

A Relational Shape Abstract Domain

Hugo Illous^{1,2}, Matthieu Lemerre¹, and Xavier Rival²

¹ CEA, LIST, Software Reliability and Security Laboratory,
P.C. 174, Gif-sur-Yvette, 91191, France

² INRIA Paris/CNRS/École Normale Supérieure/PSL Research University

Abstract. Static analyses aim at inferring semantic properties of programs. While many analyses compute an over-approximation of reachable states, some analyses compute a description of the input-output relations of programs. In the case of numeric programs, several analyses have been proposed that utilize relational numerical abstract domains to describe relations. On the other hand, designing abstractions for relations over memory states and taking shapes into account is challenging. In this paper, we propose a set of novel logical connectives to describe such relations, which are inspired by separation logic. This logic can express that certain memory areas are unchanged, freshly allocated, or freed, or that only part of the memory was modified. Using these connectives, we build an abstract domain and design a static analysis that over-approximates relations over memory states containing inductive structures. We implement this analysis and report on the analysis of a basic library of list manipulating functions.

1 Introduction

Generally, static analyses aim at computing semantic properties of programs. Two common families of analyses are *reachability analyses*, that compute an over-approximation for the *set of reachable states* of programs, and *relational analyses*, that compute an over-approximation for the relations between input and output states. In general, sets of states are easier to abstract than state relations, which often makes reachability analyses simpler to design. On the other hand, abstracting relations brings several advantages:

- First, state relations allow to make the analyses modular [10,22,17,6,3] and compositional. Indeed, to analyze a sequence of two sub-programs, relational analyses can simply analyze each sub-program separately, and compose the resulting state relations. When sub-programs are functions, relational analyses may analyze each function separately, and compute one summary per function, so that the analysis of a function call does not require re-analyzing the body of the function, which is an advantage for scalability.

- Second, some properties can be expressed on state relations but not on sets of states, which makes relational analyses intrinsically more expressive. For example, contract languages [1,21] let functions be specified by formulas that may refer both to the input and to the output states. Such properties cannot be expressed using abstractions of sets of states, thus are beyond the scope of reachability analyses.

In general, the increased expressiveness of relational analyses requires more expressive abstractions. Let us discuss, as an example the case of numeric programs. A common way to express relations between input and output states consists in defining for each variable x a primed version x' that describes the value of x in the output state whereas the non primed version denotes the value of x in the input state. In this context, non-relational numerical abstract domain such as intervals [8] cannot capture any interesting relation between input and output states. On the other hand, relational numerical abstract domains such as convex polyhedra [7] can effectively capture relations between input and output states, as shown in [22]: for instance, when applied to a program that increments x by one, this analysis can infer the relation $x' = x + 1$.

In the context of programs manipulating complex data structures, relational analysis could allow to compute interesting classes of program properties. For instance, such analyses could express and verify that some memory areas were not physically modified by a program. Reachability analyses such as [24,15,5] cannot distinguish a program that inputs a list and leaves it unmodified from a program that inputs a list, copies it into an identical version and deallocates it, whereas a relational analysis could. More generally, it is often interesting to infer that a memory region is not modified by a program.

Separation logic [23] provides an elegant description for sets of states and is at the foundation of many reachability analyses for heap properties. In particular, the separating conjunction connective $*$ expresses that two regions are disjoint and allows local reasoning. On the other hand, it cannot describe state relations.

In this paper, we propose a logic inspired by separation logics and that can describe such properties. It provides connectives to describe that a memory region has been left unmodified by a program fragment, or that memory states can be split into disjoint sub-regions that undergo different transformations. We build an abstract domain upon this logic, and apply it to design an analysis for programs manipulating simple list or tree data structures. We make the following contributions:

- In Section 2, we demonstrate the abstraction of state relations using a specific family of heap predicates;
- In Section 4, we set up a logic to describe heap state relations and lift it into an abstract domain that describe concrete relations defined in Section 3;
- In Section 5, we design static analysis algorithms to infer heap state relations from abstract pre-condition;
- In Section 6, we report on experiments on basic linked data structures (lists and trees);
- Finally, we discuss related works in Section 7 and conclude in Section 8.

2 Overview and Motivating Example

We consider the example code shown in Figure 1, which implements the insertion of an element inside a non empty singly linked list containing integer values. When applied to a pointer to an existing non empty list and an integer value, this function

```

1 typedef struct list { struct list * next; int data; } list;
2 void insert_non_empty( list *l, int v ){
3     assume(l != NULL); list *c = l;
4     while( c->next != NULL && ... ){
5         c = c->next;
6     }
7     list *e = new( {next, data} ); // allocate 2 fields block
8     e->next = c->next; c->next = e; e->data = v;
9 }

```

Fig. 1. A list insertion program

traverses it partially (based on a condition on the values stored in list elements—that is elided in the figure). It then allocates a new list element, inserts it at the selected position and copies the integer argument into the `data` field. For instance, Figure 2(a) shows an input list containing elements 0, 8, 6, 1 and an output list where value 9 is inserted as a new element in the list. We observe that all elements of the input list are left physically unmodified except the element right before the insertion point. We now discuss abstractions of the behaviors of this program using abstractions for sets of states and abstractions for state relations.

Reachability analysis. First, we consider an abstraction based on separation logics with inductive predicates as used in [15,5]. We assume that the predicate $\mathbf{list}(\alpha)$ describes heap regions that consist of a well-formed linked list starting at address α (α is a symbolic variable used in the abstraction to denote a concrete address). This predicate is intuitively defined by induction as follows: it means either the region is empty and α is the null pointer, or the region is not empty, and consists of a list element of address α and with a `next` field containing a value described by symbolic variable β and a region that can be described by $\mathbf{list}(\beta)$. Thus, the valid input states for the insertion function can be abstracted by the abstract state shown in the top of Figure 2(b). The analysis of the function needs to express that the insertion occurs somewhere in the middle of the list. This requires a list segment predicate $\mathbf{listseg}(\alpha, \alpha')$, that is defined in a similar way as for \mathbf{list} : it describes region that stores a sub list starting at address α and the last element of which has a `next` field pointing to address α' (note that the empty region can be described by $\mathbf{listseg}(\alpha, \alpha)$). Using this predicate, we can now also express an abstraction for the output states of the insertion function: the abstract state shown in the bottom of Figure 2(b) describes the states where the new element was inserted in the middle of the structure (the list starts with a segment, then the predecessor of the inserted element, then the inserted element, and finally the list tail). We observe that this abstraction allows to express and to verify that the function is memory safe, and returns a well-formed list. Indeed, it captures the fact that no null or dangling pointer is ever dereferenced. Moreover, all states described by the abstract post-condition consist of a well-formed list, made of a segment, followed by two elements and a list tail. On the other hand, it does not say anything about




Fig. 2. Abstractions

the location of the list in the output state with respect to the list in the input state. More precisely, it cannot capture the fact that the elements of addresses a_0, a_1, a_3 are left unmodified physically. This is a consequence of the fact that each abstract state in Figure 2(b) independently describes a set of concrete heaps.

Relational analysis. To abstract *state relations* instead of sets of states, we now propose to define a new structure in Figure 2(c), that partially overlays the abstractions of input and output states. First, we observe that the tail of the list is not modified at all, thus, we describe it with a single predicate $\text{Id}(\mathbf{list}(\beta))$, that denotes pairs made of input state and an output state, that are *physically equal* and can both be described by $\mathbf{list}(\beta)$. The same kind of predicate can be used to describe that the initial segment has not changed between the two states. Second, we need to define a counterpart for separating conjunction at the relation level. Indeed, the effect of the insertion function can be decomposed as its effect on the initial segment (which is left unchanged), its effect on the tail (which is also left unchanged) and its effect on the insertion point (where a new element is allocated and a `next` pointer is modified). This relation separating conjunction is noted $*_{\mathbf{R}}$. To avoid confusion, from now on, we write $*_{\mathbf{S}}$ for the usual separating conjunction. Last, the insertion function allocates a new element and modifies the value of the `next` field of an existing element. To account for this, we need a new connective $[\cdot \dashrightarrow \cdot]$ which is applied to two abstract states: if h_0^\sharp, h_1^\sharp are abstract heaps (described by formulas in the usual separation logic with inductive predicates), then $[h_0^\sharp \dashrightarrow h_1^\sharp]$ describes the transformation of an input state described by h_0^\sharp into an output state described by h_1^\sharp . This is presented with different colors in the figure. In Section 4, we formalize this logics and the abstraction that it defines. The analysis by forward abstract interpretation [8] starts with the identity relation at function entry, and computes relations between input and output states step by step. The analysis algorithms need to unfold inductive predicates to materialize cells (for instance to analyze the test at line 4), and to fold inductive predicates in order to analyze loops. In addition to this, it also needs to reason over Id , $[\cdot \dashrightarrow \cdot]$ and $*_{\mathbf{R}}$ predicates, and perform operations similar to unfolding and folding on them. Section 5 describes the analysis algorithms.

