h}}_j$ are disjoint, as previously shown.
- Let $\ell \in dom(\lambda(\mathfrak{h}_i)) \cap dom(\hat{\mathfrak{h}}')$. By definition $\ell \in L \cup \mathfrak{s}(fv(\phi))$ and $\ell \in loc(\mathfrak{h})$, thus there exists $x \in fv(\phi)$ such that $\ell = \mathfrak{s}(x)$. As ϕ is normalized, $(\mathfrak{s}, \hat{\mathfrak{h}}') \models_{\mathcal{R}}^{\mathfrak{P}} \psi'$ and $\ell \in dom(\hat{\mathfrak{h}}')$ we get $x \in roots(\psi')$. Similarly, since $(\mathfrak{s}, \lambda(\mathfrak{h}_i)) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$ and $\ell \in dom(\lambda(\mathfrak{h}_i))$ we get $x \in roots(\phi_i)$, thus $roots(\phi_i) \cap roots(\psi') \neq \emptyset$, which contradicts the hypothesis of the lemma.

Conversely, assume that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} (\phi' \circ \star_{i=1}^n \phi_i) \circ ((\star_{i=1}^n \psi_i) * \psi')$. Then there exist heaps $\mathfrak{h}_i, \hat{\mathfrak{h}}_i, \mathfrak{h}'$ and $\hat{\mathfrak{h}}'$ (for $i \in \{1, \dots, n\}$) such that $\mathfrak{h} = \mathfrak{h}' \sqcup \bigsqcup_{i=1}^n \mathfrak{h}_i \sqcup \bigsqcup_{i=1}^n \hat{\mathfrak{h}}_i \sqcup \hat{\mathfrak{h}}'$, $(\mathfrak{s}, \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} \phi'$, $(\mathfrak{s}, \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$, $(\mathfrak{s}, \hat{\mathfrak{h}}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \psi_i$, $(\mathfrak{s}, \hat{\mathfrak{h}}') \models_{\mathcal{R}}^{\mathfrak{P}} \psi'$, $dom(\mathfrak{h}') \cap dom(\bigsqcup_{i=1}^n \mathfrak{h}_i) = \emptyset$ and $dom(\mathfrak{h}' \sqcup \bigsqcup_{i=1}^n \mathfrak{h}_i) \cap dom(\bigsqcup_{i=1}^n \hat{\mathfrak{h}}_i \sqcup \hat{\mathfrak{h}}') = \emptyset$. This entails that $dom(\mathfrak{h}_i) \cap dom(\hat{\mathfrak{h}}_i) = \emptyset$ (for all $i \in \{1, \dots, n\}$), so that $(\mathfrak{s}, \mathfrak{h}_i \sqcup \hat{\mathfrak{h}}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i \circ \psi_i$. Thus $(\mathfrak{s}, \bigsqcup_{i=1}^n (\mathfrak{h}_i \sqcup \hat{\mathfrak{h}}_i)) \models_{\mathcal{R}}^{\mathfrak{P}} \star_{i=1}^n (\phi_i \circ \psi_i)$ and $(\mathfrak{s}, \bigsqcup_{i=1}^n (\mathfrak{h}_i \sqcup \hat{\mathfrak{h}}_i) \sqcup \hat{\mathfrak{h}}') \models_{\mathcal{R}}^{\mathfrak{P}} \star_{i=1}^n (\phi_i \circ \psi_i) * \psi'$. Moreover, we also have $dom(\mathfrak{h}') \cap (\bigcup_{i=1}^n (dom(\mathfrak{h}_i) \cup dom(\hat{\mathfrak{h}}_i)) \cup dom(\hat{\mathfrak{h}}')) = \emptyset$, hence $(\mathfrak{s}, \hat{\mathfrak{h}} \sqcup \bigsqcup_{i=1}^n (\mathfrak{h}_i \sqcup \hat{\mathfrak{h}}_i) \sqcup \hat{\mathfrak{h}}') \models_{\mathcal{R}}^{\mathfrak{P}} \phi' \circ \star_{i=1}^n (\phi_i \circ \psi_i) * \psi'$, i.e., $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$.

I Proof of Lemma 27

We prove the two implications separately by induction on $|\mathfrak{h}|$. For the first implication, we show for technical convenience that the entailment is valid also in the case where the hypothesis $\mathfrak{s}(x) \neq \mathfrak{s}(y) \implies \mathfrak{s}(y) \notin dom(\mathfrak{h})$ is not satisfied.

\implies Assume that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} P(x, \mathbf{y}, \mathbf{p}) \nabla P'(\mathbf{y}', \mathbf{p}')$. Then there exists a formula ξ such that $P(x, \mathbf{y}, \mathbf{p}) \nabla P'(\mathbf{y}', \mathbf{p}') \leftarrow_{\mathcal{R}} \xi$ and $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \xi$. By definition of the rules defining ∇ , ξ is of the form:

$$\exists \mathbf{u} x \xrightarrow{q} (\mathbf{v}) \circ \bigcirc_{i=1}^n (Q_i(u_i, \mathbf{y}_i, \mathbf{q}_i) \nabla Q'_i(u_i, \mathbf{y}'_i, \mathbf{q}'_i)) \circ \phi \circ \phi' \circ \psi$$

with $q = p \oplus q$, $\mathbf{u} = (u_1, \dots, u_n)$, $\mathbf{v} = (v_1, \dots, v_k)$, ϕ and ϕ' are pure, $\psi = \mathbf{v} \simeq \mathbf{v}'$ and:

$$\begin{aligned} P(x, \mathbf{y}, \mathbf{p}) &\leftarrow_{\mathcal{R}} \exists \mathbf{u} \quad x \xrightarrow{p} (\mathbf{v}) \circ \bigcirc_{i=1}^n Q_i(u_i, \mathbf{y}_i, \mathbf{q}_i) \circ \phi \\ P(x, \mathbf{y}', \mathbf{p}') &\leftarrow_{\mathcal{R}} \exists \mathbf{u} \quad x \xrightarrow{p'} (\mathbf{v}) \circ \bigcirc_{i=1}^n Q'_i(u_i, \mathbf{y}'_i, \mathbf{q}'_i) \circ \phi' \end{aligned}$$

By definition of the semantics, there exist locations ℓ and disjoint heaps $\mathfrak{h}_0, \dots, \mathfrak{h}_n$ such that $(\mathfrak{s}', \mathfrak{h}_0) \models_{\mathcal{R}}^{\mathfrak{P}} x \xrightarrow{q} (\mathbf{v})$ and $(\mathfrak{s}', \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} Q_i(u_i, \mathbf{y}_i, \mathbf{q}_i) \nabla Q'_i(u_i, \mathbf{y}'_i, \mathbf{q}'_i)$, for all $i \in \{1, \dots, n\}$. This entails that $|\mathfrak{h}_0| = 1$, so that $|\mathfrak{h}_i| < |\mathfrak{h}|$ (for all $i \in \{1, \dots, n\}$) and by the induction hypothesis, we deduce that there exist heaps \mathfrak{h}'_i and \mathfrak{h}''_i such that: $dom(\mathfrak{h}'_i) \subseteq dom(\mathfrak{h}_i)$, $dom(\mathfrak{h}''_i) \subseteq dom(\mathfrak{h}_i)$,

