# Automatic Inference of Heap Properties Exploiting Value Domains

Pietro Ferrara[1], Peter Müller[2], and Milos Novacek[2]

[1] IBM Thomas J. Watson Research Center, USA
pietroferrara@us.ibm.com
[2] Department of Computer Science, ETH Zurich, Switzerland
peter.mueller@inf.ethz.ch,milos.novacek@inf.ethz.ch

**Abstract.** Effective static analyses of heap-manipulating programs need to track precise information about the heap structures and the values computed by the program. Most existing heap analyses rely on manual annotations to precisely analyze general and, in particular, recursive, heap structures. Moreover, they either do not exploit value information to obtain more precise heap information or require more annotations for this purpose. In this paper, we present a combined heap and value analysis that infers complex invariants for recursive heap structures such as lists and trees, including relations between value fields of heap-allocated objects. Our analysis uses a novel notion of edge-local identifiers to track value information about the source and target of a pointer, even if these are summary nodes. With each potential pointer in the heap, our analysis associates value information that describes in which states the pointer may exist, and uses this information to improve the precision of the analysis by pruning infeasible heap structures. Our analysis has been implemented in the static analyzer Sample; experimental results show that it can automatically infer invariants for data structures, for which state-of-the-art analyses require manual annotations.

## 1 Introduction

Effective static analyses of heap-manipulating programs need to track precise information about the heap structures and the values computed by a program. Heap and value information is not independent: heap information determines which locations need to be tracked by a value analysis, and information about value fields may be useful to obtain more precise heap information, for instance, to rule out certain forms of aliasing. Moreover, many interesting invariants of heap-manipulating programs combine heap and value information such as the invariant that a heap structure is a sorted linked list.

Despite these connections, heap and value analyses have often been treated as orthogonal problems. Some existing heap analyses such as TVLA [17] rely on manual instrumentation to infer invariants that combine heap and value information. However, TVLA does not support general value domains, which limits, for instance, arithmetical reasoning. Recent work addresses this issue by combining TVLA with value domains, but still requires the user to provide predicates

to track and exchange information between the heap and value domains [21], or is not able to track complex invariants over recursive data structures [13]. Chang and Rival [4] present an efficient inference for combined heap and value invariants, which also relies on user-provided predicates. Other analyses do not require manual annotations [1,2], but are specific to programs that manipulate certain data structures such as singly-linked lists.

In this paper, we present a combined heap and value analysis—expressed as an abstract interpretation [7]—that infers complex invariants of heap structures. It is automatic in the sense that it uses only the information included in the program, without relying on manual annotations. Our analysis uses a graph-based abstraction of heaps, where each edge in the graph represents a *potential* pointer in the concrete heap. Each edge is associated with an abstract value state that characterizes in which concrete states this pointer might actually exist. The value states on the edges allow our analysis to represent disjunctive information in a single heap graph (like the bracketing constraints in Dillig et al.'s Fluid Updates [9]). They are also used to improve the precision of the analysis when value information implies that certain pointer chains cannot exist in concrete heaps. Our analysis can be instantiated with different value domains to obtain different trade-offs between precision and efficiency.

Like many heap analyses, we use summary nodes to abstract over sets of concrete objects. A key innovation of our analysis is to introduce *edge-local identifiers* for the source and target of each edge in the heap graph. An edge-local identifier represents a field of one particular concrete object, even when the object is abstracted by a summary node. By having identifiers per edge, the value analysis may relate the fields of the source and the target of a concrete pointer and, thus, track inductive invariants such as the sortedness of a linked list.

*Example.* Method increasingList in Fig. 1 creates and returns a linked list. If parameter v is non-positive, the list is empty, that is, result is null (invariant I1). Otherwise, the list satisfies the following properties: it is non-empty, that is, the result is non-null (I2), the first node has value 0 (I3), the values of all other nodes are one larger than their predecessor's (I4), and the value of the last node is $v-1$ (I5). Note that these invariants imply that the list is acyclic and has v nodes.

```
1 Node increasingList(int v) {
2   Node result = null;
3   int i = v;
4   while (i > 0) {
5     Node p = new Node();
6     p.next = result;
7     p.val = i - 1;
8     result = p;
9     i = i - 1;
10  }
11  return result;
12 }
```

Fig. 1: Running example.

Fig. 2 shows the abstract state that our analysis infers at the end of method increasingList. Here, we use a numerical domain such as Polyhedra [8] or Octagon [22] for the abstract states associated with each edge in the graph. The figure shows the relevant constraints from these states. They are expressed in terms of parameter v and the edge-local identifiers (Src, val) and (Trg, val), which refer to the val field of the source and target of a pointer, respectively.

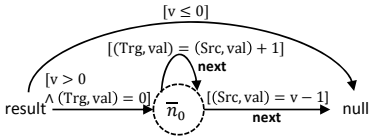


Fig. 2: The abstract heap state inferred at line 11 of Fig. 1.

The abstract state reflects the five invariants stated above. Variable result is null if the constraints on the corresponding edge hold, that is, if v is non-positive (I1). Otherwise, result points to the summary node $\overline{n}_0$, which implies that it is non-null (I2). This example illustrates that our analysis represents disjunctive information in a single graph: both possible values of result are represented by the same graph, and we use value information to determine the states in which each pointer may exist. The constraint $(\mathsf{Trg}, \mathsf{val}) = 0$ on the edge from result to the summary node $\overline{n}_0$ expresses that the first list node has value 0 (I3). Note that the edge-local identifier allows us to express properties of a single object, even if it is abstracted by a summary node. The same feature is used in the constraint $(\mathsf{Trg}, \mathsf{val}) = (\mathsf{Src}, \mathsf{val}) + 1$ on the edge from $\overline{n}_0$ to itself to express invariant I4. Finally, the constraint $(\mathsf{Src}, \mathsf{val}) = \mathsf{v} - 1$ on the edge from $\overline{n}_0$ to null expresses that the last list node has value $\mathsf{v} - 1$ (I5). All five invariants are inferred automatically by our analysis without manual annotations.

**Outline.** Sec. 2 defines the language and the concrete domain. Sec. 3 formalizes the abstract domain, while Sec. 5 defines the abstract semantics. Sec. 6 reports the experimental results, Sec. 7 discusses related work, and Sec. 8 concludes.

## 2 Programming Language and Concrete Domain

We present our analysis for the small object-based language in Fig. 3. To simplify the formalization, we model local variables as fields of a special object $\Im$, that is, treat local variables as heap locations. We distinguish *reference field* and *value field* access expressions rAE and vAE, depending on the type of the accessed field. A reference expression rexp may be null, a reference field access expression, or an object creation. A value expression vexp may be a literal, a value field access expression, or a binary expression. Since the treatment of loops and conditionals is standard, the only relevant statements in ST are value and reference assignments.

$$
\begin{aligned}
\mathsf{rAE} &::= \Im.\mathsf{f_r} \mid \mathsf{rAE}.\mathsf{f_r} \\
\mathsf{vAE} &::= \Im.\mathsf{f_v} \mid \mathsf{rAE}.\mathsf{f_v} \\
\mathsf{rexp} &::= \mathsf{null} \mid \mathsf{rAE} \mid \mathsf{new\ C} \\
\mathsf{vexp} &::= \mathsf{n} \mid \mathsf{vAE} \mid \mathsf{vexp}\ \langle \mathsf{op} \rangle\ \mathsf{vexp} \\
\mathsf{op} &::= + \mid - \mid * \mid \cdots \\
\mathsf{ST} &::= \mathsf{rAE} = \mathsf{rexp} \mid \mathsf{vAE} = \mathsf{vexp}
\end{aligned}
$$

Fig. 3: Expressions and statements.

In the concrete domain, we partition the content of heap locations into values and references. Let Ref be the set of concrete references (objects and null), with $\Im, \mathsf{null} \in \mathsf{Ref}$, and let Val be the set of values. Let $\mathsf{Field}_{\mathsf{Ref}}$ and $\mathsf{Field}_{\mathsf{Val}}$ be finite sets of reference and value fields, respectively. An execution state consists of a *value store* and a *reference store*. We model a value store as a partial map in $\mathsf{Store}_{\mathsf{Val}} = (\mathsf{Ref} \setminus \{\mathsf{null}\}) \times \mathsf{Field}_{\mathsf{Val}} \rightharpoonup \mathsf{Val}$ and a reference store as a partial map in $\mathsf{Store}_{\mathsf{Ref}} = (\mathsf{Ref} \setminus \{\mathsf{null}\}) \times \mathsf{Field}_{\mathsf{Ref}} \rightharpoonup \mathsf{Ref}$. For each reference in their domain, these maps contain an entry for every field. We will refer to entries in a reference store as *concrete edges*. We define the set of all concrete states (*concrete heaps*) as $\Sigma = \mathsf{Store}_{\mathsf{Ref}} \times \mathsf{Store}_{\mathsf{Val}}$.