3 Concrete Semantics

Before defining the abstraction, we fix notations for concrete states and programs.

We let \mathbb{X} denote the set of program variables and \mathbb{V} denote the set of values (that includes the set of numeric addresses). A field $\in \mathbb{F}$ (noted as **next**, **data**, ...) denotes both field names and offsets. A memory state $\sigma \in \mathbb{M}$ is a partial function from addresses to values. We write $\mathbf{dom}(\sigma)$ for the domain of σ , that is the set of addresses for which it is defined. Additionally, if σ_0, σ_1 are such that $\mathbf{dom}(\sigma_0) \cap \mathbf{dom}(\sigma_1) = \emptyset$, we let $\sigma_0 \otimes \sigma_1$ be the memory state obtained by merging σ_0 and σ_1 (its domain is $\mathbf{dom}(\sigma_0) \cup \mathbf{dom}(\sigma_1)$). If a_i is an address and v_i a value, we write $[a_0 \mapsto v_0; \dots; a_n \mapsto v_n]$ the memory state where a_i contains v_i (with $0 \leq i \leq n$).

In the following, we consider simple imperative programs, that include basic assignments, allocation and deallocation statements and loops (although our analysis supports a larger language, notably with conditionals and unstructured control flow). Programs are described by the grammar below:

$$\begin{array}{ll} \mathbf{L} ::= \mathbf{x} \quad (\mathbf{x} \in \mathbb{X}) \mid \mathbf{L} \rightarrow \mathbf{f} \quad (\mathbf{f} \in \mathbb{F}) & \text{l-values} \\ \mathbf{E} ::= v \quad (v \in \mathbb{V}) \mid \mathbf{L} \mid \mathbf{E} \oplus \mathbf{E} \quad (\oplus \in \{+, -, \leq, \dots\}) & \text{expressions} \\ \mathbf{P} ::= \mathbf{L} = \mathbf{E}; \mid \mathbf{L} = \mathbf{new}(\{\mathbf{f}_0, \dots\}); \mid \mathbf{free}(\mathbf{L}); \mid \mathbf{P}; \mathbf{P} \mid \mathbf{while}(\mathbf{E})\mathbf{P} & \text{programs} \end{array}$$

We assume the semantics of a program \mathbf{P} is defined as a function $\llbracket \mathbf{P} \rrbracket$ that maps a set of input states into a set of output states (thus $\llbracket \mathbf{P} \rrbracket : \mathcal{P}(\mathbb{M}) \rightarrow \mathcal{P}(\mathbb{M})$). We do not provide a full formal definition for $\llbracket \mathbf{P} \rrbracket$ as it is classical. Given a program \mathbf{P} , we define its *relational semantics* $\llbracket \mathbf{P} \rrbracket_{\mathcal{R}} : \mathbb{M} \rightarrow \mathbb{M} \times \mathbb{M}$ by:

$$\forall M \subseteq \mathbb{M}, \llbracket \mathbf{P} \rrbracket_{\mathcal{R}}(M) = \{(\sigma_0, \sigma_1) \mid \sigma_0 \in M \wedge \sigma_1 \in \llbracket \mathbf{P} \rrbracket(\{\sigma_0\})\}$$

In the following, we define an analysis to compute an over-approximation for $\llbracket \mathbf{P} \rrbracket_{\mathcal{R}}$.

4 Abstraction

In this section, we first define *abstract states*, that describe sets of memory states (as in [5]), and then we set up *abstract state relations*, that describe binary relations over memory states. Although our analysis and implementation support more general inductive predicates (such as trees and others), we consider only list inductive predicates in the body of the paper, for the sake of simplicity.

Abstract states. We assume a countable set $\mathbb{A} = \{\alpha, \beta, \dots\}$ of *symbolic addresses* that abstract values and heap addresses. An abstract state σ^\sharp consists of an abstract heap h^\sharp with a conjunction of numerical constraints such as equalities and disequalities. An abstract heap is a separating conjunction of region predicates that abstract *separate* memory regions [23] (as mentioned above, separating conjunction is denoted by $*_s$). A node $n \in \mathbb{N}$ is either a variable address $\&\mathbf{x}$ or a symbolic address α . A region predicate is either **emp** describing an empty region, or a points-to predicate $n \cdot \mathbf{f} \mapsto n'$ (that describes a heap memory cell at the base address n with the possibly null offset \mathbf{f} and with the content n'), or a summary

predicate $\mathbf{list}(n)$ describing a list structure or $\mathbf{listseg}(n, n')$ for a (possibly empty) list segment from address n to n' . The \mathbf{list} predicate is defined by induction as follows:

$$\begin{aligned} \mathbf{list}(n) &::= \mathbf{emp} \wedge n = \mathbf{0x0} \\ &\vee n \cdot \mathbf{next} \mapsto \alpha_n *_{\mathbf{S}} n \cdot \mathbf{data} \mapsto \alpha_d *_{\mathbf{S}} \mathbf{list}(\alpha_n) \wedge n \neq \mathbf{0x0} \end{aligned}$$

Segment predicate $\mathbf{listseg}$ stands for the segment version of \mathbf{list} and describes a list without a tail; it can also be defined by induction. We write $\xrightarrow{\text{unfold}}$ for the unfolding relation that syntactically transforms an instance of an inductive predicate into any of the disjuncts of that predicate.

Definition 1 (Abstract state). Abstract heaps and abstract states are defined by the grammar below:

$$\begin{aligned} c^\# &::= n \odot \mathbf{0x0} \quad (\odot \in \{=, \neq\}) \mid n = n' \mid c^\# \wedge c^\# \\ h^\#(\in \mathbb{H}) &::= \mathbf{emp} \mid n \cdot \mathbf{f} \mapsto n' \mid \mathbf{list}(n) \mid \mathbf{listseg}(n, n') \mid h^\# *_{\mathbf{S}} h^\# \\ \sigma^\#(\in \Sigma) &::= h^\# \wedge c^\# \quad n(\in \mathbb{N}) ::= \alpha \quad (\alpha \in \mathbb{A}) \mid \&\mathbf{x} \quad (\mathbf{x} \in \mathbb{X}) \end{aligned}$$

We now define the meaning of abstract heaps and abstract states using *concretization functions* [8], that associate to abstract elements the set of concrete elements they describe. To concretize an abstract heap, we also need to define how the nodes are bound into concrete values in concrete memories. We call *valuation* a function ν that maps nodes into concrete values and addresses.

Definition 2 (Concretization of abstract states). The concretization function $\gamma_{\mathbb{C}}$ maps a numeric constraint into a set of valuations whereas $\gamma_{\mathbb{H}}$ and γ_{Σ} respectively map an abstract heap and an abstract state into a set of pairs made of memory state and a valuation. They are defined by induction as follows:

$$\begin{aligned} \gamma_{\mathbb{C}}(n \odot \mathbf{0x0}) &= \{\nu \mid \nu(n) \odot \mathbf{0x0}\} \\ \gamma_{\mathbb{C}}(n = n') &= \{\nu \mid \nu(n) = \nu(n')\} \quad \gamma_{\mathbb{C}}(c_0^\# \wedge c_1^\#) = \gamma_{\mathbb{C}}(c_0^\#) \cap \gamma_{\mathbb{C}}(c_1^\#) \\ \gamma_{\mathbb{H}}(n \cdot \mathbf{f} \mapsto n') &= \{[\nu(n) + \mathbf{f} \mapsto \nu(n')], \nu\} \quad \gamma_{\mathbb{H}}(\mathbf{emp}) = \{([], \nu)\} \\ \gamma_{\mathbb{H}}(\mathbf{ind}) &= \bigcup \{\gamma_{\Sigma}(\sigma^\#) \mid \mathbf{ind} \xrightarrow{\text{unfold}} \sigma^\#\} \quad \text{if } \mathbf{ind} \text{ is } \mathbf{list}(n) \text{ or } \mathbf{listseg}(n, n') \\ \gamma_{\mathbb{H}}(h_0^\# *_{\mathbf{S}} h_1^\#) &= \{(\sigma_0 \otimes \sigma_1, \nu) \mid (\sigma_0, \nu) \in \gamma_{\mathbb{H}}(h_0^\#) \wedge (\sigma_1, \nu) \in \gamma_{\mathbb{H}}(h_1^\#)\} \\ \gamma_{\Sigma}(h^\# \wedge c^\#) &= \{(\sigma, \nu) \mid (\sigma, \nu) \in \gamma_{\mathbb{H}}(h^\#) \wedge \nu \in \gamma_{\mathbb{C}}(c^\#)\} \end{aligned}$$

Example 1 (Abstract state). The abstract pre-condition of the program of Figure 1 is $\&\mathbf{l} \mapsto \alpha *_{\mathbf{S}} \mathbf{list}(\alpha) *_{\mathbf{S}} \&\mathbf{v} \mapsto \beta$.