$\mathfrak{h}'_i \sqcup \mathfrak{h}''_i = \mathfrak{h}_i$, $(\mathfrak{s}', \mathfrak{h}'_i) \models_{\mathcal{R}}^{\mathfrak{P}} Q_i(u_i, \mathbf{y}_i, \mathbf{q}_i)$ and $(\mathfrak{s}', \mathfrak{h}''_i) \models_{\mathcal{R}}^{\mathfrak{P}} Q_i(u_i, \mathbf{y}'_i, \mathbf{q}'_i)$, with $\mathfrak{s}' \stackrel{\text{def}}{=} \mathfrak{s}\{\mathbf{u} \leftarrow \ell\}$. As $\mathfrak{s}'(q)$ is defined, necessarily $\mathfrak{s}'(p)$ and $\mathfrak{s}'(p')$ are both defined and $\mathfrak{s}'(q) = \mathfrak{s}'(p) \oplus_{\mathfrak{P}} \mathfrak{s}'(p')$. Consider the heaps $\mathfrak{h}'_0 \stackrel{\text{def}}{=} \{(\mathfrak{s}(x), \mathfrak{s}'(\mathbf{v}), \mathfrak{s}(p))\}$ and $\mathfrak{h}''_0 \stackrel{\text{def}}{=} \{(\mathfrak{s}(x), \mathfrak{s}'(\mathbf{v}), \mathfrak{s}(p'))\}$. By definition, we have $\text{dom}(\mathfrak{h}'_0) = \text{dom}(\mathfrak{h}''_0) = \text{dom}(\mathfrak{h}_0)$ and $\mathfrak{h}'_0 \sqcup \mathfrak{h}''_0 = \mathfrak{h}_0$ (since $\mathfrak{s}'(\mathbf{v}) = \mathfrak{s}'(\mathbf{v}')$ as $\mathfrak{s}' \models_{\mathcal{R}}^{\mathfrak{P}} \psi$). It is clear that the heaps $\mathfrak{h}'_0, \dots, \mathfrak{h}'_n$ are pairwise disjoint (as $\mathfrak{h}_0, \dots, \mathfrak{h}_n$ are disjoint) thus $\mathfrak{h}' \stackrel{\text{def}}{=} \mathfrak{h}'_0 \sqcup \dots \sqcup \mathfrak{h}'_n$ is defined. Moreover, as $\mathfrak{s}' \models_{\mathcal{R}}^{\mathfrak{P}} \phi$, we get $(\mathfrak{s}', \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} x \xrightarrow{p} (\mathbf{v}) \circ \bigcirc_{i=1}^n Q_i(u_i, \mathbf{y}_i, \mathbf{q}_i) \circ \phi$, hence $(\mathfrak{s}, \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} P(x, \mathbf{y}, \mathbf{p})$. By symmetry, we also have $(\mathfrak{s}', \mathfrak{h}'') \models_{\mathcal{R}}^{\mathfrak{P}} P'(x, \mathbf{y}', \mathbf{p}')$ with $\mathfrak{h}'' \stackrel{\text{def}}{=} \mathfrak{h}''_0 \sqcup \dots \sqcup \mathfrak{h}''_n$. Furthermore, $\mathfrak{h}' \sqcup \mathfrak{h}'' = (\mathfrak{h}'_0 \sqcup \dots \sqcup \mathfrak{h}'_n) \sqcup (\mathfrak{h}''_0 \sqcup \dots \sqcup \mathfrak{h}''_n) = (\mathfrak{h}'_0 \sqcup \mathfrak{h}''_0) \sqcup \dots \sqcup (\mathfrak{h}'_n \sqcup \mathfrak{h}''_n) = (\mathfrak{h}_0 \sqcup \dots \sqcup \mathfrak{h}_n) = \mathfrak{h}$. Thus $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} P(x, \mathbf{y}, \mathbf{p}) * P'(x, \mathbf{y}', \mathbf{p}')$.

\Leftarrow Let $\mathfrak{h}', \mathfrak{h}''$ be heaps such that $\mathfrak{h}' \sqcup \mathfrak{h}'' = \mathfrak{h}$, $(\mathfrak{s}, \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} P(x, \mathbf{y}, \mathbf{p})$ and $(\mathfrak{s}, \mathfrak{h}'') \models_{\mathcal{R}}^{\mathfrak{P}} P'(x, \mathbf{y}', \mathbf{p}')$. By definition, we have $P(x, \mathbf{y}, \mathbf{p}) \leftarrow_{\mathcal{R}} \xi$ and $P(x, \mathbf{y}', \mathbf{p}') \leftarrow_{\mathcal{R}} \xi'$, with $(\mathfrak{s}, \mathfrak{h}') \models_{\mathcal{R}}^{\mathfrak{P}} \xi$ and $(\mathfrak{s}, \mathfrak{h}'') \models_{\mathcal{R}}^{\mathfrak{P}} \xi'$. Since \mathcal{R} is \mathfrak{h} -regular, ξ and ξ' are of the form $\exists \mathbf{u} (x \xrightarrow{p} (\mathbf{v}) \circ \bigcirc_{i=1}^n Q_i(u_i, \mathbf{y}_i, \mathbf{q}_i) \circ \phi)$ and $\exists \mathbf{u}' (x \xrightarrow{p'} (\mathbf{v}') \circ \bigcirc_{i=1}^{n'} Q_i(u'_i, \mathbf{y}_i, \mathbf{q}'_i) \circ \phi')$, respectively, with $\mathbf{u} \subseteq \mathbf{v}$, $\mathbf{u}' \subseteq \mathbf{v}'$, $\mathbf{u} = (u_1, \dots, u_n)$, $\mathbf{u}' = (u'_1, \dots, u'_{n'})$, $\mathbf{v} = (v_1, \dots, v_k)$, $\mathbf{v}' = (v'_1, \dots, v'_{k'})$, and ϕ, ϕ' are pure. Then there exist stores $\mathfrak{s}', \mathfrak{s}''$, coinciding with \mathfrak{s} on all variables not occurring in \mathbf{u}, \mathbf{u}' , respectively, as well as sequences of pairwise disjoint heaps $\mathfrak{h}'_0, \dots, \mathfrak{h}'_n$ and $\mathfrak{h}''_0, \dots, \mathfrak{h}''_{n'}$ such that $(\mathfrak{s}', \mathfrak{h}'_0) \models_{\mathcal{R}}^{\mathfrak{P}} x \xrightarrow{p} (\mathbf{v})$, $(\mathfrak{s}'', \mathfrak{h}''_0) \models_{\mathcal{R}}^{\mathfrak{P}} x \xrightarrow{p'} (\mathbf{v}')$, $(\mathfrak{s}', \mathfrak{h}'_i) \models_{\mathcal{R}}^{\mathfrak{P}} Q_i(u_i, \mathbf{y}_i, \mathbf{q}_i)$ (for all $i \in \{1, \dots, n\}$), $(\mathfrak{s}'', \mathfrak{h}''_i) \models_{\mathcal{R}}^{\mathfrak{P}} Q'_i(u'_i, \mathbf{y}'_i, \mathbf{q}'_i)$ (for all $i \in \{1, \dots, n'\}$), $\mathfrak{s}' \models_{\mathcal{R}}^{\mathfrak{P}} \phi$ and $\mathfrak{s}'' \models_{\mathcal{R}}^{\mathfrak{P}} \phi'$. Thus $(\mathfrak{s}(x), \mathfrak{s}'(\mathbf{v}), \mathfrak{s}(p)) \in \mathfrak{h}'$ and $(\mathfrak{s}(x), \mathfrak{s}''(\mathbf{v}'), \mathfrak{s}(p')) \in \mathfrak{h}''$. As $\mathfrak{h}' \sqcup \mathfrak{h}'' = \mathfrak{h}$, we deduce that $k = k'$, $\mathfrak{s}'(\mathbf{v}) = \mathfrak{s}''(\mathbf{v}')$ and $\mathfrak{s}(p) \oplus \mathfrak{s}(p')$ is defined, so that $\mathfrak{h}_0 \stackrel{\text{def}}{=} \{(\mathfrak{s}(x), \mathfrak{s}'(\mathbf{v}), \mathfrak{s}(p) \oplus \mathfrak{s}(p'))\} = \mathfrak{h}'_0 \sqcup \mathfrak{h}''_0 \leq \mathfrak{h}$.