## 3 Abstract Domain and Operators

In this section, we present the abstract domain, the concretization function, as well as join and widening operators.

### 3.1 Abstract Domain

Let $\overline{\mathsf{Ref}}$ be the set of *abstract references* (or *abstract nodes*) with $\Im, \mathsf{null} \in \overline{\mathsf{Ref}}$ (that is, we overload the symbols $\Im$ and $\mathsf{null}$ to denote both concrete and abstract references). Each abstract node $\overline{n} \in \overline{\mathsf{Ref}}$ represents either a single concrete non-null reference (*definite node*), or a non-empty set of concrete non-null references (*summary node*) with $\Im$ and $\mathsf{null}$ being definite nodes. The functions in $\mathsf{IsSummary} = \overline{\mathsf{Ref}} \to \{\mathsf{true}, \mathsf{false}\}$ define whether a node is a summary node.

An *abstract reference store* in $\overline{\mathsf{Store}}_{\overline{\mathsf{Ref}}} = \mathcal{P}((\overline{\mathsf{Ref}} \setminus \{\mathsf{null}\}) \times \mathsf{Field}_{\mathsf{Ref}} \times \overline{\mathsf{Ref}})$ represents possible pointers between abstract nodes through reference fields. It can be interpreted as a directed graph where edges are labeled with a field name. Hence, we will refer to members of the abstract reference store as *abstract edges*. For an abstract edge $\overline{n}_1 \xrightarrow{f_r} \overline{n}_2$, we will refer to $\overline{n}_1$ as the *source* and to $\overline{n}_2$ as the *target* of the edge.

Our heap analysis is parameterized by an *abstract value domain* $\overline{\mathsf{V}}$, which tracks information about value fields, for instance, relations among numerical values. Each abstract edge is associated with an abstract value state (*abstract condition*) via a map in $\overline{\mathsf{Cond}} = (\overline{\mathsf{Ref}} \setminus \{\mathsf{null}\}) \times \mathsf{Field}_{\mathsf{Ref}} \times \overline{\mathsf{Ref}} \to \overline{\mathsf{V}}$. The abstract condition of an abstract edge approximates the concrete value stores in which the edge exists. That is, our abstract domain tracks disjunctive information by having several edges with the same source and field, and associating them with different abstract conditions.

Abstract value states in $\overline{\mathsf{V}}$ refer to memory locations via *abstract identifiers* $\overline{\mathsf{ID}} = \overline{\mathsf{Loc}} \cup \overline{\mathsf{EId}}$ where $\overline{\mathsf{Loc}} = (\overline{\mathsf{Ref}} \setminus \{\mathsf{null}\}) \times \mathsf{Field}_{\mathsf{Val}}$ and $\overline{\mathsf{EId}} = \{\mathsf{Src}, \mathsf{Trg}\} \times \mathsf{Field}_{\mathsf{Val}}$. An identifier $(\overline{n}, f_v) \in \overline{\mathsf{Loc}}$ represents the value field $f_v$ of the concrete references abstracted by the node $\overline{n}$. *Edge-local identifiers* $(\mathsf{Src}, f_v), (\mathsf{Trg}, f_v) \in \overline{\mathsf{EId}}$ represent the value field $f_v$ of the *single* concrete *source* or *target* reference of the concrete edges represented by an abstract edge. They track relations between the value fields of adjacent references in concrete heaps, which allows us to infer precise invariants on summary nodes. For instance, the constraint $(\mathsf{Src}, \mathsf{val}) \leq (\mathsf{Trg}, \mathsf{val})$ in the abstract condition of an abstract edge $\overline{n} \xrightarrow{\mathsf{next}} \overline{n}$ expresses sortedness of the concrete list that is abstracted by the summary node $\overline{n}$.

We define the set of all abstract states (*abstract heaps*) as $\overline{\Sigma} = \overline{\mathsf{Store}}_{\overline{\mathsf{Ref}}} \times \overline{\mathsf{Cond}} \times \mathsf{IsSummary}$.

*Example.* The abstract heap in Fig. 4 depicts the loop invariant of the program in Fig. 1. Many of the constraints are similar to the constraints in Fig. 2. In particular, combining the abstract heap for the loop invariant with the negation of the loop guard (that is, $\mathsf{i} \leq 0$) yields the information reflected in Fig. 2, for instance, that $\mathsf{result}$ is null iff $\mathsf{v} \leq 0$ and that the first list node has value 0.

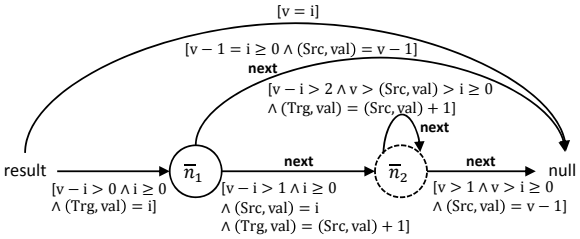


Fig. 4: The abstract heap representing the loop invariant at line 4 of the example in Fig. 1. Solid and dashed circles denote definite and summary nodes, respectively. Arrows depict abstract edges and are annotated with relevant constraints from their abstract conditions. To improve readability, we depict local reference variables as nodes and use local variables as identifiers in constraints, although the analysis models them as fields of $\Im$.

### 3.2 Concretization

In this section, we define the concretization function $\gamma : \overline{\Sigma} \to \mathcal{P}(\Sigma)$ that yields the set of concrete heaps represented by a given abstract heap.

We assume that our heap analysis is instantiated with a sound value analysis. Its concretization function $\gamma_{\overline{\mathsf{V}}} : \overline{\mathsf{V}} \to \mathcal{P}(\overline{\mathsf{ID}} \to \mathcal{P}(\mathsf{Val}))$ yields a set of maps from abstract identifiers to sets of concrete values. These maps yield sets of concrete values rather than single values since an abstract value state may contain identifiers for fields of summary nodes, and the value analysis alone cannot concretize them. Let $\mathit{references}(st_{\mathsf{Ref}})$ be the set of concrete references of a given concrete reference store $st_{\mathsf{Ref}}$, including $\Im$ and $\mathsf{null}$. We define the concretization function $\gamma$ of abstract heaps as:

$$(st_{\mathsf{Ref}}, st_{\mathsf{Val}}) \in \gamma(\overline{St}, \overline{con}, \mathit{isSum}) \Leftrightarrow \left( \begin{array}{l} \exists \alpha_{\mathsf{Ref}} \in (\mathit{references}(st_{\mathsf{Ref}}) \to \overline{\mathsf{Ref}}) \cdot \\ \quad \mathit{GraphEmbed}(\alpha_{\mathsf{Ref}}, st_{\mathsf{Ref}}, (\overline{St}, \mathit{isSum})) \wedge \\ \quad \mathit{ValueEmbed}(\alpha_{\mathsf{Ref}}, (st_{\mathsf{Ref}}, st_{\mathsf{Val}}), \overline{con}) \end{array} \right)$$

That is, a concrete heap $(st_{\mathsf{Ref}}, st_{\mathsf{Val}})$ is in the concretization of an abstract heap $(\overline{St}, \overline{con}, \mathit{isSum})$ iff there exists an embedding $\alpha_{\mathsf{Ref}}$ (a function from concrete references to abstract nodes) such that the *shape* and the *values* of the concrete heap can be embedded into the abstract heap. These embeddings are expressed via the predicates *GraphEmbed* and *ValueEmbed*, which are defined as follows.

*GraphEmbed* holds if a given concrete reference store matches the shape of a given abstract heap, ignoring the value information. This is the case if $\Im$ and $\mathsf{null}$ are the only concrete references that are abstracted to their abstract counterparts $\Im$ and $\mathsf{null}$ (1), if, whenever multiple concrete references are abstracted to a single abstract reference, that abstract reference is a summary node (2), and if every concrete edge is represented by a (unique) abstract edge in the abstract heap (3):

$$\begin{aligned} &\mathit{GraphEmbed}(\alpha_{\mathsf{Ref}}, st_{\mathsf{Ref}}, (\overline{St}, \mathit{isSum})) \Leftrightarrow \\ &\quad \alpha^{-1}_{\mathsf{Ref}}(\Im) = \{\Im\} \ \wedge \ \alpha^{-1}_{\mathsf{Ref}}(\mathsf{null}) = \{\mathsf{null}\} \ \wedge && (1) \\ &\quad (\forall \overline{n} \in \mathit{img}(\alpha_{\mathsf{Ref}}) \cdot |\alpha^{-1}_{\mathsf{Ref}}(\overline{n})| > 1 \Rightarrow \mathit{isSum}(\overline{n})) \ \wedge && (2) \\ &\quad \forall r_1 \xrightarrow{f_r} r_2 \in st_{\mathsf{Ref}} \cdot \alpha_{\mathsf{Ref}}(r_1) \xrightarrow{f_r} \alpha_{\mathsf{Ref}}(r_2) \in \overline{St} && (3) \end{aligned}$$

where $\alpha^{-1}_{\mathsf{Ref}}$ is the *preimage* of $\alpha_{\mathsf{Ref}}$ (that is, it yields the set of concrete references abstracted by a given abstract reference).