Abstract relations. An *abstract heap relation* describes a set of pairs made of an *input* memory state σ_i and an *output* memory state σ_o . Abstract heap relations are defined by the following connectives:

- the *identity relation* $\text{Id}(h^\#)$ describes pairs of memory states that are equal and are both abstracted by $h^\#$; this corresponds to the identity transformation;
- the *transformation relation* $[h_1^\# \dashrightarrow h_0^\#]$ describes pairs corresponding to the transformation of a memory state abstracted by $h_1^\#$ into a memory state abstracted by $h_0^\#$;

- the *relation separating conjunction* $r_0^\sharp *_{\mathbb{R}} r_1^\sharp$ of two heap relations r_0^\sharp, r_1^\sharp denotes a transformation that can be described by combining independently the transformations described by r_0^\sharp and r_1^\sharp on disjoint memory regions.

Definition 3 (Abstract relations). *The syntax of abstract heap relations and abstract state relations are defined by the grammar below:*

$$r^\sharp(\in \mathbb{R}) ::= \text{Id}(h^\sharp) \mid [h^\sharp \dashrightarrow h^\sharp] \mid r^\sharp *_{\mathbb{R}} r^\sharp \quad \rho^\sharp(\in \mathbb{II}) ::= r^\sharp \wedge c^\sharp$$

The concretization of relations also requires using valuations as it also needs to define the concrete values that nodes denote. It thus returns triples made of two memory states and a valuation.

Definition 4 (Concretization of abstract relations). *The concretization functions $\gamma_{\mathbb{R}}, \gamma_{\mathbb{II}}$ respectively map an abstract heap relation and an abstract state relation into elements of $\mathbb{M} \times \mathbb{M} \times (\mathbb{N} \rightarrow \mathbb{V})$. They are defined by:*

$$\begin{aligned} \gamma_{\mathbb{R}}(\text{Id}(h^\sharp)) &= \{(\sigma, \sigma, \nu) \mid (\sigma, \nu) \in \gamma_{\mathbb{H}}(h^\sharp)\} \\ \gamma_{\mathbb{R}}([h_1^\sharp \dashrightarrow h_0^\sharp]) &= \{(\sigma_i, \sigma_o, \nu) \mid (\sigma_i, \nu) \in \gamma_{\mathbb{H}}(h_1^\sharp) \wedge (\sigma_o, \nu) \in \gamma_{\mathbb{H}}(h_0^\sharp)\} \\ \gamma_{\mathbb{R}}(r_0^\sharp *_{\mathbb{R}} r_1^\sharp) &= \{(\sigma_{i,0} \otimes \sigma_{i,1}, \sigma_{o,0} \otimes \sigma_{o,1}, \nu) \mid \\ &\quad (\sigma_{i,0}, \sigma_{o,0}, \nu) \in \gamma_{\mathbb{R}}(r_0^\sharp) \wedge \mathbf{dom}(\sigma_{i,0}) \cap \mathbf{dom}(\sigma_{o,1}) = \emptyset \\ &\quad \wedge (\sigma_{i,1}, \sigma_{o,1}, \nu) \in \gamma_{\mathbb{R}}(r_1^\sharp) \wedge \mathbf{dom}(\sigma_{i,1}) \cap \mathbf{dom}(\sigma_{o,0}) = \emptyset\} \\ \gamma_{\mathbb{II}}(r^\sharp \wedge c^\sharp) &= \{(\sigma_i, \sigma_o, \nu) \mid (\sigma_i, \sigma_o, \nu) \in \gamma_{\mathbb{R}}(r^\sharp) \wedge \nu \in \gamma_{\mathbb{C}}(c^\sharp)\} \end{aligned}$$

We remark that $*_{\mathbb{R}}$ is commutative and associative.

Example 2 (Expressiveness). Let $r_0^\sharp = \text{Id}(\mathbf{list}(n))$ and $r_1^\sharp = [\mathbf{list}(n) \dashrightarrow \mathbf{list}(n)]$. We observe that r_0^\sharp describes only the identity transformation applied to a precondition where n is the address of a well-formed list, whereas r_1^\sharp describes any transformation that inputs such a list and also outputs such a list, but may modify its content, add or remove elements, or may modify the order of list elements (except for the first one which remains at address n). This means that $\gamma_{\mathbb{R}}(r_0^\sharp) \subset \gamma_{\mathbb{R}}(r_1^\sharp)$.

More generally, we have the following properties:

Theorem 1 (Properties). *Let $h^\sharp, h_0^\sharp, h_1^\sharp, h_{i,0}^\sharp, h_{i,1}^\sharp, h_{o,0}^\sharp, h_{o,1}^\sharp$ be abstract heaps. Then, we have the following properties*

1. $\gamma_{\mathbb{R}}(\text{Id}(h_0^\sharp *_{\mathbb{S}} h_1^\sharp)) = \gamma_{\mathbb{R}}(\text{Id}(h_0^\sharp) *_{\mathbb{R}} \text{Id}(h_1^\sharp))$
2. $\gamma_{\mathbb{R}}(\text{Id}(h^\sharp)) \subseteq \gamma_{\mathbb{R}}([h^\sharp \dashrightarrow h^\sharp])$ (the opposite inclusion may not hold, as observed in Example 2);
3. $\gamma_{\mathbb{R}}([h_{i,0}^\sharp \dashrightarrow h_{o,0}^\sharp] *_{\mathbb{R}} [h_{i,1}^\sharp \dashrightarrow h_{o,1}^\sharp]) \subseteq \gamma_{\mathbb{R}}([(h_{i,0}^\sharp *_{\mathbb{S}} h_{i,1}^\sharp) \dashrightarrow (h_{o,0}^\sharp *_{\mathbb{S}} h_{o,1}^\sharp)])$ (the opposite inclusion may not hold).

Example 3 (Abstract state relation). The effect of the insertion function of Figure 1 can be described by the abstract state relation $\text{Id}(h_0^\sharp) *_{\mathbb{R}} [h_1^\sharp \dashrightarrow h_2^\sharp] *_{\mathbb{R}} [\mathbf{emp} \dashrightarrow h_3^\sharp]$, where $h_0^\sharp = \&1 \mapsto \alpha_0 *_{\mathbb{S}} \&v \mapsto \beta *_{\mathbb{S}} \mathbf{listseg}(\alpha_0, \alpha_1) *_{\mathbb{S}} \mathbf{list}(\alpha_2) *_{\mathbb{S}} \alpha_1 \cdot \mathbf{data} \mapsto \beta_2$ (preserved region), $h_1^\sharp = \alpha_1 \cdot \mathbf{next} \mapsto \alpha_2$, $h_2^\sharp = \alpha_1 \cdot \mathbf{next} \mapsto \alpha$ (modified region) and $h_3^\sharp = \alpha \cdot \mathbf{next} \mapsto \alpha_2 *_{\mathbb{S}} \alpha \cdot \mathbf{data} \mapsto \beta$ (new region).

5 Analysis Algorithms

We now propose a static analysis to compute abstract state relations as described in Definition 3. It proceeds by forward abstract interpretation [8], starting from the abstract relation $\text{Id}(h^\sharp)$ where h^\sharp is a pre-condition, supplied by the user.

More generally, the analysis of a program P is a function $\llbracket P \rrbracket_{\mathcal{R}}^\sharp$ that inputs an abstract state relation describing a previous transformation \mathcal{T} done on the input *before* running P and returns a relation describing that transformation \mathcal{T} followed by the execution of P . Thus, $\llbracket P \rrbracket_{\mathcal{R}}^\sharp$ should meet the following soundness condition:

$$\begin{aligned} \forall \rho^\sharp \in \Pi, \forall (\sigma_0, \sigma_1) \in \gamma_\Pi(\rho^\sharp), \forall \sigma_2 \in \mathbb{M}, \\ (\sigma_1, \sigma_2) \in \llbracket P \rrbracket_{\mathcal{R}}^\sharp \implies (\sigma_0, \sigma_2) \in \gamma_\Pi(\llbracket P \rrbracket_{\mathcal{R}}^\sharp(\rho^\sharp)) \end{aligned}$$

5.1 Basic abstract post-conditions

We start with the computation of abstract post-condition for assignments, allocation and deallocation, on abstract relations that do not contain inductive predicates. As an example, we consider the analysis of an assignment $L = E$, starting from an abstract pre-condition relation r^\sharp . To compute the effect of this assignment on r^\sharp , the analysis should update it so as to reflect the modification of L in the output states of the pairs denoted by r^\sharp . We first consider the case where r^\sharp is a transformation relation.