Now assume (for the sake of contradiction) that there exists $i \in \{1, \dots, n\}$ such that $v_i \notin \mathbf{u}$ and $v'_i \in \mathbf{u}'$. Then $v'_i \in \text{roots}(\xi')$, we get $\mathfrak{s}''(v'_i) \in \text{dom}(\mathfrak{h}'')$. As $\mathfrak{s}'(\mathbf{v}) = \mathfrak{s}''(\mathbf{v}')$ and $\text{dom}(\mathfrak{h}'') \subseteq \text{dom}(\mathfrak{h})$ we deduce that $\mathfrak{s}'(v_i) \in \text{dom}(\mathfrak{h})$ hence $\mathfrak{s}(v_i) \in \text{dom}(\mathfrak{h})$ (since $v_i \notin \mathbf{u}$), thus $\mathfrak{s}(v_i) = \mathfrak{s}(x)$, by the hypothesis of the lemma. But then $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h}'_0) \cap \text{dom}(\mathfrak{h}''_i)$, which contradicts the fact that \mathfrak{h}'_0 and \mathfrak{h}''_i are disjoint. Similarly, there is no $i \in \{1, \dots, n'\}$ such that $v'_i \notin \mathbf{u}'$ and $v_i \in \mathbf{u}$. This entails that $\{i \in \{1, \dots, k\} \mid v_i \in \mathbf{u}\} = \{i \in \{1, \dots, k\} \mid v'_i \in \mathbf{u}'\}$ and, as $\mathbf{u} \subseteq \mathbf{v}$ and $\mathbf{u}' \subseteq \mathbf{v}'$, we deduce that $n = n'$ (since $\mathfrak{s}'(\mathbf{v}) = \mathfrak{s}''(\mathbf{v}')$ and \mathfrak{s}' and \mathfrak{s}'' are injective on u_1, \dots, u_n and $u'_1, \dots, u'_{n'}$, respectively). By α -renaming, we may then assume that $\mathbf{u} = \mathbf{u}'$ and $\mathfrak{s}' = \mathfrak{s}''$, thus $\mathfrak{s}' \models_{\mathcal{R}}^{\mathfrak{P}} \mathbf{v} \simeq \mathbf{v}'$. We have $\mathfrak{h}' \sqcup \mathfrak{h}'' = \mathfrak{h}$, i.e., $(\mathfrak{h}'_0 \sqcup \dots \sqcup \mathfrak{h}'_n) \sqcup (\mathfrak{h}''_0 \sqcup \dots \sqcup \mathfrak{h}''_n) = \mathfrak{h}$, and $\mathfrak{h} = (\mathfrak{h}'_0 \sqcup \mathfrak{h}''_0) \sqcup \dots \sqcup (\mathfrak{h}'_n \sqcup \mathfrak{h}''_n)$. Let $\mathfrak{h}_i \stackrel{\text{def}}{=} \mathfrak{h}'_i \sqcup \mathfrak{h}''_i$ for all $i \in \{0, \dots, n\}$, so that $\mathfrak{h} = \mathfrak{h}_0 \sqcup \dots \sqcup \mathfrak{h}_n$. Note that $\mathfrak{s}'(u_i) \in \text{dom}(\mathfrak{h}'_i)$ (by Proposition 8), and $\mathfrak{s}(x) \in \text{dom}(\mathfrak{h}'_0)$ which entails, as $\mathfrak{h}'_0, \dots, \mathfrak{h}'_n$ are disjoint and $\text{dom}(\mathfrak{h}'_i) \subseteq \text{dom}(\mathfrak{h})$, that the entailment $\mathfrak{s}'(y) \in \text{dom}(\mathfrak{h}'_i) \implies \mathfrak{s}'(y) = \mathfrak{s}'(u_i)$ holds for all $i \in \{1, \dots, n\}$. By symmetry, $\mathfrak{s}'(y) \in \text{dom}(\mathfrak{h}''_i) \implies \mathfrak{s}'(y) = \mathfrak{s}'(u_i)$ also holds for all $i \in \{1, \dots, n\}$, hence $\mathfrak{s}'(y) \in \text{dom}(\mathfrak{h}_i) \implies \mathfrak{s}'(y) = \mathfrak{s}'(u_i)$. Since $|\mathfrak{h}_0| = 1$, necessarily $|\mathfrak{h}_i| < |\mathfrak{h}|$, for all $i \in \{1, \dots, n\}$. By the induction hypothesis, we deduce that $(\mathfrak{s}', \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} Q_i(u_i, \mathbf{y}_i, \mathbf{p}) \nabla Q'_i(u_i, \mathbf{y}'_i, \mathbf{p}'')$, so that

$(\mathfrak{s}', \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} x \xrightarrow{q} (\mathbf{v}) \circ \bigcirc_{i=1}^n Q_i(u_i, \mathbf{y}_i, \mathbf{p}) \nabla Q'_i(u_i, \mathbf{y}'_i, \mathbf{p}'') \circ \phi \circ \phi' \circ (\mathbf{v} \simeq \mathbf{v}')$. By definition of the rules of $P(x, \mathbf{y}, \mathbf{p}) \nabla P'(x, \mathbf{y}', \mathbf{p}')$, this entails that $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} P(x, \mathbf{y}, \mathbf{p}) \nabla P'(x, \mathbf{y}', \mathbf{p}')$.

J Proof of Proposition 30

The direct implication is immediate to prove. We now establish the converse. If $\rho_{\mathbf{a}}$ is \mathfrak{P} -satisfiable then there exists a store \mathfrak{s} such that $\mathfrak{s} \models^{\mathfrak{P}} \rho_{\mathbf{a}}$ (i.e., $(\mathfrak{s}, \emptyset) \models^{\mathfrak{P}} \rho_{\mathbf{a}}$). Since $\sim_{\mathbf{a}}$ is an equivalence relation on $V_{\mathbf{a}} \cap \mathcal{V}_1$, it is clear that there exists a store $\hat{\mathfrak{s}}$ such that $\hat{\mathfrak{s}}(x) = \mathfrak{s}(x)$ holds for all $x \in \mathcal{V}_p$, and for all $y, z \in V_{\mathbf{a}} \cap \mathcal{V}_1$: $\hat{\mathfrak{s}}(y) = \hat{\mathfrak{s}}(z) \iff y \sim_{\mathbf{a}} z$ (it suffices to associate all equivalence classes of $\sim_{\mathbf{a}}$ with pairwise distinct arbitrarily chosen locations). As $V_{\mathbf{a}}$ is finite, there exists a heap $\hat{\mathfrak{h}}$ such that $\text{dom}(\hat{\mathfrak{h}}) = \hat{\mathfrak{s}}(V_{\mathbf{a}})$ (e.g., $\hat{\mathfrak{h}} = \{(\ell, \pi) \mid \ell \in \text{dom}(\hat{\mathfrak{h}})\}$ where π is some arbitrary permission in $\mathcal{P}_{\mathfrak{P}}$). As the truth value of \mathbf{a} only depends on the interpretation of variables of sort p , we have $\hat{\mathfrak{s}} \models^{\mathfrak{P}} \mathbf{a}$, so that $(\hat{\mathfrak{s}}, \hat{\mathfrak{h}}) \models^{\mathfrak{P}} \mathbf{a}$.

K Proof of Lemma 31

We need to establish a slightly more general property: Let ϕ be a \circ -formula, let $(\mathfrak{s}, \mathfrak{h})$ be a structure and let $V \supseteq \text{fv}(\phi)$. We show that the two following assertions hold: (i) If $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$ then there exists a heap abstraction $(V, \sim, A, \rho) \in \mathfrak{A}(\phi)$ such that $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} (V, \sim, A, \rho)$. (ii) If $(V, \sim, A, \rho) \in \mathfrak{A}(\phi)$, $\sim = \{(u, v) \in (V \cap \mathcal{V}_1)^2 \mid \mathfrak{s}(u) = \mathfrak{s}(v)\}$ and $\mathfrak{s} \models^{\mathfrak{P}} \rho$, then for every store $\hat{\mathfrak{s}}$ such that $\hat{\mathfrak{s}}(x) = \mathfrak{s}(x)$ for all $x \in \mathcal{V}_p$ and $\hat{\mathfrak{s}}(x) = \hat{\mathfrak{s}}(y) \iff \mathfrak{s}(x) = \mathfrak{s}(y)$ for all $x, y \in V \cap \mathcal{V}_1$, and for every infinite subset L of \mathcal{L} , there exists a heap $\hat{\mathfrak{h}}$ such that $(\hat{\mathfrak{s}}, \hat{\mathfrak{h}}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$, with $\text{dom}(\hat{\mathfrak{h}}) \subseteq L \cup \hat{\mathfrak{s}}(A)$.