*ValueEmbed* expresses that, for a given concrete reference store $st_{\mathsf{Ref}}$, the value store $st_{\mathsf{Val}}$ matches all relevant abstract conditions in the abstract heap. Here, an abstract condition is *relevant* if it is associated with an abstract edge that corresponds to a concrete edge in $st_{\mathsf{Ref}}$. In the definition below, we relate the concrete value store $st_{\mathsf{Val}}$ to each relevant abstract condition via a map $s$ from abstract identifiers to sets of concrete values. For each concrete edge $r_1 \xrightarrow{f_r} r_2$ in the concrete reference store $st_{\mathsf{Ref}}$, there is a map $s$ in the concretization of the abstract condition of the corresponding abstract edge (4). The map $s$ may constrain a concrete location $(r, f_v)$ in three ways: via the abstract identifier $(\alpha_{\mathsf{Ref}}(r), f_v)$, via the edge-local identifier $(\mathsf{Src}, f_v)$ if $r$ is the source of the concrete edge, that is, $r = r_1$, and via the edge-local identifier $(\mathsf{Trg}, f_v)$ if $r$ is the target of the concrete edge, that is, $r = r_2$. In all three cases, the map $s$ must yield a set that contains the value $v$ stored in the concrete value store for $(r, f_v)$ (5). Finally, any concrete value store matches the relevant abstract conditions only if the conditions do not contradict each other, even on abstract locations that are not included in a given concrete heap. To ensure there are no such contradictions, $s$ must be in the concretization of *all* relevant conditions, ignoring edge-local identifiers, which may denote different locations for different abstract edges. We use the operator $\downarrow_{\overline{\mathsf{Loc}}}$ to project to the identifiers in $\overline{\mathsf{Loc}}$, that is, to remove edge-local identifiers (6).

$$
\begin{aligned}
&\mathit{ValueEmbed}(\alpha_{\mathsf{Ref}}, (st_{\mathsf{Ref}}, st_{\mathsf{Val}}), \overline{con}) \Leftrightarrow \\
&\quad \forall r_1 \xrightarrow{f_r} r_2 \in st_{\mathsf{Ref}} \cdot \exists s \in \gamma_{\overline{\mathsf{V}}}(\overline{con}(\alpha_{\mathsf{Ref}}(r_1) \xrightarrow{f_r} \alpha_{\mathsf{Ref}}(r_2))) \cdot && (4)\\
&\qquad \forall((r, f_v) \mapsto v) \in st_{\mathsf{Val}} \cdot \begin{pmatrix} v \in s(\alpha_{\mathsf{Ref}}(r), f_v) \wedge \\ r = r_1 \Rightarrow v \in s(\mathsf{Src}, f_v) \wedge \\ r = r_2 \Rightarrow v \in s(\mathsf{Trg}, f_v) \end{pmatrix} \wedge && (5)\\
&\qquad s \downarrow_{\overline{\mathsf{Loc}}} \in \gamma_{\overline{\mathsf{V}}} \left( \prod_{r'_1 \xrightarrow{f'_r} r'_2 \in st_{\mathsf{Ref}}} \left( \overline{con}(\alpha_{\mathsf{Ref}}(r'_1) \xrightarrow{f'_r} \alpha_{\mathsf{Ref}}(r'_2)) \downarrow_{\overline{\mathsf{Loc}}} \right) \right) && (6)
\end{aligned}
$$

*Example.* Fig. 5 shows the reference and value stores of two concrete heaps. The heap of the left) is in the concretization of the abstract heap in Fig. 2. For the embedding $\alpha_{\mathsf{Ref}} = [\Im \mapsto \Im, \mathsf{null} \mapsto \mathsf{null}, r_1 \mapsto \overline{n}_0, r_2 \mapsto \overline{n}_0]$, *GraphEmbed* holds since $\overline{n}_0$ is a summary node and all three concrete edges have corresponding abstract edges. *ValueEmbed* also holds since the concrete value store satisfies the three relevant abstract conditions, and these conditions do not contradict each other.

In contrast, the heap on the right is *not* in the concretization of the abstract heap in Fig. 2. The graph embedding forces the embedding to be $\alpha_{\mathsf{Ref}} = [\Im \mapsto \Im, \mathsf{null} \mapsto \mathsf{null}, r_1 \mapsto \overline{n}_0]$. Therefore, both edge-local identifiers $(\mathsf{Src}, \mathsf{val})$ and $(\mathsf{Trg}, \mathsf{val})$ on the abstract edge from $\overline{n}_0$ to $\overline{n}_0$ correspond to $(r_1, \mathsf{val})$, such that there is no value for $(r_1, \mathsf{val})$ that satisfies the constraint $(\mathsf{Trg}, \mathsf{val}) = (\mathsf{Src}, \mathsf{val})+1$.

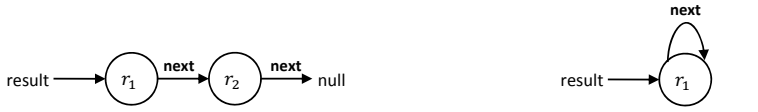


Fig. 5: Concrete heaps, consisting of a reference store, displayed on top, and a value store, displayed underneath. The heap on the left is in the concretization of the abstract heap in Fig. 2, whereas the heap on the right is not because it violates the condition that list nodes store increasing values.

In other words, any map $s$ in the concretization of this constraint assigns different values to these edge-local identifiers and, thus, does not satisfy condition (5).

### 3.3 Join

The *join* operator $\sqcup_{\overline{\Sigma}}$ first computes an abstract reference store for the joined heaps and then the abstract conditions for the edges in this store.

**Abstract Reference Store.** An abstract heap can be viewed as a directed graph in which vertices are labeled as $\Im$, $\mathsf{null}$, definite node other than $\Im$ and $\mathsf{null}$, or summary node; edges are labeled with reference fields. The vertex labels are used to avoid matching nodes in two heaps that cannot correspond (for instance, a summary node and a definite node). A *labeled heap graph* is a triple $g = (V, E, \eta) \in \mathsf{Graph}$, where $V \subseteq \overline{\mathsf{Ref}}$ is a set of vertices, $E \subseteq V \times \mathsf{Field}_{\mathsf{Ref}} \times V$ is a set of edges labeled with a reference field, and $\eta : V \to \{\Im, \mathsf{Null}, \mathsf{Def}, \mathsf{Sum}\}$ is a labeling function on vertices. We assume a *strict total order* $<_{\mathsf{G}}$ on graphs that ensures in particular that $g_1 <_{\mathsf{G}} g_2$ if $g_1$ has fewer vertices than $g_2$ or the same number of vertices but fewer edges.

To improve performance, we define the join of two abstract reference stores such that it minimizes the size of the resulting store. Its structure is the minimum common supergraph of the two joined stores. Let $g_1$ and $g_2$ be graphs. Graph $g$ is a *common supergraph* of $g_1$ and $g_2$ iff $g_1$ and $g_2$ are subgraph isomorphic to $g$ with the isomorphisms $\mathcal{I}_1$ and $\mathcal{I}_2$, respectively. We call $g$ the *minimum common supergraph* (MCS) of $g_1$ and $g_2$ if there exists no other common supergraph that is smaller in the ordering $<_{\mathsf{G}}$. The procedure $MCS(g_1, g_2)$ yields the (unique) minimum common supergraph $g$ of $g_1$ and $g_2$ as well as the corresponding subgraph isomorphisms $\mathcal{I}_1$ and $\mathcal{I}_2$ between $g_1$ and $g$, and $g_2$ and $g$, respectively. The problem of computing $MCS$ can be reduced to the well-studied problem of finding the maximum common subgraph [3]. See Appendix A for the definitions of graph/subgraph isomorphism and maximum common subgraph. Intuitively, we can compute $(g, \mathcal{I}_1, \mathcal{I}_2) = MCS(g_1, g_2)$ by "gluing" to the maximum common subgraph of $g_1$ and $g_2$ those parts of $g_1$ and $g_2$ that are not in their maximum common subgraph.