Case of a transformation relation. We assume $r^\sharp = [h_0^\sharp \dashrightarrow h_1^\sharp]$. Then, if h_2^\sharp is an abstract state that describes the memory states after the assignment $L = E$, when it is executed on a state that is in $\gamma_{\text{H}}(h_1^\sharp)$, then a valid definition for $\llbracket L = E \rrbracket_{\mathcal{R}}^\sharp(r^\sharp)$ is $[h_0^\sharp \dashrightarrow h_2^\sharp]$. An algorithm for computing such a h_2^\sharp can be found in [5]. It first evaluates L into a points-to predicate $n \cdot \mathbf{f} \mapsto n'$ describing the cell that L represents, then evaluates E into a node n'' describing the value of the right hand side and finally replaces $n \cdot \mathbf{f} \mapsto n'$ with $n \cdot \mathbf{f} \mapsto n''$. As a consequence, we have the following definitions for the two main cases of assignments :

$$\begin{aligned} \llbracket \mathbf{x} = \mathbf{y} \rightarrow \mathbf{f} \rrbracket_{\mathcal{R}}^\sharp([h_0^\sharp \dashrightarrow (h_1^\sharp *_{\text{S}} \&\mathbf{x} \mapsto \alpha_0 *_{\text{S}} \&\mathbf{y} \mapsto \alpha_1 *_{\text{S}} \alpha_1 \cdot \mathbf{f} \mapsto \alpha_2)]) \\ = [h_0^\sharp \dashrightarrow (h_1^\sharp *_{\text{S}} \&\mathbf{x} \mapsto \alpha_2 *_{\text{S}} \&\mathbf{y} \mapsto \alpha_1 *_{\text{S}} \alpha_1 \cdot \mathbf{f} \mapsto \alpha_2)] \\ \llbracket \mathbf{x} \rightarrow \mathbf{f} = \mathbf{y} \rrbracket_{\mathcal{R}}^\sharp([h_0^\sharp \dashrightarrow (h_1^\sharp *_{\text{S}} \&\mathbf{x} \mapsto \alpha_0 *_{\text{S}} \alpha_0 \cdot \mathbf{f} \mapsto \alpha_1 *_{\text{S}} \&\mathbf{y} \mapsto \alpha_2)]) \\ = [h_0^\sharp \dashrightarrow (h_1^\sharp *_{\text{S}} \&\mathbf{x} \mapsto \alpha_0 *_{\text{S}} \alpha_0 \cdot \mathbf{f} \mapsto \alpha_2 *_{\text{S}} \&\mathbf{y} \mapsto \alpha_2)] \end{aligned}$$

Case of a separating conjunction relation. We now assume that $r^\sharp = r_0^\sharp *_{\text{R}} r_1^\sharp$. If the assignment can be fully analyzed on r_0^\sharp (i.e., it does not read or modify r_1^\sharp), then the following definition provides a sound transfer function, that relies on the same principle as the Frame rule [23] for separation logic:

$$\text{if } \llbracket L = E \rrbracket_{\mathcal{R}}^\sharp(r_0^\sharp) \text{ is defined, then } \llbracket L = E \rrbracket_{\mathcal{R}}^\sharp(r_0^\sharp *_{\text{R}} r_1^\sharp) = \llbracket L = E \rrbracket_{\mathcal{R}}^\sharp(r_0^\sharp) *_{\text{R}} r_1^\sharp$$

When $L = E$ writes in r_0^\sharp and reads in r_1^\sharp , we get a similar definition as above. For instance:

$$\begin{aligned} \llbracket \mathbf{x} = \mathbf{y} \rightarrow \mathbf{f} \rrbracket_{\mathcal{R}}^\sharp([h_0^\sharp \dashrightarrow (\&\mathbf{x} \mapsto \alpha_0)] *_{\text{R}} [h_1^\sharp \dashrightarrow (\&\mathbf{y} \mapsto \alpha_1 *_{\text{S}} \alpha_1 \cdot \mathbf{f} \mapsto \alpha_2)]) \\ [h_0^\sharp \dashrightarrow (\&\mathbf{x} \mapsto \alpha_2)] *_{\text{R}} [h_1^\sharp \dashrightarrow (\&\mathbf{y} \mapsto \alpha_1 *_{\text{S}} \alpha_1 \cdot \mathbf{f} \mapsto \alpha_2)] \end{aligned}$$

Case of an identity relation. We now assume that $r^\sharp = \text{Id}(h^\sharp)$. As observed in Theorem 1, $\gamma_{II}(\text{Id}(h^\sharp)) \subseteq \gamma_{II}([h^\sharp \dashrightarrow h^\sharp])$. We derive from the previous two paragraphs and from this principle the following definitions:

$$\begin{aligned} & \llbracket \mathbf{x} = \mathbf{y} \dashrightarrow \mathbf{f} \rrbracket_{\mathcal{R}}^\sharp (\text{Id}(h^\sharp *_{\mathcal{S}} \& \mathbf{x} \mapsto \alpha_0 *_{\mathcal{S}} \& \mathbf{y} \mapsto \alpha_1 *_{\mathcal{S}} \alpha_1 \cdot \mathbf{f} \mapsto \alpha_2)) \\ & = \text{Id}(h^\sharp *_{\mathcal{S}} \& \mathbf{y} \mapsto \alpha_1 *_{\mathcal{S}} \alpha_1 \cdot \mathbf{f} \mapsto \alpha_2) *_{\mathcal{R}} [(\& \mathbf{x} \mapsto \alpha_0) \dashrightarrow (\& \mathbf{x} \mapsto \alpha_2)] \\ & \llbracket \mathbf{x} \dashrightarrow \mathbf{f} = \mathbf{y} \rrbracket_{\mathcal{R}}^\sharp (\text{Id}(h^\sharp *_{\mathcal{S}} \& \mathbf{x} \mapsto \alpha_0 *_{\mathcal{S}} \alpha_0 \cdot \mathbf{f} \mapsto \alpha_1 *_{\mathcal{S}} \& \mathbf{y} \mapsto \alpha_2)) \\ & = \text{Id}(h^\sharp *_{\mathcal{S}} \& \mathbf{x} \mapsto \alpha_0 *_{\mathcal{S}} \& \mathbf{y} \mapsto \alpha_2) *_{\mathcal{R}} [(\alpha_0 \cdot \mathbf{f} \mapsto \alpha_1) \dashrightarrow (\alpha_0 \cdot \mathbf{f} \mapsto \alpha_2)] \end{aligned}$$

Other transfer functions. Condition tests boil down to numeric constraints intersections. The analysis of allocation needs to account for the creation of cells in the right side of relations whereas deallocation needs to account for the deletion of cells that were present before. Thus, for instance:

$$\begin{aligned} & \llbracket \mathbf{x} = \mathbf{new}(\{\mathbf{f}_0, \dots, \mathbf{f}_n\}) \rrbracket_{\mathcal{R}}^\sharp (r^\sharp *_{\mathcal{R}} [h^\sharp \dashrightarrow (\& \mathbf{x} \mapsto \alpha)]) \\ & = r^\sharp *_{\mathcal{R}} [h^\sharp \dashrightarrow (\& \mathbf{x} \mapsto \beta)] *_{\mathcal{R}} [\mathbf{emp} \dashrightarrow (\beta \cdot \mathbf{f}_0 \mapsto \beta_0 *_{\mathcal{S}} \dots *_{\mathcal{S}} \beta \cdot \mathbf{f}_n \mapsto \beta_n)] \\ & \quad \text{where } \beta, \beta_0, \dots, \beta_n \text{ are fresh} \\ & \llbracket \mathbf{free}(\mathbf{x}) \rrbracket_{\mathcal{R}}^\sharp (r^\sharp *_{\mathcal{R}} \text{Id}(\& \mathbf{x} \mapsto \alpha *_{\mathcal{S}} \alpha \cdot \mathbf{f}_0 \mapsto \alpha_0) *_{\mathcal{R}} [h_1^\sharp \dashrightarrow (h_0^\sharp *_{\mathcal{S}} \alpha \cdot \mathbf{f}_1 \mapsto \alpha_1)]) \\ & = r^\sharp *_{\mathcal{R}} \text{Id}(\& \mathbf{x} \mapsto \alpha) *_{\mathcal{R}} [(\alpha \cdot \mathbf{f}_0 \mapsto \alpha_0) \dashrightarrow \mathbf{emp}] *_{\mathcal{R}} [h_1^\sharp \dashrightarrow h_0^\sharp] \end{aligned}$$

5.2 Materialization and general abstract post-conditions

In Section 5.1, we considered only abstract states without inductive predicates, to first provide a simpler definition of abstract post-conditions. We now lift this restriction. For example, the analysis of the program in Figure 1 starts with $\text{Id}(\& \mathbf{l} \mapsto \alpha *_{\mathcal{S}} \mathbf{list}(\alpha) *_{\mathcal{S}} \& \mathbf{v} \mapsto \beta)$, and then has to analyze a reading of $\mathbf{l} \dashrightarrow \mathbf{next}$.