1. The proof is by induction on the satisfiability relation. Let $\sim = \{(x, y) \in (V \cap \mathcal{V}_1)^2 \mid \mathfrak{s}(x) = \mathfrak{s}(y)\}$.
 - If $\phi = x \xrightarrow{p} (y_1, \dots, y_k)$ then necessarily $\text{dom}(\mathfrak{h}) = \{\mathfrak{s}(x)\}$ and $\mathfrak{s}(p)$ is defined, so that $\mathfrak{s} \models^{\mathfrak{P}} \text{def}(p)$. Moreover $(V, \sim, \{y \in V \mid y \sim x\}, \text{def}(p)) \in \mathfrak{A}(\phi)$, by definition of $\mathfrak{A}(\phi)$. By definition of \sim , we have $\forall x, y \in V_{\mathbf{a}} \cap \mathcal{V}_1$: $x \sim y \iff \mathfrak{s}(x) = \mathfrak{s}(y)$ and $\{y \in V \mid y \sim x\} = \{y \in V \mid \mathfrak{s}(y) = \mathfrak{s}(x)\} = \{y \in V \mid \mathfrak{s}(y) \in \text{dom}(\mathfrak{h})\}$, thus $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} (V, \sim, \{y \in V \mid y \sim x\}, \text{def}(p))$.
 - If $\phi = x \simeq y$ (resp. $x \not\approx y$) with $x, y \in \mathcal{V}_1$, then $\mathfrak{h} = \emptyset$. Moreover $\mathfrak{s}(x) = \mathfrak{s}(y)$ (resp. $\mathfrak{s}(x) \neq \mathfrak{s}(y)$) so that $x \sim y$ (resp. $x \not\sim y$) and $(V, \sim, \emptyset, \text{emp}) \in \mathfrak{A}(\phi)$, by definition of $\mathfrak{A}(\phi)$. As $\sim = \{(x, y) \in (V \cap \mathcal{V}_1)^2 \mid \mathfrak{s}(x) = \mathfrak{s}(y)\}$, $\text{dom}(\mathfrak{h}) = \emptyset$, and $\mathfrak{s} \models^{\mathfrak{P}} \text{emp}$ we have $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} (V, \sim, \emptyset, \text{emp})$.
 - If ϕ is a permission formula then $\mathfrak{h} = \emptyset$, and $(V, \sim, \emptyset, \phi) \in \mathfrak{A}(\phi)$ (by definition of $\mathfrak{A}(\phi)$). As $\sim = \{(x, y) \in (V \cap \mathcal{V}_1)^2 \mid \mathfrak{s}(x) = \mathfrak{s}(y)\}$, $\text{dom}(\mathfrak{h}) = \emptyset$ and $(\mathfrak{s}, \emptyset) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$, we get $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} (V, \sim, \emptyset, \phi)$.
 - Assume that $\phi = \exists x \psi$. By α -renaming, we assume that $x \notin V$. There exists a store \mathfrak{s}' coinciding with \mathfrak{s} on all the variables distinct from x such that $(\mathfrak{s}', \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \psi$. Let $V' = V \cup \{x\}$, so that $\text{fv}(\psi) \subseteq V'$. By

the induction hypothesis, $\mathfrak{A}(\psi)$ contains a tuple (V', \sim', A', ρ) such that $(\mathfrak{s}', \mathfrak{h}) \models^{\mathfrak{P}} (V', \sim', A', \rho)$. Then $\sim' = \{(y, z) \in (V' \cap \mathcal{V}_1) \mid \mathfrak{s}'(y) = \mathfrak{s}'(z)\}$, $A' = \{y \in V' \mid \mathfrak{s}'(y) \in \text{dom}(\mathfrak{h})\}$ and $\mathfrak{s}' \models^{\mathfrak{P}} \rho$. By definition of $\mathfrak{A}(\phi)$, we deduce that the heap abstraction $\mathfrak{a} \stackrel{\text{def}}{=} (V' \setminus \{x\}, \{(y, z) \in V' \mid y \sim' z \wedge y, z \neq x\}, A' \setminus \{x\}, \rho)$ is in $\mathfrak{A}(\psi)$. By definition of V' , we have $V' \setminus \{x\} = V$, $\{(y, z) \in (V' \cap \mathcal{V}_1)^2 \mid y \sim' z \wedge y \neq x \wedge z \neq x\} = \{(y, z) \in (V \cap \mathcal{V}_1)^2 \mid \mathfrak{s}'(y) = \mathfrak{s}'(z)\}$ and $A' \setminus \{x\} = \{y \in V \mid \mathfrak{s}'(y) \in \text{dom}(\mathfrak{h})\}$, so that $\mathfrak{a} = (V, \{(y, z) \in (V \cap \mathcal{V}_1)^2 \mid \mathfrak{s}'(y) = \mathfrak{s}'(z)\}, \{y \in V \mid \mathfrak{s}'(y) \in \text{dom}(\mathfrak{h})\}, \rho)$. As \mathfrak{s}' and \mathfrak{s} coincides on all variables $y \neq x$ (hence on all variables in V), and since ρ does not contain x (since existential variables are of sort 1) we get $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} \mathfrak{a}$.

- Assume that $\phi = \phi_1 \circ \phi_2$. Then there exist disjoint heaps $\mathfrak{h}_1, \mathfrak{h}_2$ with $(\mathfrak{s}, \mathfrak{h}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$ and $\mathfrak{h} = \mathfrak{h}_1 \sqcup \mathfrak{h}_2$. Let $A_i = \{x \in V \mid \mathfrak{s}(x) \in \text{dom}(\mathfrak{h}_i)\}$. We have $\text{fv}(\phi_i) \subseteq \text{fv}(\phi) \subseteq V$, thus, by the induction hypothesis, there exist $(V, \sim_i, A_i, \rho_i) \in \mathfrak{A}(\phi_i)$ such that $(\mathfrak{s}, \mathfrak{h}_i) \models^{\mathfrak{P}} (V, \sim_i, A_i, \rho_i)$, so that $\sim_1 = \sim_2 = \{(x, y) \in (V \cap \mathcal{V}_1)^2 \mid \mathfrak{s}(x) = \mathfrak{s}(y)\}$, $A_i = \{x \in V \mid \mathfrak{s}(x) \in \text{dom}(\mathfrak{h}_i)\}$ and $\mathfrak{s} \models^{\mathfrak{P}} \rho_i$. As $\text{dom}(\mathfrak{h}_1) \cap \text{dom}(\mathfrak{h}_2) = \emptyset$ necessarily $A_1 \cap A_2 = \emptyset$ thus $(V, \sim_1, A_1 \cup A_2, \rho_1 \circ \rho_2) \in \mathfrak{A}(\phi)$. Then $\mathfrak{s} \models^{\mathfrak{P}} \rho_1 \circ \rho_2$, and since $\text{dom}(\mathfrak{h}) = \text{dom}(\mathfrak{h}_1) \cup \text{dom}(\mathfrak{h}_2)$, we get $\{x \mid \mathfrak{s}(x) \in \text{dom}(\mathfrak{h})\} = \{x \mid \mathfrak{s}(x) \in \text{dom}(\mathfrak{h}_1)\} \cup \{x \mid \mathfrak{s}(x) \in \text{dom}(\mathfrak{h}_2)\} = A_1 \cup A_2$, thus $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} (V, \sim_1, A_1 \cup A_2, \rho_1 \circ \rho_2)$.
 - Assume that ϕ is a predicate atom. Then there exists ψ such that $\phi \Leftarrow_{\mathcal{R}} \psi$ and $(\mathfrak{s}, \mathfrak{h}) \models_{\mathcal{R}}^{\mathfrak{P}} \psi$. We have $\text{fv}(\psi) \subseteq \text{fv}(\phi) \subseteq V$. By the induction hypothesis, there exists a heap abstraction $(V, \sim, A, \rho) \in \mathfrak{A}(\psi)$ such that $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} (V, \sim, A, \rho)$. By definition of $\mathfrak{A}(\phi)$, (V, \sim, A, ρ) is in $\mathfrak{A}(\phi)$.
2. The proof is by induction on $\mathfrak{A}(\phi)$.
- Assume that $\phi = x \stackrel{\mathfrak{p}}{=} (y_1, \dots, y_k)$. Since $(V, \sim, A, \rho) \in \mathfrak{A}(\phi)$, we must have $A = \{y \in (V \cap \mathcal{V}_1) \mid y \sim x\}$ (thus $x \in A$) and $\rho = \text{def}(p)$, by definition of $\mathfrak{A}(\phi)$. Since $\mathfrak{s} \models^{\mathfrak{P}} \rho$ necessarily $\hat{\mathfrak{s}} \models^{\mathfrak{P}} \rho$, as $\hat{\mathfrak{s}}$ and \mathfrak{s} coincide on all variables of sort \mathfrak{p} . Consequently, $\hat{\mathfrak{s}}(p)$ must be defined. Consider the heap $\hat{\mathfrak{h}} \stackrel{\text{def}}{=} \{\hat{\mathfrak{s}}(x), \hat{\mathfrak{s}}(y_1), \dots, \hat{\mathfrak{s}}(y_k), \hat{\mathfrak{s}}(p)\}$. By definition of the semantics, $(\hat{\mathfrak{s}}, \hat{\mathfrak{h}}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$, moreover $\text{dom}(\mathfrak{h}) = \{\hat{\mathfrak{s}}(x)\} \subseteq \hat{\mathfrak{s}}(A)$, since $x \in A$.
 - Assume that $\phi = x \simeq y$. As $(V, \sim, A, \rho) \in \mathfrak{A}(\phi)$, we must have $x \sim y$, by definition of $\mathfrak{A}(\phi)$. Since $\sim = \{(u, v) \in (V \cap \mathcal{V}_1)^2 \mid \mathfrak{s}(u) = \mathfrak{s}(v)\}$, this entails that $\mathfrak{s}(x) = \mathfrak{s}(y)$, so that $\hat{\mathfrak{s}}(x) = \hat{\mathfrak{s}}(y)$, and by letting $\hat{\mathfrak{h}} = \emptyset$, we get $(\hat{\mathfrak{s}}, \hat{\mathfrak{h}}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$.
 - The proof is similar if $\phi = x \not\sim y$.
 - Assume that ϕ is a permission formula. In this case, we must have $\rho = \phi$, hence $\mathfrak{s} \models^{\mathfrak{P}} \phi$. Since $\hat{\mathfrak{s}}$ coincides with \mathfrak{s} on all variables of sort \mathfrak{p} , this entails that $\hat{\mathfrak{s}} \models^{\mathfrak{P}} \phi$, so that $(\hat{\mathfrak{s}}, \emptyset) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$.
 - Assume that $\phi = \exists x \psi$. As $(V, \sim, A, \rho) \in \mathfrak{A}(\phi)$, we deduce, by definition of $\mathfrak{A}(\phi)$, that $(V', \sim', A', \rho') \in \mathfrak{A}(\psi)$ with $V = V' \setminus \{x\}$, $\sim = \{(u, v) \mid u \sim' v \wedge u \neq x \wedge v \neq x\}$, $A = A' \setminus \{x\}$ and $\rho = \rho'$. Moreover, since $(\mathfrak{s}, \mathfrak{h}) \models^{\mathfrak{P}} (V, \sim, A, \rho)$, we deduce that $y \sim z \implies \mathfrak{s}(y) = \mathfrak{s}(z)$, for