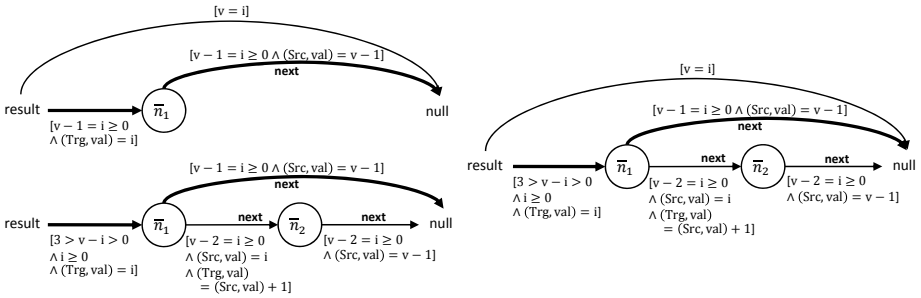


Fig. 6: The abstract heap graphs on the left occur before and after the second iteration of the fixed-point computation of the loop invariant (line 4) in Fig. 1. Joining them results in the heap on the right. Bold arrows indicate edges of the maximum common subgraph.

Let $\Pi : \overline{\mathsf{Store}}_{\overline{\mathsf{Ref}}} \times \mathsf{IsSummary} \to \mathsf{Graph}$ be a *bijective* function from abstract stores to heap graphs. The abstract store and the IsSummary function of the join of $(\overline{St}_1, \overline{con}_1, isSum_1)$ and $(\overline{St}_2, \overline{con}_2, isSum_2)$ are $(\overline{St}, isSum) = \Pi^{-1}(MCS(\Pi(\overline{St}_1, isSum_1), \Pi(\overline{St}_2, isSum_2)) \downarrow_1)$, where $\downarrow_1$ denotes projection of a tuple on the first component, that is, the graph returned by *MCS*. Note that $\Pi(\overline{St}, isSum)$ includes both $\Pi(\overline{St}_1, isSum_1)$ and $\Pi(\overline{St}_2, isSum_2)$. Hence, the abstract reference store $(\overline{St}, isSum)$ subsumes the abstract reference stores $(\overline{St}_1, isSum_1)$ and $(\overline{St}_2, isSum_2)$.

*Example.* Fig. 6 shows on the left the abstract heap graphs $g_1$ and $g_2$ before and after the second iteration of the fixed-point computation of the loop invariant (line 4) in Fig. 1, and their join $g$ on the right. Besides the special nodes $\Im$ and null, the maximum common subgraph includes node $\overline{n}_1$ as well as the edges from result to $\overline{n}_1$ and from $\overline{n}_1$ to null. To this common subgraph, we add the remainder of $g_1$ (the edge from result to null) and the remainder of $g_2$ ($\overline{n}_2$ with its edges). Note that both $g_1$ and $g_2$ are subgraph isomorphic to $g$, where the isomorphism is the identity function.

**Abstract Conditions.** Consider an edge in the abstract store resulting from the join of two abstract heaps. We determine its abstract condition as follows. If the edge is in the maximum common subgraph of the joined heap graphs, its abstract condition is the join of the abstract conditions in the two heaps. Otherwise, the condition is the same as in the heap that contributed the edge, after applying the subgraph isomorphism.

As explained above, computing the minimum common supergraph $(g, \mathcal{I}_1, \mathcal{I}_2) = MCS(\Pi(\overline{St}_1, isSum_1), \Pi(\overline{St}_2, isSum_2))$ yields the subgraph isomorphisms $\mathcal{I}_1$ and $\mathcal{I}_2$ from $\Pi(\overline{St}_1, isSum_1)$ to $g$ and from $\Pi(\overline{St}_2, isSum_2)$ to $g$, respectively. We define the function $rename_{\mathsf{ISO}} : (\overline{\mathsf{Ref}} \to \overline{\mathsf{Ref}}) \times \overline{\mathsf{V}} \to \overline{\mathsf{V}}$ to rename the identifiers in $\overline{\mathsf{Loc}}$ of a given abstract value state according to an isomorphism. Using this renaming, we define the join operator $\sqcup_{\overline{\Sigma}} : \overline{\Sigma} \times \overline{\Sigma} \to \overline{\Sigma}$ as

$$(\overline{St}_1, \overline{con}_1, isSum_1) \sqcup_{\overline{\Sigma}} (\overline{St}_2, \overline{con}_2, isSum_2) = (\overline{St}, \overline{con}, isSum)$$

where:

$$(g, \mathcal{I}_1, \mathcal{I}_2) = MCS(\Pi(\overline{St}_1, isSum_1), \Pi(\overline{St}_2, isSum_2)) \wedge$$
$$(\overline{St}, isSum) = \Pi^{-1}(g) \wedge$$
$$\overline{con} = \left[ \overline{e} \mapsto \overline{s} \;\middle|\; \overline{e} \in \overline{St} \;\wedge\; \overline{s} = \bigsqcup \left\{ \overline{s}' \;\middle|\; \begin{array}{l} \exists\, i \in \{1,2\} \cdot \exists\, (\overline{n}_1, f_r, \overline{n}_2) \in \overline{St}_i \cdot \\ \overline{e} = (\mathcal{I}_i(\overline{n}_1), f_r, \mathcal{I}_i(\overline{n}_2)) \wedge \\ \overline{s}' = rename_{\mathsf{ISO}}(\mathcal{I}_i, \overline{con}_i(\overline{n}_1, f_r, \overline{n}_2)) \end{array} \right\} \right]$$

Computing the maximum common subgraph is NP-complete; however, most code fragments change only small portions of the abstract heap. Our implementation exploits this fact to compute the isomorphisms incrementally, usually in linear time.

*Example.* Consider the edge from result to $\overline{n}_1$ in the heap on the right of Fig. 6, which is in the maximum common subgraph of the heaps on the left. Hence, its abstract condition is the join of the conditions for those heaps (assuming a relational numerical domain). Since the constraint $\mathsf{v}-1 = \mathsf{i}$ in the top left abstract heap implies the constraint $3 > \mathsf{v} - \mathsf{i} > 0$ in the lower heap, the latter constraint is tracked by the result of the join operation; the other constraints of the joined conditions are identical and, thus, carried over to the result. Conversely, the edges from result to null, from $\overline{n}_1$ to $\overline{n}_2$, and from $\overline{n}_2$ to null are not in the maximum common subgraph; their conditions come from the heap contributing the edges.

### 3.4 Widening

The above join operator does not guarantee the convergence of the analysis. In fact, the size of the abstract heap may grow at each application of join, and the abstract conditions may not stabilize. Therefore, we define a *widening* operator $\nabla_{\overline{\Sigma}} : \overline{\Sigma} \times \overline{\Sigma} \to \overline{\Sigma}$ that guarantees that the analysis reaches a fixed point in finite time (that is, terminates). In order to do so, the widening operator must bound the size of the abstract heap, which means that it has to merge nodes into summary nodes. This merging is controlled via a finite set of field access expressions $\mathcal{W}$, which is a parameter of the analysis and denotes references that the analysis should track separately. By default, $\mathcal{W}$ is the set of local reference variables, but it can be extended to any set of field access expressions if desired. For all examples in our evaluation (Sec. 6), the analysis uses the default.

We perform widening in two steps. First, in both input heaps, we merge nodes that (i) are denoted by the same set of field access expressions from $\mathcal{W}$, *and* (ii) are reachable (via some access path) from the same set of local variables. Second, if the two heaps are isomorphic, we apply edge-wise widening to the abstract conditions; otherwise, we join them.

**Merging Nodes.** Our widening operator is parametric in a finite set of access expressions $\mathcal{W} \subseteq \texttt{rAE}$, which denote objects that the user wants the analysis to track separately. Using access expressions for obtaining a bounded heap representation is standard [17,19,23]. Different heuristics and data-flow analyses can

be used to automatically determine $\mathcal{W}$. For instance, for all examples in our evaluation (Sec. 6), we used the set of local reference variables.

Based on this set of access expressions, we partition the nodes of an abstract heap into equivalence classes such that two nodes are in the same equivalence class iff they are reached via the same set of access expressions from $\mathcal{W}$ *and* they are reachable from the same set of local variables. The latter condition improves precision, for instance, by preventing the widening from merging the tails of two separate lists into a single summary node. We define special equivalence classes for $\Im$ and $\mathsf{null}$.