If we consider an abstract state relation of the form $[h^\sharp \dashrightarrow \mathbf{list}(\mathbf{n})]$, and an assignment that reads or writes a field at base address \mathbf{n} , the inductive predicate $\mathbf{list}(\mathbf{n})$ should first be *unfolded* [5]: before the post-condition operators of Section 5.1 can be applied, this predicate first needs to be substituted with the disjunction of cases it is made of, as defined in Section 4. This process is known in reachability shape analyses as a technique to materialize cells [24,15,5]. It results in disjunctive abstract states. For instance, the concretization of the abstract state relation $[h^\sharp \dashrightarrow \mathbf{list}(\mathbf{n})]$ is included in the union of the concretizations of $[h^\sharp \dashrightarrow \mathbf{emp}] \wedge \mathbf{n} = \mathbf{0x0}$ and $[h^\sharp \dashrightarrow (\mathbf{n} \cdot \mathbf{next} \mapsto \alpha_n *_{\mathcal{S}} \mathbf{n} \cdot \mathbf{data} \mapsto \alpha_d *_{\mathcal{S}} \mathbf{list}(\alpha_n))] \wedge \mathbf{n} \neq \mathbf{0x0}$. This disjunctive abstract states allows to analyze a read or write into a field at address \mathbf{n} .

However, this naive extension of unfolding may be imprecise here. Let us consider the unfolding at node \mathbf{n} in the abstract state relation $[\mathbf{n} \cdot \mathbf{next} \mapsto \alpha *_{\mathcal{S}} \mathbf{n} \cdot \mathbf{data} \mapsto \beta \dashrightarrow \mathbf{list}(\mathbf{n})]$. The above technique will generate two disjuncts, including one where $\mathbf{n} = \mathbf{0x0}$. However, \mathbf{n} cannot be equal to the null pointer here, since \mathbf{n} is the base address of a regular list element in the left side of the $[\cdot \dashrightarrow \cdot]$ abstract relation. Therefore, unfolding should take into account information in both sides of abstract relations for the sake of analysis precision.

In the following, we let $\mathbf{unfold}_{\Sigma}(\mathbf{n}, \sigma^\sharp)$ denote the set of disjuncts produced by unfolding an inductive predicate at node \mathbf{n} in abstract state σ^\sharp , if any. For instance, $\mathbf{unfold}_{\Sigma}(\mathbf{n}, \mathbf{list}(\mathbf{n}))$ is $\{(\mathbf{emp} \wedge \mathbf{n} = \mathbf{0x0}), (\mathbf{n} \cdot \mathbf{next} \mapsto \alpha_n *_{\mathcal{S}} \mathbf{n} \cdot \mathbf{data} \mapsto \alpha_d *_{\mathcal{S}}$

$\mathbf{list}(\alpha_n) \wedge n \neq \mathbf{0x0}$). If there is no inductive predicate attached to node n in σ^\sharp , we let $\mathbf{unfold}_\Sigma(\alpha, \sigma^\sharp) = \{\sigma^\sharp\}$. This operator is sound in the sense that, $\gamma_\Sigma(\sigma^\sharp)$ is included in $\cup\{\gamma_\Sigma(\sigma_u^\sharp) \mid \sigma_u^\sharp \in \mathbf{unfold}_\Sigma(n, \sigma^\sharp)\}$.

Using \mathbf{unfold}_Σ , we define the function \mathbf{unfold}_Π that performs unfolding at a given node and in an abstract state relation as follows:

- $\mathbf{unfold}_\Pi(n, \mathbf{Id}(h_u^\sharp)) = \{\mathbf{Id}(h_u^\sharp) \wedge c_u^\sharp \mid (h_u^\sharp \wedge c_u^\sharp) \in \mathbf{unfold}_\Sigma(n, h^\sharp)\}$;
- if the node n carries inductive predicate in r_0^\sharp then $\mathbf{unfold}_\Pi(n, r_0^\sharp *_R r_1^\sharp) = \{(r_{0,u}^\sharp *_R r_1^\sharp) \wedge c_{0,u}^\sharp \mid (r_{0,u}^\sharp \wedge c_{0,u}^\sharp) \in \mathbf{unfold}_\Pi(n, r_0^\sharp)\}$;
- $\mathbf{unfold}_\Pi(n, [h_i^\sharp \dashrightarrow h_o^\sharp]) = \{[h_{i,u}^\sharp \dashrightarrow h_{o,u}^\sharp] \wedge (c_{i,u}^\sharp \wedge c_{o,u}^\sharp) \mid (h_{i,u}^\sharp \wedge c_{i,u}^\sharp) \in \mathbf{unfold}_\Sigma(n, h_i^\sharp) \wedge (h_{o,u}^\sharp \wedge c_{o,u}^\sharp) \in \mathbf{unfold}_\Sigma(n, h_o^\sharp)\}$;
- $\mathbf{unfold}_\Pi(n, r^\sharp \wedge c^\sharp) = \{r_u^\sharp \wedge (c^\sharp \wedge c_u^\sharp) \mid (r_u^\sharp \wedge c_u^\sharp) \in \mathbf{unfold}_\Pi(n, r^\sharp)\}$.

We note that conjunctions of numerical constraints over node may yield to unfeasible elements being discarded in the last two cases: for instance, in the $[\cdot \dashrightarrow \cdot]$ case, unfolding will only retain disjuncts where both sides of the arrow express compatible conditions over n .

We can prove by case analysis that this unfolding operator is sound:

$$\gamma_\Pi(\rho^\sharp) \subseteq \bigcup\{\gamma_\Pi(\rho_u^\sharp) \mid \rho_u^\sharp \in \mathbf{unfold}_\Pi(n, \rho^\sharp)\}$$

Example 4 (Abstract state relation unfolding and post-condition). Let us consider the analysis of the insertion function of Figure 1. This function should be applied to states where l is a non null list pointer (the list should have at least one element), thus, the analysis should start from $\mathbf{Id}(\&l \mapsto \alpha *_S \mathbf{list}(\alpha)) \wedge \alpha \neq \mathbf{0x0}$ (in this example, we omit v for the sake of concision). Before the loop entry, the analysis computes the abstract state relation $\mathbf{Id}(\&l \mapsto \alpha *_S \mathbf{list}(\alpha)) *_R [\mathbf{emp} \dashrightarrow (\&c \mapsto \alpha)] \wedge \alpha \neq \mathbf{0x0}$. To deal with the test $c \dashrightarrow \mathbf{next} \neq \mathbf{NULL}$ (and the assignment $c = c \dashrightarrow \mathbf{next}$), the analysis should materialize the cell at node α . This unfolding is performed under the \mathbf{Id} connective, and produces:

$$\begin{aligned} &\mathbf{Id}(\&l \mapsto \alpha *_S \alpha \cdot \mathbf{next} \mapsto \alpha_0 *_S \alpha \cdot \mathbf{data} \mapsto \beta_0 *_S \mathbf{list}(\alpha_0)) \\ &*_R [\mathbf{emp} \dashrightarrow (\&c \mapsto \alpha)] \wedge \alpha \neq \mathbf{0x0} \end{aligned}$$

In turn, the effect of the condition test and of the assignment in the loop body can be precisely analyzed from this abstract state relation.

5.3 Folding and lattice operations

Like classical shape analyses [15,5], our analysis needs to *fold* inductive predicates so as to (conservatively) decide inclusion and join abstract states. We present folding algorithms in the following paragraphs.

Conservative inclusion checking. Inclusion checking is used to verify logical entailment, to check the convergence of loop iterates, and to support the join / widening algorithm. It consists of a conservative function \mathbf{isle}_Π over abstract states and a conservative function \mathbf{isle}_R over abstract state relations, that either return **true**