all $y, z \in (V \cap \mathcal{V}_1)$. By the hypothesis of the lemma, this entails that the implication $\forall y, z \in (V \cap \mathcal{V}_1) (y \sim z \implies \hat{\mathfrak{s}}(y) = \hat{\mathfrak{s}}(z))$ also holds. Consider any stores \mathfrak{s}' and $\hat{\mathfrak{s}}'$ coinciding with \mathfrak{s} and $\hat{\mathfrak{s}}$ (respectively) on all variables distinct from x and such that:

- If $x \sim' y$ for some variable $y \in V$ then $\mathfrak{s}'(x) \stackrel{\text{def}}{=} \mathfrak{s}(y)$ and $\hat{\mathfrak{s}}'(x) \stackrel{\text{def}}{=} \hat{\mathfrak{s}}(y)$ (note that $\mathfrak{s}(y)$ and $\hat{\mathfrak{s}}(y)$ does not depend on the choice of $y \in V$ in the same equivalence class for \sim).

- Otherwise, we let $\mathfrak{s}'(x) = \hat{\mathfrak{s}}'(x) = \ell$, where ℓ is an arbitrarily chosen location in $L \setminus \hat{\mathfrak{s}}(V)$. Note that such a location exists since L is infinite. By construction, it is clear that the equivalence $\mathfrak{s}'(y) = \hat{\mathfrak{s}}'(z) \iff y \sim' z \iff \hat{\mathfrak{s}}'(y) = \hat{\mathfrak{s}}'(z)$ holds for all $y, z \in V' \cap \mathcal{V}_1$. Moreover, \mathfrak{s}' coincides with \mathfrak{s} on all permission variables (as x is of sort 1, since quantification over permission variables is not allowed), thus $\mathfrak{s}' \models^{\mathfrak{P}} \rho' = \rho$. Let $L' = L \setminus \{\hat{\mathfrak{s}}'(x)\}$. By the induction hypothesis, applied on (V', \sim', A', ρ') , ψ , \mathfrak{s}' , $\hat{\mathfrak{s}}'$ and L' , there exists a heap $\hat{\mathfrak{h}}$ such that $(\hat{\mathfrak{s}}', \hat{\mathfrak{h}}) \models_{\mathcal{R}}^{\mathfrak{P}} \psi$ and $\text{dom}(\hat{\mathfrak{h}}) \subseteq L' \cup \hat{\mathfrak{s}}'(A')$. This entails that $(\hat{\mathfrak{s}}, \hat{\mathfrak{h}}) \models_{\mathcal{R}}^{\mathfrak{P}} \exists x \psi$. It only remains to show that $\text{dom}(\hat{\mathfrak{h}}) \subseteq L \cup \hat{\mathfrak{s}}(A)$. Assume, for the sake of contradiction, that $\text{dom}(\hat{\mathfrak{h}})$ contains a location ℓ not occurring in $L \cup \hat{\mathfrak{s}}(A)$. As $\ell \in L' \cup \hat{\mathfrak{s}}'(A')$, with $L' = L \setminus \{\hat{\mathfrak{s}}'(x)\}$ and $A = A' \setminus \{x\}$, this entails that $\ell = \hat{\mathfrak{s}}'(x)$ with $x \in A'$ and $\hat{\mathfrak{s}}'(x) \notin L$. By definition of $\hat{\mathfrak{s}}'$, the latter assertion entails that there exists a variable $y \in \mathcal{V}_1 \cap V$ such that $x \sim y$ and $\hat{\mathfrak{s}}'(x) = \hat{\mathfrak{s}}(y)$. We must have $y \in A'$ (since $x \in A'$ and A' is closed under \sim) thus $y \in A$ (as $x \neq y$, since $x \notin V$), hence $\hat{\mathfrak{s}}(y) = \ell \in \hat{\mathfrak{s}}(A)$, which contradicts our assumption.

- Assume that $\phi = \phi_1 \circ \phi_2$. As $(V, \sim, A, \rho) \in \mathfrak{A}(\phi)$, we deduce, by definition of $\mathfrak{A}(\phi)$, that for all $i = 1, 2$, $\mathfrak{A}(\phi_i)$ contains a heap abstraction (V, \sim, A_i, ρ_i) with $A = A_1 \cup A_2$, $\rho = \rho_1 \circ \rho_2$ and $A_1 \cap A_2 = \emptyset$. As $\mathfrak{s} \models^{\mathfrak{P}} \rho$, we get $\mathfrak{s} \models^{\mathfrak{P}} \rho_i$. Moreover, $\mathfrak{s}(A_1) \cap \mathfrak{s}(A_2) = \emptyset$. Indeed, if $\mathfrak{s}(x_i) \in \mathfrak{s}(A_i)$ with $x_i \in A_i$, then $x_1 \sim x_2$ (as $\sim = \{(u, v) \in (V \cap \mathcal{V}_1)^2 \mid \mathfrak{s}(u) = \mathfrak{s}(v)\}$) which entails (by Definition 28) that $x_2 \in A_1$ and $x_1 \in A_2$, contradicting the fact that $A_1 \cap A_2 = \emptyset$. Let L_1, L_2 be disjoint infinite subsets of L also disjoint from $\hat{\mathfrak{s}}(V)$ (such sets exist since L is infinite). By the induction hypothesis, there exist heaps $\hat{\mathfrak{h}}_i$ such that $(\hat{\mathfrak{s}}, \hat{\mathfrak{h}}_i) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_i$ and $\text{dom}(\hat{\mathfrak{h}}_i) \subseteq L_i \cup \hat{\mathfrak{s}}(A_i)$. As $L_1 \cap L_2 = \emptyset$, $\mathfrak{s}(A_1) \cap \mathfrak{s}(A_2) = \emptyset$ and $(L_1 \cup L_2) \cap \hat{\mathfrak{s}}(V) = \emptyset$, we have $\text{dom}(\hat{\mathfrak{h}}_1) \cap \text{dom}(\hat{\mathfrak{h}}_2) = \emptyset$, so that $\hat{\mathfrak{h}}_1$ and $\hat{\mathfrak{h}}_2$ are disjoint and $(\hat{\mathfrak{s}}, \hat{\mathfrak{h}}_1 \sqcup \hat{\mathfrak{h}}_2) \models_{\mathcal{R}}^{\mathfrak{P}} \phi_1 \circ \phi_2 = \phi$. Moreover, $\text{dom}(\hat{\mathfrak{h}}_1 \sqcup \hat{\mathfrak{h}}_2) \subseteq L \cup \hat{\mathfrak{s}}(A)$.
- Assume that ϕ is a predicate atom. Since $(V, \sim, A, \rho) \in \mathfrak{A}(\phi)$, necessarily $(V, \sim, A, \rho) \in \mathfrak{A}(\xi)$ with $\phi \leftarrow_{\mathcal{R}} \xi$. By the induction hypothesis there exists a heap $\hat{\mathfrak{h}}$ such that $(\hat{\mathfrak{s}}, \hat{\mathfrak{h}}) \models_{\mathcal{R}}^{\mathfrak{P}} \xi$ and $\text{dom}(\hat{\mathfrak{h}}) \subseteq L \cup \hat{\mathfrak{s}}(A)$. By definition of the semantics, we also have $(\hat{\mathfrak{s}}, \hat{\mathfrak{h}}) \models_{\mathcal{R}}^{\mathfrak{P}} \phi$.