Let $\mathit{pointedBy}_{\mathcal{W}} : (\overline{\mathsf{Ref}} \times \overline{\mathsf{Store}_{\overline{\mathsf{Ref}}}}) \to \mathcal{P}(\mathcal{W})$ return the set of access expressions in $\mathcal{W}$ leading to a given node in a given abstract reference store. Let $\mathit{reach} : \overline{\mathsf{Ref}} \times \overline{\mathsf{Store}_{\overline{\mathsf{Ref}}}} \to \mathcal{P}(\mathsf{Field}_{\mathsf{Ref}})$ be a function that returns the set of reference fields representing local variables from which a give node is reachable in a given abstract store. Formally:

$$\mathit{reach}(\overline{n}, \overline{St}) = \left\{ f \;\middle|\; \begin{array}{c} \exists\ ae \in \mathtt{AE}_{\Im}, \langle \overline{e}_1, \cdots, \overline{e}_n \rangle \in \mathit{paths}(ae, \overline{St})\ . \\ \mathit{target}(\overline{e}_n) = \overline{n}\ \wedge \mathit{field}(e)_1 = f \end{array} \right\}$$

Let $\mathit{partition} : \overline{\mathsf{Store}_{\overline{\mathsf{Ref}}}} \to \mathcal{P}(\mathcal{P}(\overline{\mathsf{Ref}}))$ be the a function that partitions the nodes of a given abstract store into equivalence classes as described above. The partitioning satisfies the following properties: (1) the partitions cover all nodes in the abstract store and there are separate partitions for $\Im$ and $\mathsf{null}$, (2) the partitions are non-empty and do not overlap, (3) all and only the nodes (other than $\Im$ and $\mathsf{null}$) nodes pointed by the same access expressions from $\mathcal{W}$ and reachable from the same local variables are in the same partition. Note, that the above conditions define a *unique* partitioning and the number of partitions is bounded by the size of $\mathcal{W}$ and the number of local variables. Formally: $\mathit{partition}(\overline{St}) = \{\overline{N}_1, \cdots, \overline{N}_k\}$ where

$$\bigcup_{1 \le i \le k} \overline{N}_i = \mathit{heapNodes}(\overline{St})\ \wedge\ \overline{N}_1 = \{\Im\}\ \wedge\ \overline{N}_2 = \{\mathsf{null}\}\ \wedge \tag{1}$$

$$(\forall\ 1 \le i, j \le k\ .\ \overline{N}_i \neq \emptyset\ \wedge\ (i \neq j \Rightarrow \overline{N}_i \cap \overline{N}_j = \emptyset))\ \wedge \tag{2}$$

$$\forall\ 3 \le i, j \le k, \overline{n}_1 \in \overline{N}_i, \overline{n}_2 \in \overline{N}_j\ .$$

$$i = j \iff \begin{pmatrix} \mathit{pointedBy}_{\mathcal{W}}(\overline{n}_1, \overline{St}) = \mathit{pointedBy}_{\mathcal{W}}(\overline{n}_2, \overline{St})\ \wedge \\ \mathit{reach}(\overline{n}_1, \overline{St}) = \mathit{reach}(\overline{n}_2, \overline{St}) \end{pmatrix} \tag{3}$$

After partitioning the nodes of an abstract heap, we merge all nodes in the same partition into a single (if necessary, summary) node. The merging of nodes may also replace several edges by a single edge from or to the merged node. The abstract condition for this edge is the join of the conditions of the edges it replaces.

We merge the nodes in the same partition into a single node, namely the smallest node of the partition, guaranteeing that the nodes from two different partitions are not merged into the same node. Let $\mathit{min} : \mathcal{P}(\overline{\mathsf{Ref}}) \to \overline{\mathsf{Ref}}$ return the smallest node from a given set of nodes. We also *fold* abstract identifiers in the abstract conditions of the heap corresponding to the merged nodes. Let

$getFolds : \overline{\mathsf{Store}_{\overline{\mathsf{Ref}}}} \to \mathcal{P}(\mathcal{P}(\overline{\mathsf{Loc}}) \times \overline{\mathsf{Loc}})$ be a function that describes how the abstract identifiers should be folded. Recall, that $fold(R, \overline{s})$ returns a value state in which for any $(\overline{I}, \overline{i}) \in R$, identifiers in $\overline{I}$ are replaced with $\overline{i}$ and the value of $\overline{i}$ is the least upper-bound of the values of identifiers in $\overline{I}$. Note, that we do not need to take care of the edge-local identifiers. Formally:

$$\begin{aligned} getFolds(\overline{St}) = \{(R, (min(\overline{N}_i), f_v)) \mid \{\overline{N}_1, \cdots, \overline{N}_k\} = partition(\overline{St}) \wedge \\ f_v \in \mathsf{Field}_{\mathsf{Val}} \ \wedge \ 1 \leq i \leq k \ \wedge \ R = \{(\overline{n}_i, f_v) \in \overline{\mathsf{Loc}} \mid \overline{n}_i \in \overline{N}_i\} \ \} \end{aligned}$$

Now, we are in the position to define a function $merge_{\overline{\Sigma}} : \overline{\Sigma} \to \overline{\Sigma}$ that formalizes what it means to merge nodes of an abstract heap with respect to given $\mathcal{W}$. First, we compute the map $\phi$ that describes how the nodes are merged, and subsequently, we merge the nodes in the abstract store. Then, we compute which abstract identifiers of abstract conditions should be merged and merge them. Merging nodes may also lead to merging of several edges. Therefore, we need to join the conditions of merged edges. Finally, we update the map defines nodes as summary or definite. If the partition contained only a single definite node, the node stays definite. If we have merged several nodes or the partition already contained a summary node, the node into which we have merged the partition is a summary node. Formally: $merge_{\overline{\Sigma}}(\overline{St}, \overline{con}, isSum) = (\overline{St}', \overline{con}', isSum')$ where

$$\begin{aligned} \phi &= \left[ \overline{n} \mapsto min(\overline{N}_i) \,\middle|\, \begin{array}{l} \{\overline{N}_1, \cdots, \overline{N}_k\} = partition(\overline{St}) \wedge \\ \exists\ 1 \leq i \leq k \ .\ \overline{n} \in \overline{N}_i \end{array} \right] \ \wedge \\ \overline{St}' &= \{(\phi(\overline{n}_1), f_r, \phi(\overline{n}_2)) \mid (\overline{n}_1, f_r, \overline{n}_2) \in \overline{St}\} \wedge \\ \overline{con}' &= [(\overline{n}'_1, f_r, \overline{n}'_2) \mapsto \overline{s} \mid (\overline{n}'_1, f_r, \overline{n}'_2) \in \overline{St}' \ \wedge \\ &\qquad \overline{s} = \bigsqcup\{fold(getFolds(\overline{St}), \overline{con}(\overline{n}_1, f_r, \overline{n})) \mid \\ &\qquad\quad (\overline{n}_1, f_r, \overline{n}_2) \in \overline{St} \ \wedge \ \phi(\overline{n}_1) = \overline{n}'_1 \ \wedge \ \phi(\overline{n}_2) = \overline{n}'_2\}] \wedge \\ isSum' &= \left[ \overline{n} \mapsto \neg \left( \begin{array}{c} \exists\ 1 \leq i \leq k \ .\{\overline{n}\} = \overline{N}_i \\ \wedge\ \neg isSum(\overline{n}) \end{array} \right) \,\middle|\, \begin{array}{l} \overline{n} \in heapNodes(\overline{St}') \wedge \\ \{\overline{N}_1, \cdots, \overline{N}_k\} = partition(\overline{St}) \end{array} \right] \end{aligned}$$

**From Merge to Widening.** Once we merged nodes in the heaps to which we want to apply widening, we need to perform edge-wise widening of the abstract conditions of the merged heaps. The widening of the conditions on the edges of the heaps can be performed only if the merged heaps are structurally the same (isomorphic), otherwise there are edges in one heap that do not correspond to any edge in the other heap, hence, their conditions can not be widened. Therefore, we need to be able to realize whether two merged heaps are isomorphic.