$$\begin{array}{c}
 \frac{\mathbf{h}^\sharp \text{ is of the form } n \cdot \mathbf{f} \mapsto n' \text{ or } \mathbf{list}(n) \text{ or } \mathbf{listseg}(n, n')}{\mathbf{h}^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{h}^\sharp} \quad (\sqsubseteq=) \\
 \\
 \frac{\mathbf{h}^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{list}(n')}{\mathbf{listseg}(n, n') *_{\mathbb{S}} \mathbf{h}^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{list}(n)} \quad (\sqsubseteq_{\text{seg}}) \qquad \frac{\mathbf{h}_{0,0}^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{h}_{1,0}^\sharp \quad \mathbf{h}_{0,1}^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{h}_{1,1}^\sharp}{\mathbf{h}_{0,0}^\sharp *_{\mathbb{S}} \mathbf{h}_{0,1}^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{h}_{1,0}^\sharp *_{\mathbb{S}} \mathbf{h}_{1,1}^\sharp} \quad (\sqsubseteq_{*_{\mathbb{S}}}) \\
 \frac{\mathbf{r}_u^\sharp \in \mathbf{unfold}_{\Pi}(n, \mathbf{r}_1^\sharp) \quad \mathbf{r}_0^\sharp \sqsubseteq_{\mathbb{R}} \mathbf{r}_u^\sharp \quad \mathbf{r}_1^\sharp \text{ contains } \mathbf{list}(n) \text{ or } \mathbf{listseg}(n, n')}{\mathbf{r}_0^\sharp \sqsubseteq_{\mathbb{R}} \mathbf{r}_1^\sharp} \quad (\sqsubseteq_{\text{unfold}}) \\
 \\
 \frac{\mathbf{h}_0^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{h}_1^\sharp}{\mathbf{Id}(\mathbf{h}_0^\sharp) \sqsubseteq_{\mathbb{R}} \mathbf{Id}(\mathbf{h}_1^\sharp)} \quad (\sqsubseteq_{\text{Id}}) \qquad \frac{\mathbf{h}_{i,0}^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{h}_{i,1}^\sharp \quad \mathbf{h}_{o,0}^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{h}_{o,1}^\sharp}{[\mathbf{h}_{i,0}^\sharp \dashrightarrow \mathbf{h}_{o,0}^\sharp] \sqsubseteq_{\mathbb{R}} [\mathbf{h}_{i,1}^\sharp \dashrightarrow \mathbf{h}_{o,1}^\sharp]} \quad (\sqsubseteq_{\dashrightarrow\text{-intro}}) \\
 \frac{\mathbf{r}^\sharp *_{\mathbb{R}} [\mathbf{h}^\sharp \dashrightarrow \mathbf{h}^\sharp] \sqsubseteq_{\mathbb{R}} [\mathbf{h}_i^\sharp \dashrightarrow \mathbf{h}_o^\sharp]}{\mathbf{r}^\sharp *_{\mathbb{R}} \mathbf{Id}(\mathbf{h}^\sharp) \sqsubseteq_{\mathbb{R}} [\mathbf{h}_i^\sharp \dashrightarrow \mathbf{h}_o^\sharp]} \quad (\sqsubseteq_{\text{Id-weak}}) \qquad \frac{\mathbf{r}_{0,0}^\sharp \sqsubseteq_{\mathbb{R}} \mathbf{r}_{1,0}^\sharp \quad \mathbf{r}_{0,1}^\sharp \sqsubseteq_{\mathbb{R}} \mathbf{r}_{1,1}^\sharp}{\mathbf{r}_{0,0}^\sharp *_{\mathbb{R}} \mathbf{r}_{0,1}^\sharp \sqsubseteq_{\mathbb{R}} \mathbf{r}_{1,0}^\sharp *_{\mathbb{R}} \mathbf{r}_{1,1}^\sharp} \quad (\sqsubseteq_{*_{\mathbb{R}}}) \\
 \frac{\mathbf{r}^\sharp *_{\mathbb{R}} [\mathbf{h}_{i,0}^\sharp *_{\mathbb{S}} \mathbf{h}_{i,1}^\sharp \dashrightarrow \mathbf{h}_{o,0}^\sharp *_{\mathbb{S}} \mathbf{h}_{o,1}^\sharp] \sqsubseteq_{\mathbb{R}} [\mathbf{h}_i^\sharp \dashrightarrow \mathbf{h}_o^\sharp]}{\mathbf{r}^\sharp *_{\mathbb{R}} [\mathbf{h}_{i,0}^\sharp \dashrightarrow \mathbf{h}_{o,0}^\sharp] *_{\mathbb{R}} [\mathbf{h}_{i,1}^\sharp \dashrightarrow \mathbf{h}_{o,1}^\sharp] \sqsubseteq_{\mathbb{R}} [\mathbf{h}_i^\sharp \dashrightarrow \mathbf{h}_o^\sharp]} \quad (\sqsubseteq_{\dashrightarrow\text{-weak}})
 \end{array}$$

Fig. 3. Inclusion checking rules

(meaning that the inclusion of concretizations holds) or **false** (meaning that the analysis cannot conclude whether inclusion holds).

Their definition relies on a conservative algorithm, that implements a proof search, based on the rules shown in Figure 3 (for clarity, we omit the numerical constraints inclusion checking). In this system of rules, if $\mathbf{h}_0^\sharp \sqsubseteq_{\mathbb{H}} \mathbf{h}_1^\sharp$ (resp., $\mathbf{r}_0^\sharp \sqsubseteq_{\mathbb{R}} \mathbf{r}_1^\sharp$), then $\gamma_{\mathbb{H}}(\mathbf{h}_0^\sharp) \subseteq \gamma_{\mathbb{H}}(\mathbf{h}_1^\sharp)$ (resp., $\gamma_{\mathbb{R}}(\mathbf{r}_0^\sharp) \subseteq \gamma_{\mathbb{R}}(\mathbf{r}_1^\sharp)$). The rules $(\sqsubseteq=)$, $(\sqsubseteq_{\text{seg}})$ and $(\sqsubseteq_{*_{\mathbb{S}}})$ are specific to reasoning of abstract states, and are directly inspired from [5] (they allow to reason over equal abstract regions, over segments, and over separating conjunction). The rule $(\sqsubseteq_{\text{unfold}})$ allows to reason by unfolding of inductive predicates, at the level of relations. Finally, the rules $(\sqsubseteq_{\text{Id}})$, $(\sqsubseteq_{\dashrightarrow\text{-intro}})$, $(\sqsubseteq_{\text{Id-weak}})$, $(\sqsubseteq_{*_{\mathbb{R}}})$ and $(\sqsubseteq_{\dashrightarrow\text{-weak}})$ allow to derive inclusion over abstract state relations, and implement the properties observed in Theorem 1. The proof search algorithm starts from the goal to prove and attempt to apply these rules so as to complete an inclusion derivation. We observe that abstract states are equivalent up to a renaming of the internal nodes (the nodes that are not of the form $\&\mathbf{x}$), thus, the implementation also takes care of this renaming, although the rules of Figure 3 do not show it, as this issue is orthogonal to the reasoning over abstract state relations which is the goal of this paper (indeed, this requires complex renaming functions that are made fully explicit in [5]). The rules can be proved sound one by one, thus they define a sound inclusion checking procedure:

Theorem 2 (Soundness of inclusion checking). *If $\mathbf{h}_0^\sharp, \mathbf{h}_1^\sharp \in \mathbb{H}$ and $\mathbf{r}_0^\sharp, \mathbf{r}_1^\sharp \in \mathbb{R}$ then:*

$$\begin{aligned}
 \mathbf{isle}_{\mathbb{H}}(\mathbf{h}_0^\sharp, \mathbf{h}_1^\sharp) = \mathbf{true} &\implies \gamma_{\mathbb{H}}(\mathbf{h}_0^\sharp) \subseteq \gamma_{\mathbb{H}}(\mathbf{h}_1^\sharp) \\
 \mathbf{isle}_{\mathbb{R}}(\mathbf{r}_0^\sharp, \mathbf{r}_1^\sharp) = \mathbf{true} &\implies \gamma_{\mathbb{R}}(\mathbf{r}_0^\sharp) \subseteq \gamma_{\mathbb{R}}(\mathbf{r}_1^\sharp)
 \end{aligned}$$

Example 5 (Inclusion checking). Let us consider the following abstract state relations, and discuss the computation of $\mathbf{isle}_{\mathbb{R}}(r_0^\#, r_1^\#)$:

$$\begin{aligned} r_0^\# &= \text{Id}(n \cdot \mathbf{next} \mapsto \alpha_0 *_{\mathbb{S}} \mathbf{list}(\alpha_0)) *_{\mathbb{R}} [n \cdot \mathbf{data} \mapsto \alpha_1 \dashrightarrow n \cdot \mathbf{data} \mapsto \alpha_2] \\ r_1^\# &= [\mathbf{list}(n) \dashrightarrow \mathbf{list}(n)] \end{aligned}$$

Using first rule ($\sqsubseteq_{\text{Id-weak}}$) then rule ($\sqsubseteq_{\dashrightarrow\text{-weak}}$), this goal gets reduced into checking the inclusion $[h_0^\# \dashrightarrow h_1^\#] \sqsubseteq_{\mathbb{R}} r_1^\#$, where $h_0^\# = n \cdot \mathbf{next} \mapsto \alpha_0 *_{\mathbb{S}} \mathbf{list}(\alpha_0) *_{\mathbb{S}} n \cdot \mathbf{data} \mapsto \alpha_1$ and $h_1^\# = n \cdot \mathbf{next} \mapsto \alpha_0 *_{\mathbb{S}} \mathbf{list}(\alpha_0) *_{\mathbb{S}} n \cdot \mathbf{data} \mapsto \alpha_2$. In turn, this inclusion follows from rule ($\sqsubseteq_{\text{unfold}}$).