L Proof of Theorem 32

We assume for simplicity that ϕ contains no quantifier and no points-to atom¹².

¹² It is clear that this is not restrictive: any existential variable may be replaced by a fresh free variable, and any points-to atom $x \dot{\rightarrow} (y_1, \dots, y_k)$ may be replaced

Step 1 (Normalization). Using Lemma 23, one first compute a set of normalized formulas Φ_2 such that ϕ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable iff there exists $\phi_2 \in \Phi_2$ such that ϕ_2 is $(\mathcal{R}, \mathfrak{P})$ -satisfiable, and all formulas in Φ_2 are separating conjunctions of \circ -formulas (all the existential variables occurring in formulas in Φ_2 are replaced by fresh free variables).

Step 2 (Elimination of $$).* By definition, every formula $\phi_2 \in \Phi_2$ may be written on the form $\phi' \circ (\bigstar_{i=1}^n \chi_i)$ (\dagger), where $\phi', \chi_1, \dots, \chi_n$ are \circ -formulas (with initially $\phi' = \mathbf{emp}$). We show that, if $n > 1$, then the above formula can be reduced into a strictly smaller sat-equivalent formula that is still of the form (\dagger). Let x be any variable occurring in $\mathit{roots}(\chi_i)$, for some $i \in \{1, \dots, n\}$ (if no such x exists then $\bigstar_{i=1}^n \chi_i$ is pure and the symbol $*$ may be replaced by \circ , yielding a formula of the form (\dagger) with $n = 1$). Let I be the set of indices in $\{1, \dots, n\}$ such that $x \in \mathit{roots}(\chi_i)$. By symmetry, we assume that $I = \{1, \dots, m\}$ for some $m \in \{1, \dots, n\}$. Every formula χ_i must be of the form $\delta_i * \chi'_i$, where δ_i is a predicate atom, $\mathit{roots}(\delta_i) = \{x\}$ and $x \notin \mathit{roots}(\chi'_i)$ (if χ_i contains two distinct atoms with the same root x then by Proposition 8, χ_i is $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable and the entire formula can be dismissed). By Lemma 26, applied with $\phi_i = \delta_i$, $\psi_i = \chi'_i$ and $\psi' = \bigstar_{i=m+1}^n \chi_i$, we deduce that $\phi' \circ (\bigstar_{i=1}^n \chi_i)$ is sat-equivalent to $(\phi' \circ \bigstar_{i=1}^m \delta_i) \circ (\bigstar_{i=1}^m \chi'_i * \bigstar_{i=m+1}^n \chi_i)$. Let $\phi'_2 = \phi' \circ (\delta_1 \nabla \dots \nabla \delta_m) \circ (\bigstar_{i=1}^m \chi'_i * \bigstar_{i=m+1}^n \chi_i)$. By Lemma 27, $\bigstar_{i=1}^m \delta_i \equiv_{\mathfrak{P}}^{\mathcal{R}} \delta_1 \nabla \dots \nabla \delta_m$ (the hypothesis of the lemma is satisfied since all formulas are normalized), thus ϕ'_2 and $\phi' \circ (\bigstar_{i=1}^n \chi_i)$ are sat-equivalent. As the weight of the symbol $*$ is strictly greater than that of \circ , it is clear that $|\phi'_2| < |\phi' \circ (\bigstar_{i=1}^n \chi_i)|$. Moreover, ϕ'_2 is also of form (\dagger). By repeating the above transformation, we eventually obtain a formula ϕ_3 of the form (\dagger) above, with $n = 1$, i.e., a \circ -formula. We thus get a set of \circ -formulas Φ_3 such that ϕ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable iff there exists $\phi_3 \in \Phi_3$ such that ϕ_3 is $(\mathcal{R}, \mathfrak{P})$ -satisfiable.

Step 3 (Abstractions). It only remains to check that one of \circ -formulas $\phi_3 \in \Phi_3$ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable. By Lemma 31, it is sufficient to compute the set $\mathfrak{A}(\phi_3)$ and test whether it contains a heap abstraction (V, \sim, A, ρ) that is \mathfrak{P} -satisfiable. By Proposition 30, we only have to check that ρ is \mathfrak{P} -satisfiable, which is decidable by the hypothesis of the lemma.

We now analyze the complexity of the algorithm and we show that it runs in exponential time.

1. At Step 1, one replacement is performed for each atom δ and for each variable. Moreover, each replacement may in turn introduce new atoms and new variables in the formula (however, as shown above, the application of the replacement operation on these new variables do not add further variables). Note that the replacement does not increase the number of \circ -formulas occurring as operands of $*$: only the number of atoms inside \circ -formulas may increase. Thus the number of \circ -formulas is at most $|\phi|$, as

by a predicate atom $P(x, y_1, \dots, y_k, z)$, with the rule $P(u, v_1, \dots, v_k, w) \Leftarrow u \xrightarrow{w} (v_1, \dots, v_k)$.

initially ϕ contains at most $|\phi|$ spatial predicate atoms and no occurrence of \circ . Moreover, the replacement cannot be applied twice with the same variable and the same \circ -formula. We denote by N be the maximal number of variables introduced by one single replacement. We get a total of at most $fv(\phi) + (N \times fv(\phi) \times |\phi|) \leq (N + 1) \times |\phi|^2$ variables. As all the spatial atoms occurring in the same \circ -formula must have distinct roots (otherwise the formula is $(\mathcal{R}, \mathfrak{P})$ -unsatisfiable by Proposition 8 and can be dismissed), this yields a total of at most $(N + 1) \times |\phi|^2$ atoms in each \circ -formula, hence of at most $(N + 1) \times |\phi|^3$ atoms in every formula $\phi_2 \in \Phi_2$. Each replacement may produce a new derived predicate, which may be of the form $\delta[x]^-$ or $(P(x, \mathbf{z}, p) \circ \bigcirc_{i=1}^n P_i(x_i, \mathbf{y}_i, p)) \rightarrow Q(\mathbf{y}, p)$ (with $\delta = Q(\mathbf{y}, p)$). In the former case, no new variable is introduced and the maximal arity increases by at most 1, whereas in the latter case, new variables x_1, \dots, x_n (with $n \leq N$) are added into the formula and the maximal arity of the predicates occurring in the formula is increased by at most n (see Definition 17). By Lemma 20, we have $\{x_1, \dots, x_n\} \subseteq (x, \mathbf{z})|_{\gamma_{\mathcal{R}}(P)}$, thus $n \leq \text{card}(\gamma_{\mathcal{R}}(P))$. By Propositions 14 and 18, the computation of derived predicates cannot increase the maximal value of $\text{card}(\gamma_{\mathcal{R}}(P))$, so that $N \leq |\mathcal{R}|$ and $\text{card}(\gamma_{\mathcal{R}}(P)) \leq |\mathcal{R}|$, for every derived predicate P . Thus, in every formula $\phi_2 \in \Phi_2$, the number of variables added in the formula, the number of atoms and the number of derived predicates and their maximal arity are all polynomial w.r.t. $|\phi| + |\mathcal{R}|$, so that $|\phi_2|$ is polynomial w.r.t. $|\phi|$. Moreover, the total number of derived predicates is at most exponential w.r.t. $|\phi| + |\mathcal{R}|$. Indeed, it is easy to see that these derived predicates may be uniquely determined by choosing the initial atom δ in ϕ from which the derived predicate is computed and the sets of atoms that are removed from the call tree of δ (inside the initial set of predicates). At most $N + 1$ such atoms are introduced at each replacement, and there are at most $(N + 1) \times fv(\phi)$ replacements, yielding at most $2^{|\mathcal{R}| \times (N+1)^2 \times fv(\phi)}$ possible choices. Therefore, the total number of formulas in Φ_2 is at most exponential w.r.t. $|\phi| + |\mathcal{R}|$. It is clear that each formula ϕ_2 can be computed in polynomial time.