We observe that if $\mathcal{I}$ is an isomorphism from a heap $(\overline{St}_2, \overline{con}_2, isSum_2)$ to a heap $(\overline{St}_1, \overline{con}_1, isSum_1)$, then for any $\overline{n} \in dom(\mathcal{I})$, the nodes $\overline{n}$ and $\mathcal{I}(\overline{n})$ must be reachable from the same set of local variables and the same set of access expressions in $\mathcal{W}$ must lead to them in their respective abstract stores. Furthermore, it must be the case that $isSum_2(\overline{n}) = isSum_1(\mathcal{I}(\overline{n}))$ and that if $\overline{n} \in \{\Im, \mathsf{null}\}$ then $\overline{n} = \mathcal{I}(\overline{n})$. Due to the uniqueness of the partitioning of nodes of an abstract heap, any two different nodes of the abstract heap $merge_{\overline{\Sigma}}(\overline{\sigma})$

other than $\Im$ and $\mathsf{null}$ are reachable from different sets of local variables, or different sets of access expressions in $\mathcal{W}$ lead to them. Therefore, we can find a candidate for an isomorphism between two heaps in which we have merged the nodes just by looking at the set of access expressions from $\mathcal{W}$ leading to the node of the heaps and local variables from which the nodes are reachable. Let $\mathit{find}_{\mathsf{ISO}} : ((\overline{\mathsf{Store}_{\overline{\mathsf{Ref}}}} \times \mathsf{IsSummary}) \times (\overline{\mathsf{Store}_{\overline{\mathsf{Ref}}}} \times \mathsf{IsSummary})) \to (\overline{\mathsf{Ref}} \to \overline{\mathsf{Ref}})$ return a possible witness (candidate) for the isomorphism from the second abstract store to the first. Formally:

$$\mathit{find}_{\mathsf{ISO}}((\overline{St}_1, \mathit{isSum}_1), (\overline{St}_2, \mathit{isSum}_2)) = \left[ \overline{n}_2 \mapsto \overline{n}_1 \;\middle|\; \begin{array}{l} \overline{n}_1 \in \mathit{heapNodes}(\overline{St}_1) \;\wedge\; \overline{n}_2 \in \mathit{heapNodes}(\overline{St}_2) \;\wedge \\ \mathit{isSum}_1(\overline{n}_1) = \mathit{isSum}_2(\overline{n}_2) \;\wedge\; (\overline{n}_2 \in \{\Im, \mathsf{null}\} \Rightarrow \overline{n}_2 = \overline{n}_1) \;\wedge \\ \mathit{pointedBy}_{\mathcal{W}}(\overline{n}_1, \overline{St}) = \mathit{pointedBy}_{\mathcal{W}}(\overline{n}_2, \overline{St}) \;\wedge \\ \mathit{reach}(\overline{n}_1, \overline{St}) = \mathit{reach}(\overline{n}_2, \overline{St}) \end{array} \right]$$

In contrast with NP-complete problem of searching for isomorphisms between two arbitrary graphs, the complexity of $\mathit{find}_{\mathsf{ISO}}$ is linear in the number of nodes in given abstract stores and this number is bound by the size of $\mathcal{W}$ and the number of local variables. Hence, once we fix $\mathcal{W}$ and the set of local variables, $\mathit{find}_{\mathsf{ISO}}$ is of constant complexity. Furthermore, if $\mathit{find}_{\mathsf{ISO}}((\overline{St}_1, \mathit{isSum}_1), (\overline{St}_2, \mathit{isSum}_2))$ is not the isomorphism from $\overline{St}_2$ to $\overline{St}_1$ (i.e. not every node from one heap is mapped to a node in the other heap), then there exists no such isomorphism.

Even if the abstract stores of the merged heaps are structurally the same, we might not be able to perform widening on the abstract conditions of the corresponding edges, as the widening of value states is defined only for states that contain the same set of abstract identifiers. Hence, we define a function $\mathit{compatible} : (\overline{\Sigma} \times \overline{\Sigma}) \to \{\mathsf{true}, \mathsf{false}\}$ that checks whether the merged heaps are isomorphic and whether the abstract conditions of corresponding edges contain the same abstract identifiers with respect to the isomorphism. Formally:

$$\begin{array}{l} \mathit{compatible}((\overline{St}_1, \overline{\mathit{con}}_1, \mathit{isSum}_1), (\overline{St}_2, \overline{\mathit{con}}_2, \mathit{isSum}_2)) = \exists\, \mathcal{I} : \overline{\mathsf{Ref}} \to \overline{\mathsf{Ref}} \;. \\ \quad \mathcal{I} = \mathit{find}_{\mathsf{ISO}}((\overline{St}_1, \mathit{isSum}_1), (\overline{St}_2, \mathit{isSum}_2)) \;\wedge\; \mathit{dom}(\mathcal{I}) = \mathit{heapNodes}(\overline{St}_2) \;\wedge \\ \quad \mathit{img}(\mathcal{I}) = \mathit{heapNodes}(\overline{St}_1) \;\wedge \\ \quad (\forall\, \overline{e}_2 = (\overline{n}_2, f, \overline{n}_2') \in \overline{St}_2 \;.\; \exists\, \overline{e}_1 = (\mathcal{I}(\overline{n}_2), f, \mathcal{I}(\overline{n}_2')) \in \overline{St}_1 \;. \\ \qquad \mathit{get}_{\overline{\mathsf{ID}}}(\overline{\mathit{con}}(\overline{e}_1)) = \mathit{get}_{\overline{\mathsf{ID}}}(\mathit{rename}_{\mathsf{ISO}}(\mathcal{I}, \overline{\mathit{con}}(\overline{e}_2)))) \;\wedge \\ \quad (\forall\, \overline{e}_1 = (\overline{n}_1, f, \overline{n}_1') \in \overline{St}_1 \;.\; \exists\, \overline{e}_2 = (\mathcal{I}^{-1}(\overline{n}_1), f, \mathcal{I}^{-1}(\overline{n}_1')) \in \overline{St}_2 \;. \\ \qquad \mathit{get}_{\overline{\mathsf{ID}}}(\overline{\mathit{con}}(\overline{e}_1)) = \mathit{get}_{\overline{\mathsf{ID}}}(\mathit{rename}_{\mathsf{ISO}}(\mathcal{I}, \overline{\mathit{con}}(\overline{e}_2)))) \end{array}$$

where $\mathit{get}_{\overline{\mathsf{ID}}} : \overline{\mathsf{V}} \to \mathcal{P}(\overline{\mathsf{ID}})$ returns the set of abstract identifiers of a given abstract value state.

If the merged heaps are compatible, we can apply edge-wise widening to the conditions of the heaps. Let $\mathit{widen}_{\overline{\mathsf{Cond}}} : (\overline{\mathsf{Cond}} \times \overline{\mathsf{Cond}}) \to \overline{\mathsf{Cond}}$ be a function that applies widening edge-wise to conditions of compatible heaps. Formally,

$$\mathit{widen}_{\overline{\mathsf{Cond}}}(\overline{\mathit{con}}_1, \overline{\mathit{con}}_2) = [\overline{e} \mapsto \overline{\mathit{con}}_1(\overline{e}) \nabla \overline{\mathit{con}}_2(\overline{e}) \mid \overline{e} \in \mathit{dom}(\overline{\mathit{con}}_1) \cap \mathit{dom}(\overline{\mathit{con}}_2)]$$

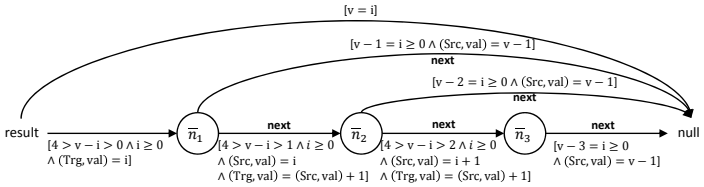


Fig. 7: Heap before the fourth iteration of the fixed-point computation of the loop invariant in Fig. 1.

Now, we are in the position to define the widening operator $\nabla_{\overline{\Sigma}} : (\overline{\Sigma} \times \overline{\Sigma}) \to \overline{\Sigma}$. First, we compute an intermediate heap that is the join of the heaps we want to widen. We do this in order to guarantee that all the edges from both heaps are considered when we subsequent merge of nodes is performed. After merging the nodes in the first heap and the intermediate heap, we check whether these heaps are compatible for widening of their abstract conditions. This means that the heaps are isomorphic and the identifiers of the abstract conditions are stabilized. If they are, we rename identifiers in the conditions of the intermediate merged heap according to the isomorphism $\mathcal{I}$ that is the witness to the compatibility of the merged heaps, and we widen the conditions on the edges. On the other hand, if the merged heaps are not compatible, we return the intermediate heap as it is as upper bound on the heaps we want to widen. Formally,

$$(\overline{\sigma}_1)\nabla_{\overline{\Sigma}}(\overline{\sigma}_2) = \begin{cases} (\overline{St}'_1, \mathit{widen}_{\overline{\mathsf{Cond}}}(\overline{\mathit{con}}'_1, \overline{\mathit{con}}''_2), \mathit{isSum}'_1) & \text{if } \mathit{compatible}(\overline{\sigma}'_1, \overline{\sigma}'_2) \\ \overline{\sigma}'_2 & \text{otherwise} \end{cases}$$

where $\overline{\sigma}'_1 = (\overline{St}'_1, \overline{\mathit{con}}'_1, \mathit{isSum}'_1) = \mathit{merge}_{\overline{\Sigma}}(\overline{\sigma}_1)$, $\overline{\sigma}'_2 = (\overline{St}'_2, \overline{\mathit{con}}'_2, \mathit{isSum}'_2) = \mathit{merge}_{\overline{\Sigma}}(\overline{\sigma}_1 \sqcup \overline{\sigma}_2)$, in case of the merged heaps being compatible, $\mathcal{I}$ is the witness and $\overline{\mathit{con}}''_2 = [(\mathcal{I}(\overline{n}_1), f_r, \mathcal{I}(\overline{n}_2)) \mapsto \mathit{rename}_{\mathsf{ISO}}(\mathcal{I}, \overline{\mathit{con}}'_2(\overline{n}_1, f_r, \overline{n}_2)) \mid (\overline{n}_1, f_r, \overline{n}_2) \in \overline{St}'_2]$.