Join / widening operators. In the following, we define abstract operators $\mathbf{wid}_{\mathbb{H}}$, $\mathbf{wid}_{\mathbb{R}}$ that respectively operate over abstract states and abstract state relations, and compute an over-approximation for concrete unions. They also ensure termination and serve as widening. The algorithm to compute these two functions heavily relies on the inclusion checking that was discussed in the previous paragraph. Indeed, the widening functions compute results that are more approximate than their arguments. To achieve this, they search for syntactic patterns in their arguments and produce outputs that inclusion checking proves more general. This process is performed region by region on both arguments of the widening, as formalized in [5, Figure 7]. We discuss in the following a list of such widening rules:

- when both arguments of widening are equal to a same base predicate, widening is trivial, and returns the same base predicate, thus for instance:

$$\begin{aligned} \mathbf{wid}_{\mathbb{H}}(n \cdot f \mapsto \alpha, n \cdot f \mapsto \alpha) &= n \cdot f \mapsto \alpha \\ \mathbf{wid}_{\mathbb{H}}(\mathbf{list}(\alpha), \mathbf{list}(\alpha)) &= \mathbf{list}(\alpha) \end{aligned}$$

- when applied to two abstract relations that consist of the same connective, the widening functions simply calls themselves recursively on the sub-components:

$$\begin{aligned} \mathbf{wid}_{\mathbb{R}}(\text{Id}(h_0^\#), \text{Id}(h_1^\#)) &= \text{Id}(\mathbf{wid}_{\mathbb{H}}(h_0^\#, h_1^\#)) \\ \mathbf{wid}_{\mathbb{R}}([h_{i,0}^\# \dashrightarrow h_{o,0}^\#], [h_{i,1}^\# \dashrightarrow h_{o,1}^\#]) &= [\mathbf{wid}_{\mathbb{H}}(h_{i,0}^\#, h_{i,1}^\#) \dashrightarrow \mathbf{wid}_{\mathbb{H}}(h_{o,0}^\#, h_{o,1}^\#)] \\ \mathbf{wid}_{\mathbb{R}}(r_{0,0}^\# *_{\mathbb{R}} r_{0,1}^\#, r_{1,0}^\# *_{\mathbb{R}} r_{1,1}^\#) &= \mathbf{wid}_{\mathbb{R}}(r_{0,0}^\#, r_{1,0}^\#) *_{\mathbb{R}} \mathbf{wid}_{\mathbb{R}}(r_{0,1}^\#, r_{1,1}^\#) \end{aligned}$$

- when applied to an $\text{Id}(\cdot)$ predicate and another abstract relation, widening first tries to maintain the $\text{Id}(\cdot)$ predicate, and, if this fails, tries to weaken it into an $[\cdot \dashrightarrow \cdot]$ predicate:

$$\begin{aligned} &\text{if } \mathbf{isle}_{\mathbb{H}}(h_0^\#, h^\#) = \mathbf{true} \text{ then,} \\ \mathbf{wid}_{\mathbb{R}}(\text{Id}(h_0^\#), r^\#) &= \begin{cases} \text{Id}(h^\#) & \text{if } \mathbf{isle}_{\mathbb{R}}(r^\#, \text{Id}(h^\#)) = \mathbf{true} \\ [h^\# \dashrightarrow h^\#] & \text{otherwise, if } \mathbf{isle}_{\mathbb{R}}(r^\#, [h^\# \dashrightarrow h^\#]) = \mathbf{true} \end{cases} \end{aligned}$$

- when applied to an $[\cdot \dashrightarrow \cdot]$ predicate, the widening tries to weaken the other argument accordingly:

$$\begin{aligned} &\text{if } \mathbf{isle}_{\mathbb{H}}(h_{i,0}^\#, h_i^\#) = \mathbf{true} \text{ and } \mathbf{isle}_{\mathbb{H}}(h_{o,0}^\#, h_o^\#) = \mathbf{true} \\ &\text{and } \mathbf{isle}_{\mathbb{R}}(r^\#, [h_i^\# \dashrightarrow h_o^\#]) = \mathbf{true} \text{ then,} \\ &\quad \mathbf{wid}_{\mathbb{R}}([h_{i,0}^\# \dashrightarrow h_{o,0}^\#], r^\#) = [h_i^\# \dashrightarrow h_o^\#] \end{aligned}$$

Each of these operations is sound, and the results computed by widening are also sound:

Theorem 3 (Soundness of widening). *If $h_0^\sharp, h_1^\sharp \in \mathbb{H}$ and $r_0^\sharp, r_1^\sharp \in \mathbb{R}$ then:*

$$\gamma_{\mathbb{H}}(h_0^\sharp) \cup \gamma_{\mathbb{H}}(h_1^\sharp) \subseteq \gamma_{\mathbb{H}}(\mathbf{wid}_{\mathbb{H}}(h_0^\sharp, h_1^\sharp)) \quad \gamma_{\mathbb{R}}(r_0^\sharp) \cup \gamma_{\mathbb{R}}(r_1^\sharp) \subseteq \gamma_{\mathbb{R}}(\mathbf{wid}_{\mathbb{R}}(r_0^\sharp, r_1^\sharp))$$

Furthermore, termination of widening follows from an argument similar to [5].

Example 6 (Widening). We consider the analysis of the program of Figure 1, and more specifically, the widening after the first abstract iteration over the loop:

$$\begin{aligned} & \mathbf{wid}_{\mathbb{R}}(\text{Id}(\&l \mapsto \alpha *_{\mathbb{S}} \mathbf{list}(\alpha) *_{\mathbb{S}} \&v \mapsto \beta) *_{\mathbb{R}} [\mathbf{emp} \dashrightarrow \&c \mapsto \alpha], \\ & \text{Id}(\&l \mapsto \alpha *_{\mathbb{S}} \alpha \cdot \mathbf{data} \mapsto \alpha_{\mathbb{d}} *_{\mathbb{S}} \alpha \cdot \mathbf{next} \mapsto \alpha_{\mathbb{n}} *_{\mathbb{S}} \mathbf{list}(\alpha_{\mathbb{n}}) *_{\mathbb{S}} \&v \mapsto \beta) \\ & \quad *_{\mathbb{R}} [\mathbf{emp} \dashrightarrow \&c \mapsto \alpha_{\mathbb{n}}]) \\ & = \text{Id}(\&l \mapsto \alpha *_{\mathbb{S}} \mathbf{listseg}(\alpha, \alpha') *_{\mathbb{S}} \mathbf{list}(\alpha') *_{\mathbb{S}} \&v \mapsto \beta) *_{\mathbb{R}} [\mathbf{emp} \dashrightarrow \&c \mapsto \alpha'] \end{aligned}$$

This abstract widening performs some generalization and introduces a list segment inductive predicate, that over-approximates an empty segment in the left argument, and a segment of length one. It also involves some renaming of symbolic nodes (as observed in the previous paragraph, the concretization of an abstract states is unchanged under symbolic nodes renaming).

5.4 Analysis

The abstract semantics $\llbracket \cdot \rrbracket_{\mathcal{R}}^\sharp$ relies on the abstract operations defined in Section 5.1, on the unfolding of Section 5.2 to analyze basic statements, and on the folding operations defined in Section 5.3 to cope with control flow joins and loop invariants computation. Soundness follows from the soundness of the basic operations.

Theorem 4 (Soundness). *The analysis is sound in the sense that, for all program P and for all abstract state relation ρ^\sharp :*

$$\forall (\sigma_0, \sigma_1) \in \gamma_{\Pi}(\rho^\sharp), \forall \sigma_2 \in \mathbb{M}, (\sigma_1, \sigma_2) \in \llbracket P \rrbracket_{\mathcal{R}} \implies (\sigma_0, \sigma_2) \in \gamma_{\Pi}(\llbracket P \rrbracket_{\mathcal{R}}^\sharp(\rho^\sharp))$$

6 Experimental Evaluation

In this section, we report on the implementation of our analysis and try to evaluate:

1. whether it can prove precise and useful relational properties, and
2. how it compares with a more classical reachability shape analysis.

Our implementation supports built-in inductive predicates to describe singly linked lists and binary trees. It provides both the analysis described in this paper, and a basic reachability shape analysis in the style of [5], and supporting the same inductive predicates. It was implemented as a Frama-C [19] plugin consisting of roughly 7800 lines of OCaml. We have ran both the reachability shape analysis and relational shape analysis on series of small programs manipulating lists and trees listed in Table 1. These tests are selected to test specifically the relational domain (and not a full analysis). This allows us to not only assess the results of

Structure	Function	Time (in ms)		Loop iterations	Relational Property
		Reach	Relat.		
sll	allocation	0.53	1.27	2	yes
sll	deallocation	0.34	0.99	2	yes
sll	traversal	0.53	0.83	2	yes
sll	insertion (head)	0.32	0.33	0	yes
sll	insertion (random pos)	1.98	2.75	2	yes
sll	insertion (random)	2.33	3.94	2	yes
sll	reverse	0.52	2.36	2	partial
sll	map	0.66	1.17	2	partial
tree	allocation	0.94	2.21	2	yes
tree	search	1.06	1.76	2	yes

Table 1. Experiment results (sll: singly linked lists; tree: binary trees; time in milliseconds averaged over 1000 runs on a laptop with Intel Core i7 running at 2.3 GHz, with 16 Gb RAM, for the reachability and relational analyses; the last column states whether the relational shape analysis computed the expected abstract relation)

the analysis computing abstract state relations, but also to compare them with an analysis that infers abstract states.