2. It is straightforward to check that for each formula $\phi_2 \in \Phi_2$, the corresponding \circ -formula $\phi_3 \in \Phi_3$ can be computed in polynomial time w.r.t. $|\phi_2|$ (each step of the transformation strictly decreases the size of the formula). Hence Step 2 can be performed in exponential time w.r.t. $|\phi|$.
3. At Step 3, for each formula $\phi_3 \in \Phi_3$, the test $\mathfrak{A}(\phi_3)$ may be computed by using a standard fixpoint computation algorithm, following the inductive rules in Definition 28. Observe that we only need to compute $\mathfrak{A}(\psi)$ for formulas ψ that occur in either ϕ_3 or in \mathcal{R} (up to a renaming of variables), as other formulas does not interfere with the computation of $\mathfrak{A}(\phi_3)$. Moreover, we only need to consider heap abstractions (V, \sim, A, ρ) such that V contains only variables in $fv(\psi)$ and existential variables in \mathcal{R} , and ρ contains only permission terms and permission predicates in ψ . It is clear that the number of such heap abstractions is simply exponential w.r.t. $|\phi_3|$, hence the computation can be performed in exponential time.

The EXPTIME-hardness proof goes by an easy reduction from the halting problem for alternating Turing machines (ATM) running in polynomial space. As APSPACE (the class of languages decidable in polynomial space by alternating Turing machines) is identical to EXPTIME, we get the result. Let $M = (Q, \Gamma, \delta, q_0, g)$ be an ATM, where Q is a finite set of states, Γ is a finite tape alphabet, $\delta \subseteq (Q \times \Gamma \times Q \times \Gamma \times \{\leftarrow, \rightarrow\})$ is a transition relation, $q_0 \in Q$ is the initial state and $g : Q \rightarrow \{\vee, \wedge\}$ specifies the type of each state. We assume, w.l.o.g., that $\Gamma \subseteq \mathcal{V}_1$, and we denote by γ any sequence containing all the elements of Γ . For every move $\mu \in \{\leftarrow, \rightarrow\}$, $\mu(i)$ is defined as $i - 1$ if $\mu = \leftarrow$ and $i + 1$ otherwise. Assume that the considered ATM runs in time at most $n \in \mathbb{N}$ on some word w (where n is polynomial w.r.t. $|w|$.) Let w_i ($1 \leq i \leq n$) be the i -th character in w , where w_i is a blank symbol if $i > |w|$. For every natural number i with $1 \leq i \leq n$ and every $q \in Q$, we consider the predicate q^i of arity $2 + n + |\gamma|$ associated with the rules:

$$q^i(x, y_1, \dots, y_n, \gamma, z) \Leftarrow \exists x' \quad x \xrightarrow{z} (x') \\ \circ q^{\mu(i)}(x', y_1, \dots, y_{i-1}, b, y_{i+1}, \dots, y_n, \gamma, z) \circ y_i \simeq a \quad (9)$$

if $g(q) = \vee$, for all transitions $(q, a, q', b, \mu) \in \delta$ such that $1 \leq \mu(i) \leq n$.

$$q^i(x, y_1, \dots, y_n, \gamma, z) \Leftarrow \exists x'_1, \dots, x'_k \quad x \xrightarrow{z} (x'_1, \dots, x'_k) \\ \circ \bigcirc_{j=1}^k q^{\mu_j(i)}(x'_j, y_1, \dots, y_{i-1}, b_j, y_{i+1}, \dots, y_n, \gamma, z) \circ y_i \simeq a \quad (10)$$

if $g(q) = \wedge$, for all $a \in \Gamma$, where (q, a, q'_j, b_j, μ_j) (for $j \in \{1, \dots, k\}$) is the set of transitions of the form $(q, a, q', b, \mu) \in \delta$ with $1 \leq \mu(i) \leq n$.

It is clear that $p^1(x, w_1, \dots, w_n, \gamma, z)$ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable iff the considered ATM terminates on the word w and accepts w (every model of $p^1(x, w_1, \dots, w_n, \gamma, z)$ encodes an accepting derivation from the initial state p and the tape w_1, \dots, w_n , where the head is at position 1). Moreover, the rules are \exists -restricted, with $\gamma_{\mathcal{R}}(q^i) = \emptyset$ and $\mathcal{P}^* = \emptyset$.

M Proof of Theorem 33

As for Theorem 9, the proof is by reduction from the PCP (the same notations are used). Potential witnesses ω are encoded in a similar way than in the proof of Theorem 9, except that the use of the heaps \mathfrak{h}^μ and \mathfrak{h}^ν is avoided, instead the indices $I^\mu(\eta^\mu(j))$ and $I^\nu(\eta^\nu(j))$ are stored directly in the list \mathfrak{h}' . Moreover, an additional location ℓ'_j , called a *mark*, is added in each tuple, that is allocated and always refers to (\cdot) . Finally, each cell has a double link to the next one (this will be useful so that the tail can be generated twice in the body of the rule). We get a heap of the following form:

$$\mathfrak{h}' = \{(\ell_j, \mathfrak{s}(I^\mu(\eta^\mu(j))), \mathfrak{s}(I^\nu(\eta^\nu(j))), \ell'_j, \ell_{j+1}, \ell_{j+1}, \pi'), (\ell'_j, \pi') \mid 1 \leq j \leq |\omega|\}$$

Moreover, the list is cyclic, i.e., we assume that $\ell_{|\omega|+1} = \ell_1$. The rules defining the predicate generating the lists \mathfrak{h}' corresponding to potential witnesses are very similar to those given in the proof of Theorem 9, except that Q is not used (a predicate D is used instead to allocate marks). We assume that \mathbf{v} contains all natural numbers in $\{1, \dots, n\}$ and a variable u , denoting the first cell in the list. The predicate $P'(x'', y, \mathbf{v}, z)$ is called (taking advantage of the double link to the next element and from the fact that $*$ is used instead of \circ) to ensure that all the marks are pairwise distinct (it allocates a list in which all marks are distinct from y). This ensures that the mapping $\ell_j \mapsto \ell'_j$ is bijective, so that the cells can be unambiguously denoted by their marks.

$$\begin{aligned} P(x, \mathbf{v}, z) &\Leftarrow \exists x', x'', y \quad x \xrightarrow{z} (i, i, y, x', x'') \\ &\quad * P'(x'', y, \mathbf{v}, z) * x'' \simeq x' * P_{2,2,i,i}(x', \mathbf{v}, z) * D(y, z) \\ &\quad \text{if } \mu_i|_1 = \nu_i|_1 \end{aligned} \quad (11)$$

$$D(x, z) \Leftarrow x \xrightarrow{z} () \quad (12)$$

$$\begin{aligned} P_{j_\mu, j_\nu, i_\mu, i_\nu}(x, \mathbf{v}, z) &\Leftarrow \exists x', x'', y \quad x \xrightarrow{z} (i_\mu, i_\nu, y, x', x'') * P'(x'', y) * x'' \simeq x' * \\ &\quad * D(y, z) * P_{j_\mu+1, j_\nu+1, i_\mu, i_\nu}(x', \mathbf{v}, z) \\ &\quad \text{if } j_\mu \leq |\mu_{i_\mu}|, j_\nu \leq |\nu_{i_\nu}| \text{ and } \mu_{i_\mu}|_{j_\mu} = \nu_{i_\nu}|_{j_\nu} \end{aligned} \quad (13)$$