**Termination.** Since $\mathcal{W}$ is finite and there are only finitely many local variables, there are only finitely many equivalence classes for the merging of nodes. Therefore, the abstract heap resulting from merging has a bounded number of nodes. When the two heaps after merging of nodes are isomorphic, termination of widening follows from the fact that there are only finitely many edges and that widening of the value analysis is assumed to terminate. When they are not isomorphic, the heaps are joined. However, the join can be applied only finitely many times, since there are only finitely many non-isomorphic structures that can be constructed from a fixed number of nodes, and after every application of join, the size of the abstract structure grows, as join is a monotone operator. Furthermore, also the set of identifiers tracked by conditions on edges grows with every application of join to conditions and there are only finitely many identifiers (since there are only finitely many nodes in the heap and $\mathsf{Field}_{\mathsf{Val}}$ is finite).

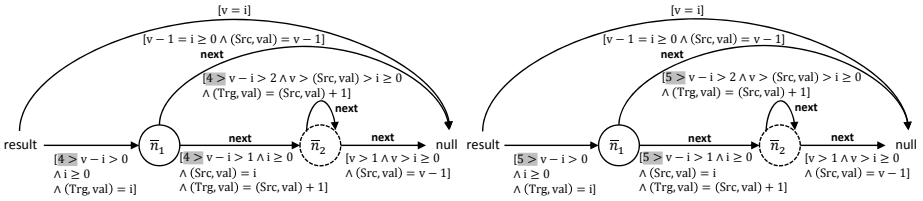


Fig. 8: Heaps with merged nodes before the fourth (left) and fifth (right) iteration of the fixed-point computation. The heaps differ only in the highlighted constraints.

*Example.* Suppose we widen the abstract heap before the fourth iteration of the fixed-point computation for the loop in Fig. 1 with the heap before the fifth iteration. The abstract heap before the fourth iteration is displayed in Fig. 7; the heap before the fifth iteration looks similar, but has four definite nodes.

In the first step, widening merges nodes using the default $\mathcal{W} = \{\mathsf{result}\}$. In the heap from Fig. 7, we merge $\overline{n}_2$ and $\overline{n}_3$ into a single summary node $\overline{n}_2$ since they are (i) denoted by the same set of field access expressions from $\mathcal{W}$ (the empty set since result denotes neither $\overline{n}_2$ nor $\overline{n}_3$), and (ii) they are reachable from the same set of local variables ($\{\mathsf{result}\}$). However, $\overline{n}_1$ is denoted by a different set of field access expressions ($\{\mathsf{result}\}$), and therefore not merged with $\overline{n}_2$ and $\overline{n}_3$. The edges from $\overline{n}_2$ and $\overline{n}_3$ to null are also merged, and their conditions are joined. The resulting heap is shown on the left of Fig. 8. Merging nodes in the heap before the fifth iteration (not shown) results in the heap on the right of Fig. 8. Note that these heaps are isomorphic, that is, the heap shape has stabilized.

In the second step, since the heaps after merging are isomorphic, we apply edge-wise widening to the abstract conditions. This step removes the upper bound on $\mathsf{v} - \mathsf{i}$, but leaves all other constraints unaffected, that is, the abstract conditions have stabilized. The resulting heap is shown in Fig. 4; it represents the loop invariant of the program from Fig. 1.

## 4 Evaluation of Expressions

### 4.1 Evaluation of Reference Expressions

Here, we formalize the evaluation of reference expressions (rexp). A reference expression is either null, an allocation of a new object (new C), or a reference field access expression (rAE). We first formalize the evaluation of new C and the evaluation of rAE. Finally, we define the evaluation function of reference expressions.

We assume that the underlying value analysis $\overline{\mathsf{V}}$ provides the following operators based on well-known definitions [15]: (*i*) *add* : $\mathcal{P}(\overline{\mathsf{ID}}) \times \overline{\mathsf{V}} \to \overline{\mathsf{V}}$, which introduces a given set of identifiers with value $\top$ into a given state, (*ii*) *drop* : $\mathcal{P}(\overline{\mathsf{ID}}) \times \overline{\mathsf{V}} \to \overline{\mathsf{V}}$, which projects out a given set of identifiers from a state, and (*iii*) *expand* : $\mathcal{P}(\overline{\mathsf{ID}} \times \overline{\mathsf{ID}}) \times \overline{\mathsf{V}} \to \overline{\mathsf{V}}$, which extends a given state such that for any $(\overline{id}_1, \overline{id}_2)$ of a given set of pairs of identifiers, the abstract value of $\overline{id}_2$ is the

same as the abstract value of $\overline{id}_1$, and the values of the identifiers present in the original state remain unchanged.

We partition the set of edge-local identifiers into *source* edge-local identifiers $\overline{\mathsf{EId}}_{\mathsf{Src}} = \{\mathsf{Src}\} \times \mathsf{Field}_{\mathsf{Val}}$ and *target* edge-local identifiers $\overline{\mathsf{EId}}_{\mathsf{Trg}} = \{\mathsf{Trg}\} \times \mathsf{Field}_{\mathsf{Val}}$.

In the evaluation of both reference and value expressions, we will have to determine a value state that approximates the value stores of all heaps represented by a given abstract heap. We call such a value state the *general condition* of an abstract heap. It is defined as follows. *Each* abstract condition associated with an edge of an abstract heap approximates the value store of *all* concrete heaps in which the edge exists. Therefore, it is sufficient to consider any abstract reference $\overline{n}$ that abstracts a concrete reference that exists in every concrete heap (such as $\Im$) and any field $f$. The join of the conditions of the edges $(\overline{n}, f, \_)$ is a value state that approximates the value stores of all concrete stores (since $\overline{n}.f$ must have some value in each of them). In the definition below, we take the join over all fields for simplicity, and drop the edge-local identifiers since the general condition is not associated with any particular edge and, hence, the edge-local identifiers have no meaning. Formally, the general condition $general_{\overline{\mathsf{V}}} : \overline{\Sigma} \to \overline{\mathsf{V}}$ of an abstract heap is defined as $general_{\overline{\mathsf{V}}}(\overline{St}, \overline{con}, \_) = \bigsqcup_{\overline{e} \in \overline{St}} \{drop(\overline{\mathsf{EId}}, \overline{con}(\overline{e})) \mid \Im = source(\overline{e})\}$.

**Object Allocation.** An abstract heap $\overline{\sigma}'$ soundly approximates an allocation of a new object in abstract heap $\overline{\sigma}$ if and only if for each concrete heap $h \in \gamma(\overline{\sigma})$ and for each concrete heap $h'$ resulting from the allocation of a new object in $h$, the concretization of $\overline{\sigma}'$ contains $h'$, that is, $h' \in \gamma(\overline{\sigma}')$. Let $alloc_{\overline{\Sigma}} : \overline{\Sigma} \to (\mathcal{P}((\overline{\mathsf{Ref}} \times \overline{\mathsf{V}})) \times \overline{\Sigma})$ be a function that for a given abstract heap returs (i) a singleton set[3] containing the abstract reference representing the newly allocated object and the abstract condition approximating value stores of heaps resulting from the allocation, and (ii) an abstract heap which soundly approximates object allocation in the given heap. Formally, $alloc_{\overline{\Sigma}}(\overline{St}, \overline{con}, isSum) = (\{(\overline{n}, \overline{s})\}, (\overline{St}', \overline{con}', isSum'))$ where

$$\overline{n} \in (\overline{\mathsf{Ref}} \setminus (\{\overline{n}' \mid (\overline{n}', \_, \_) \in \overline{St} \ \vee \ (\_, \_, \overline{n}') \in \overline{St}\} \cup \{\mathsf{null}, \Im\})) \wedge \quad (1)$$

$$isSum' = isSum[\overline{n} \mapsto \mathsf{false}] \wedge \quad (2)$$

$$\overline{I} = \{\overline{n}\} \times \mathsf{Field}_{\mathsf{Val}} \wedge \quad (3)$$

$$\overline{s} = add(\overline{I} \cup \overline{\mathsf{EId}}, general_{\overline{\mathsf{V}}}(\overline{St}, \overline{con}, isSum)) \wedge \quad (4)$$

$$\overline{St}' = \overline{St} \cup \{(\overline{n}, f_r, \mathsf{null}) \mid f_r \in \mathsf{Field}_{\mathsf{Ref}}\} \wedge \quad (5)$$

$$\overline{con}' = \left[ \overline{e} \mapsto \begin{cases} add(\overline{I}, \overline{con}(e)) & \text{if } \overline{e} \in \overline{St} \\ drop(\overline{\mathsf{EId}}_{\mathsf{Trg}}, \overline{s}) & \text{otherwise} \end{cases} \middle| \ \overline{e} \in \overline{St}' \right] \quad (6)$$