First, we discuss whether the analysis computing abstract state relations computes the expected relations, that describes the most precisely the transformation implemented by the analyzed function. As an example, in the case of an insertion at the head of a list, we expect the abstract relation below, that expresses that the body of the list was not modified:

$$[\&l \mapsto \alpha \dashrightarrow \&l \mapsto \beta] *_{\text{R}} [\mathbf{emp} \dashrightarrow \beta \cdot \mathbf{next} \mapsto \alpha *_{\text{S}} \beta \cdot \mathbf{data} \mapsto \delta] *_{\text{R}} \text{Id}(\mathbf{list}(\alpha))$$

We observe that the state relation computed in all test cases except the list reverse and map are the most precise. For example, with the function map that traverses a list and modifies only its `data` fields, the relation obtained is:

$$\text{Id}(\&l \mapsto \alpha) *_{\text{R}} [(\mathbf{listseg}(\alpha, \beta)) \dashrightarrow (\mathbf{listseg}(\alpha, \beta))]$$

This relation shows that both input and output lists start at the address α and end at the address β . This is not enough to prove that the lists contain the same addresses linked in the same order.

Second, we compare the runtime of the relational analysis and of the reachability analysis. We observe that the slow-down is at most 4x (reverse), and is about 2x in most cases. An exception is the list head insertion, which incurs no slowdown. This is due to the fact this analysis does not require computing an abstract join. While these test cases are not large, these results show that the analysis computing abstract state relations has a reasonable overhead compared to a classical analysis, yet it computes stronger properties. Furthermore, it would be more adapted to a modular interprocedural analysis.

7 Related Works

Our analysis computes an abstraction of the relational semantics of programs so as to capture the effect of a function or other blocks of code using an element of some specifically designed abstract domain. This technique has been applied to other abstractions in the past, and often applied to design *modular* static analyses [10], where program components can be analyzed once and separately. For numerical domains, it simply requires duplicating each variable into two instances respectively describing the old and the new value, and using a relational domain to the inputs and outputs. For instance, [22] implements this idea using convex polyhedra and so as to infer abstract state relations for numerical programs. It has also been applied to shape analyses based on Three Valued Logic [24] in [17]. This work is probably the closest to ours, but it relies on a very different abstraction using a TVLA whereas we use a set of abstract predicates based on separation logic. It uses the same variable duplication trick as mentioned above. Our analysis also has a notion of overlaid old / new predicates, but these are described heap regions, inside separation logic formulas. Desynchronized separation [11] also introduces a notion of overlaid state in separation logic, but does not support inductive predicates as our analysis does. Instead, it allows to reason on abstractions of JavaScript open objects seen as dictionaries. Also, [13,14] can express relations between heaps in different states using temporal logic extensions and automatas. In the context of functional languages, [18] allows to write down relations between function inputs and outputs, and relies on a solver to verify that constraints hold and [25] computes shape specifications by learning. Modular analyses that compute invariants by separate analysis of program components [6,12,4] use various sorts of abstractions for the behavior of program components. A common pattern is to use tables of couples made of an abstract pre-condition and a corresponding abstract post-condition, effectively defining a sort of cardinal power abstraction [9]. This technique has been used in several shape analyses based on separation logic [3,16,20,2]. We believe this tabular approach could benefit from abstractions of relations such as ours to infer stronger properties, and more concise summaries.

8 Conclusion

In this paper, we have introduced a set of logical connectives inspired by separation logic, to describe state relations rather than states. We have built upon this logic an abstract domain, and a static analysis based on abstract interpretation that computes conservative state relations. Experiments prove it effective for the analysis of basic data structure library functions.

Acknowledgements. We thank Arlen Cox for fruitful discussions, and Francois Berenger, Huisong Li, Jiangchao Liu and the anonymous reviewers for their comments on an earlier version of this paper. This work has received funding from the European Research Council under the EU’s seventh framework programme (FP7/2007-2013), grant agreement 278673, Project MemCAD, and from Bpifrance, grant agreement P3423-189738, FUI Project P-RC2.

References

1. Patrick Baudin, Jean-Christophe Filliâtre, Claude Marché, Benjamin Monate, Yannick Moy, and Virgile Prevosto. *Acsl: Ansi c specification language*, 2008.
2. Cristiano Calcagno, Dino Distefano, Peter O’Hearn, and Hongseok Yang. Footprint analysis : A shape analysis that discovers preconditions. In *Static Analysis Symposium (SAS)*, pages 402–418. Springer, 2007.
3. Cristiano Calcagno, Dino Distefano, Peter O’Hearn, and Hongseok Yang. Compositional shape analysis by means of bi-abduction. In *Symposium on Principles of Programming Languages (POPL)*, pages 289–300. ACM, 2009.
4. Ghila Castelnuovo, Mayur Naik, Noam Rinetzky, Mooly Sagiv, and Hongseok Yang. Modularity in lattices: A case study on the correspondence between top-down and bottom-up analysis. In *Static Analysis Symposium (SAS)*, pages 252–274. Springer, 2015.
5. Bor-Yuh Evan Chang and Xavier Rival. Relational inductive shape analysis. In *Symposium on Principles of Programming Languages (POPL)*, pages 247–260. ACM, 2008.
6. Ramkrishna Chatterjee, Barbara G Ryder, and William A Landi. Relevant context inference. In *Symposium on Principles of Programming Languages (POPL)*, pages 133–146. ACM, 1999.
7. P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Symposium on Principles of Programming Languages (POPL)*, pages 84–97. ACM, 1978.
8. Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Symposium on Principles of Programming Languages (POPL)*, 1977.
9. Patrick Cousot and Radhia Cousot. Systematic design of program analysis frameworks. In *Symposium on Principles of Programming Languages (POPL)*. ACM, 1979.
10. Patrick Cousot and Radhia Cousot. Modular static program analysis. In *Conference on Compiler Construction*, pages 159–179. Springer, 2002.
11. Arlen Cox, Bor-Yuh Evan Chang, and Xavier Rival. Desynchronized multi-state abstractions for open programs in dynamic languages. In *European Symposium on Programming (ESOP)*, pages 483–509. Springer, 2015.
12. Isil Dillig, Thomas Dillig, Alex Aiken, and Mooly Sagiv. Precise and compact modular procedure summaries for heap manipulating programs. In *Conference on Programming Language Design and Implementation (PLDI)*, pages 567–577. ACM, 2011.
13. Dino Distefano, Joost-Pieter Katoen, and Arend Rensik. Who is pointing when to whom? In *Foundations of Software Technology and Theoretical (FSTTCS)*, pages 250–262. Springer, 2004.
14. Dino Distefano, Joost-Pieter Katoen, and Arend Rensik. Safety and liveness in concurrent pointer programs. In *Formal Methods for Components and Objects (FMCO)*, pages 280–312. Springer, 2005.
15. Dino Distefano, Peter O’Hearn, and Hongseok Yang. A local shape analysis based on separation logic. In *Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 287–302. Springer, 2006.
16. Bhargav S Gulavani, Supratik Chakraborty, Ganesan Ramalingam, and Aditya V Nori. Bottom-up shape analysis. In *Static Analysis Symposium (SAS)*, pages 188–204. Springer, 2009.

17. Bertrand Jeannet, Alexey Loginov, Thomas Reps, and Mooly Sagiv. A relational approach to interprocedural shape analysis. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 32(2):5, 2010.
18. Gowtham Kaki and Suresh Jagannathan. A relational framework for higher-order shape analysis. In *International Colloquium on Function Programming*, pages 311–324. ACM, 2014.
19. Florent Kirchner, Nikolai Kosmatov, Virgile Prevosto, Julien Signoles, and Boris Yakobowski. Frama-c: A software analysis perspective. *Formal Aspects of Computing*, 27(3):573–609, 2015.
20. Quang Loc Le, Cristian Gherghina, Shengchao Qin, and Wei-Ngan Chin. Shape analysis via second-order bi-abduction. In *Conference on Computer Aided Verification (CAV)*, pages 52–68. Springer, 2014.
21. Gary T Leavens, Albert L Baker, and Clyde Ruby. Jml: a java modeling language. In *Formal Underpinnings of Java Workshop (at OOPSLA '98)*, pages 404–420, 1998.
22. Corneliu Popeea and Wei-Ngan Chin. Inferring disjunctive postconditions. In *ASIAN*, pages 331–345. Springer, 2006.
23. John Reynolds. Separation logic: A logic for shared mutable data structures. In *Symposium on Logics In Computer Science (LICS)*, pages 55–74. IEEE, 2002.
24. Mooly Sagiv, Thomas Reps, and Reinhard Wilhelm. Parametric shape analysis via 3-valued logic. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 24(3):217–298, 2002.
25. He Zhu, Gustavo Petri, and Suresh Jagannathan. Automatically learning shape specifications. In *Conference on Programming Language Design and Implementation (PLDI)*, pages 491–507. ACM, 2016.