$$\begin{aligned} P_{j_\mu, j_\nu, i_\mu, i_\nu}(x, \mathbf{v}, z) &\Leftarrow \exists x', x'', y \quad x \xrightarrow{z} (i'_\mu, i_\nu, y, x', x'') * P'(x'', y) * x'' \simeq x' * \\ &\quad * D(y, z) * P_{2, j_\nu+1, i'_\mu, i_\nu}(x', \mathbf{v}, z) \\ &\quad \text{if } j_\mu = |\mu_{i_\mu}| + 1, j_\nu \leq |\nu_{i_\nu}| \text{ and } \mu_{i'_\mu}|_1 = \nu_{i_\nu}|_{j_\nu} \end{aligned} \quad (14)$$

$$\begin{aligned} P_{j_\mu, j_\nu, i_\mu, i_\nu}(x, \mathbf{v}, z) &\Leftarrow \exists x', x'', y \quad x \xrightarrow{z} (i_\mu, i'_\nu, y, x', x'') * P'(x'', y) * x'' \simeq x' * \\ &\quad * D(y, z) * P_{j_\mu+1, 2, i_\mu, i'_\nu}(x', \mathbf{v}, z) \\ &\quad \text{if } j_\mu \leq |\mu_{i_\mu}|, j_\nu = |\nu_{i_\nu}| + 1 \text{ and } \mu_{i_\mu}|_{j_\mu} = \nu_{i'_\nu}|_1 \end{aligned} \quad (15)$$

$$\begin{aligned} P_{j_\mu, j_\nu, i_\mu, i_\nu}(x, \mathbf{v}, z) &\Leftarrow \exists y \quad x \xrightarrow{z} (0, 0, y, u, u) * D(y, z) \\ &\quad \text{if } j_\mu = |\mu_{i_\mu}| \text{ and } j_\nu = |\nu_{i_\nu}| \end{aligned} \quad (16)$$

$$P'(x, y', \mathbf{v}, z) \Leftarrow \exists x', y \quad x \xrightarrow{z} (i_\mu, i_\nu, y, x', x') * y \not\approx y' * P'(x', y', \mathbf{v}, z) * D(y, z) \quad (17)$$

$$P'(x, y', \mathbf{v}, z) \Leftarrow \exists y \quad x \xrightarrow{z} (i_\mu, i_\nu, y, u, u) * y \not\approx y' * D(y, z) \quad (18)$$

To check that the sequences I^μ and I^ν are identical, we proceed as follows, exploiting the fact that \mathfrak{h} is cyclic and can be generated an unbounded number of times (since by hypothesis, for all $n \in \mathbb{N}$, there is a permission π such that π^n is defined). By construction, we already know that $I^\mu(1) = I^\nu(1)$, we need to check that $I^\mu(i) = I^\nu(i)$ holds for all $i \in \{2, \dots, \kappa\}$. We first guess an index ι (which is intended to denote the value of I^μ and I^ν) and we keep track of the mark of the first cell by storing it into parameters y_μ and y_ν . These parameters are intended to denote the mark of the start of the current word in the witness (note that it would be much simpler to keep track of the cells themselves, instead of their marks, but this would make the rules non $\ast\text{-}\exists$ -restricted). At this point we have $y_\mu = y_\nu$ as the words $\mu_{I^\mu(1)}$ and $\nu_{I^\nu(1)}$ always start at the same cell u , later y_μ and y_ν will take different values. Then we check that $I^\mu(2) = I^\nu(2) = \iota$. This is done by skipping the word $\mu_{I^\mu(1)}$ to find the start of the word $\mu_{I^\mu(2)}$. To this aim, we define a predicate $R_{\iota, i, j}^\lambda$ skipping the characters $j, \dots, |\lambda_i|$ in the word λ_i (rule 23). Once this is done, we check (rule 24) that $I^\mu(2) = \iota$. Simultaneously, we keep track of the mark of the cell that corresponds to the start of $\mu_{I^\mu(2)}$ by storing it into y_μ . Then, we call a predicate R_ι^ν which performs the same operation for the sequence ν . To this aim, we must go through the entire list (rule 20) to go back to the cell marked with y_ν (rule 21). Note that this is feasible as the list is cyclic. Afterwards, we may call the predicate $R_{\iota, i, 2}^\nu$ to check that $I^\nu(2) = \iota$ (rule 25) and find the start of the word $\nu_{I^\nu(2)}$ as done for μ . We also keep track of the mark of the cell that correspond to the start of $\mu_{I^\nu(2)}$ by storing it into y_ν . We then guess a new index ι'' and call the predicate $R_{\iota''}^\mu$ that performs exactly the same operations, but starting at cells marked y_μ and y_ν , respectively. This will check that the indices of the next words after the ones marked with y_μ and y_ν are identical, i.e., that $I^\mu(3) = I^\nu(3)$. We repeat this operation until y_μ and y_ν are equal (rule 22). In the rules below, λ ranges over the set $\{\mu, \nu\}$ and i, i', ι range over $\{1, \dots, n\}$:

$$R(x, \mathbf{v}, z) \Leftarrow \exists x', y \quad x \xrightarrow{\tilde{z}} (i, i', y, x', x') * D(y, z) * R_{\iota, i, 2}^\mu(x', y, y, \mathbf{v}, z) \quad (19)$$

$$R_\iota^\lambda(x, y_\mu, y_\nu, \mathbf{v}, z) \Leftarrow \exists x', y \quad x \xrightarrow{\tilde{z}} (i, i', y, x', x') * y \not\cong y_\lambda * D(y, z) * R_\iota^\lambda(x', y_\mu, y_\nu, \mathbf{v}, z) \quad (20)$$

$$R_\iota^\lambda(x, y_\mu, y_\nu, \mathbf{v}, z) \Leftarrow \exists x', y \quad x \xrightarrow{\tilde{z}} (i, i', y, x', x') * y \simeq y_\lambda * D(y, z) * R_{\iota, i, 2}^\lambda(x', y_\mu, y_\nu, \mathbf{v}, z) \quad (21)$$

$$R_\iota^\mu(x, y_\mu, y_\nu, \mathbf{v}, z) \Leftarrow \exists x', y \quad x \xrightarrow{\tilde{z}} (0, 0, y, u, u) * y_\mu \simeq y_\nu * D(y, z) \quad (22)$$

$$R_{\iota, i, j}^\lambda(x, y_\mu, y_\nu, \mathbf{v}, z) \Leftarrow \exists x', y \quad x \xrightarrow{\tilde{z}} (i, i', y, x', x') * D(y, z) * R_{\iota, i, j+1}^\lambda(x', y_\mu, y_\nu, \mathbf{v}, z) \\ \text{if } j \leq |\lambda_i| \quad (23)$$

$$R_{\iota,i,j}^{\mu}(x, y_{\mu}, y_{\nu}, \mathbf{v}, z) \Leftarrow \exists x', y \quad x \xrightarrow{z} (\iota, i', y, x', x') * D(y, z) * R_{\iota}^{\nu}(x', y, y_{\nu}, \mathbf{v}, z) \\ \text{if } j = |\mu_i| + 1 \quad (24)$$

$$R_{\iota,i,j}^{\nu}(x, y_{\mu}, y_{\nu}, \mathbf{v}, z) \Leftarrow \exists x', y \quad x \xrightarrow{z} (i', \iota, y, x', x') * D(y, z) * R_{\iota}^{\mu}(x', y_{\mu}, y, \mathbf{v}, z) \\ \text{if } j = |\nu_i| + 1 \quad (25)$$

It is easy to see that the formula $P(x, \mathbf{v}, z) * R(x, \mathbf{v}, z)$ is $(\mathcal{R}, \mathfrak{P})$ -satisfiable iff the corresponding instance of the PCP admits a solution. Moreover, the rules are $*\exists$ -restricted, with $\mathcal{P}^* = \{D\}$, $\gamma_{\mathcal{R}}(R) = \gamma_{\mathcal{R}}(P) = \gamma_{\mathcal{R}}(P_{j_{\mu}, j_{\nu}, i_{\mu}, i_{\nu}}) = \emptyset$, $\gamma_{\mathcal{R}}(P') = \{2\}$ and $\gamma_{\mathcal{R}}(R_{\iota}^{\lambda}) = \gamma_{\mathcal{R}}(R_{\iota,i,j}^{\lambda}) = \{2, 3\}$.