The newly allocated object is represented by a fresh non-$\mathsf{null}$ node $\overline{n}$ (1), which is definite (2). The analysis builds up the value state $\overline{s}$ approximating the value store of all concrete heaps in which the new object has been allocated (using the general condition defined above). Consequently, $\overline{s}$ must contain the identifiers $\overline{I}$

[3] Returning a singleton set rather than a single pair simplifies the definition of $\overline{eval}_{\mathsf{rexp}}$ in Sec. 4.1.

for the value fields of $\overline{n}$ (3). Since the allocation of an object can appear only on the right-hand side of an assignment, $\overline{s}$ will participate in the condition of an edge that represents the assignment (see line (6) of the abstract semantics of reference assignment in Sec. 5.1). For this purpose, we also introduce edge-local identifiers into $\overline{s}$ (4). We assume that the concrete semantics of object allocation assigns $\mathsf{null}$ to reference fields of the newly allocated object. Accordingly, the analysis updates the abstract store such that the reference fields of $\overline{n}$ point to $\mathsf{null}$ (5). Finally, the abstract conditions are updated: for each edge that is already present, we add identifiers $\overline{I}$ to their existing conditions. For every new edge, we add the condition $\overline{s}$ in which the target edge-local identifiers are not present (since the target of the edge is $\mathsf{null}$) (6).

**Reference Field Access Expressions.** Since the abstract store includes disjunctive information, a reference field access expression $\Im.f_1.\cdots.f_n \in \mathsf{rAE}$ may correspond to multiple paths in the abstract heap. When evaluating $\Im.f_1.\cdots.f_n$, all these paths must be considered. Let $\mathit{paths} : (\mathsf{rAE} \times \overline{\mathsf{Store}_{\overline{\mathsf{Ref}}}}) \to \mathcal{P}((\overline{\mathsf{Ref}} \times \mathsf{Field}_{\mathsf{Ref}} \times \overline{\mathsf{Ref}})^*)$ be a function that returns the set of all paths in a given abstract store for a given reference field access expression. Formally:

$$\mathit{paths}(\Im.f_1.\cdots.f_n, \overline{St}) = \{\langle \overline{e}_1, \cdots, \overline{e}_n\rangle \mid \forall\ 1 \le i \le n\ \cdot \overline{e}_i \in \overline{St}\ \wedge\ \mathit{field}(\overline{e}_i) = f_i \wedge \mathit{source}(\overline{e}_1) = \Im\ \wedge\ ((1 < i) \Rightarrow (\mathit{source}(\overline{e}_i) = \mathit{target}(\overline{e}_{i-1})))\}$$

Each edge in a path is associated with an abstract condition that approximates the value stores of all concrete heaps in which the edge exists. Therefore, the path is feasible only in those concrete heaps whose value stores are approximated by the abstract conditions associated with *all* edges of the path. However, since the abstract condition of an edge contains edge-local identifiers for the source and the target of this particular edge, it is not generally sound to simply take the greatest lower bound of all abstract conditions of the path. If $\overline{e}_i$ and $\overline{e}_{i+1}$ are two consecutive edges in the path then the target edge-local identifiers in the abstract condition of $\overline{e}_i$ correspond to the source edge-local identifiers in the abstract condition of $\overline{e}_{i+1}$. To adjust the edge-local identifiers, we define a function $\mathsf{Trg}\mathit{To}\mathsf{Src} : \overline{\mathsf{V}} \to \overline{\mathsf{V}}$, which converts the target edge-local identifiers of a given abstract condition to source edge-local identifiers and sets the target edge-local identifiers to $\top$. Formally:

$$\mathsf{Trg}\mathit{To}\mathsf{Src}(\overline{s}) = \mathit{add}(\overline{\mathsf{EId}}_{\mathsf{Trg}}, \mathit{drop}(\overline{\mathsf{EId}}_{\mathsf{Trg}}, \mathit{expand}(R, \mathit{drop}(\overline{\mathsf{EId}}_{\mathsf{Src}}, \overline{s}))))$$

where $R = \{((\mathsf{Trg}, f_v), (\mathsf{Src}, f_v)) \mid f_v \in \mathsf{Field}_{\mathsf{Val}}\}$. In the above definition, we convert the target edge-local identifiers to source local-identifiers by first removing the source edge-local identifiers from the value state (via *drop*), and then re-introducing them with the abstract values of the corresponding target edge-local identifiers (via *expand*). Since the source edge-local identifiers assumed the role of the target edge-local identifiers, the target edge-local identifiers are removed from the state and then re-introduced with the value $\top$ (via *add*).

We now define a function $\mathit{pathCondition} : (\overline{\mathsf{V}} \times (\overline{\mathsf{Ref}} \times \mathsf{Field}_{\mathsf{Ref}} \times \overline{\mathsf{Ref}})^* \times \overline{\mathsf{Cond}}) \to \overline{\mathsf{V}}$, which, for a given abstract condition and a given path, returns the

abstract condition under which the path is feasible and does not contradict the condition. Furthermore, in case the last node of the given path is non-null, the target edge-local identifiers of the resulting abstract condition refer to the value fields of the last node of the path. Formally:

$$\begin{array}{l} \mathit{pathCondition}(\overline{s}, \langle \overline{e}_1, \cdots, \overline{e}_n \rangle, \overline{\mathit{con}}) = \\ \quad \begin{cases} \mathit{add}(\overline{\mathsf{EId}}_{\mathsf{Src}}, \mathit{drop}(\overline{\mathsf{EId}}_{\mathsf{Src}}, \overline{s})) & \text{if } \langle \overline{e}_1, \cdots \overline{e}_n \rangle = \varepsilon \\ \mathit{pathCondition}(\mathsf{Trg}\,\mathit{To}\mathsf{Src}(\overline{s}) \sqcap \overline{\mathit{con}}(\overline{e}_1), \langle \overline{e}_2, \cdots \overline{e}_n \rangle, \overline{\mathit{con}}) & \text{otherwise} \end{cases} \end{array}$$

where $\varepsilon$ represents an empty path. In the above definition, $\overline{s}$ represents the value information accumulated from the abstract conditions of already processed edges. In the last line of the definition, we obtain the next value of this accumulator by taking the greatest lower bound of the current $\overline{s}$ (with appropriate adjustments of edge local identifiers) and the abstract condition of the next unprocessed edge.

Now we define an evaluation function for reference field access expressions $\overline{\mathit{eval}}_{\mathsf{rAE}} : (\mathsf{rAE} \times \overline{\Sigma}) \to \mathcal{P}(\overline{\mathsf{Ref}} \times \overline{\mathsf{V}})$, which returns a set of pairs of abstract references and abstract value states. In each pair, the reference is the last node of a paths that may represent the given reference field access expressions, and the value state describes under which condition this path is feasible. In each pair, the target edge-local identifiers of the abstract condition refer to value fields of the abstract reference. Note that if the abstract condition is $\bot$, the path is not feasible and such a path may be discarded. Formally:

$$\begin{array}{l} \overline{\mathit{eval}}_{\mathsf{rAE}}(\mathit{ap}, (\overline{St}, \overline{\mathit{con}}, \mathit{isSum})) = \{(\overline{n}, \overline{s}) \mid \exists\ \langle \overline{e}_1, \cdots \overline{e}_n \rangle \in \mathit{paths}(\mathit{ap}, \overline{St}) \cdot \\ \quad \mathit{target}(\overline{e}_n) = \overline{n}\ \wedge\ \overline{s} = \mathit{pathCondition}(\top, \langle \overline{e}_1, \cdots \overline{e}_n \rangle, \overline{\mathit{con}})\ \wedge\ \overline{s} \neq \bot\} \end{array}$$

**Reference Expressions.** We define an evaluation function $\overline{\mathit{eval}}_{\mathsf{rexp}} : ((\mathsf{rexp} \cup \{\Im\}) \times \overline{\Sigma}) \to (\mathcal{P}(\overline{\mathsf{Ref}} \times \overline{\mathsf{V}}) \times \overline{\Sigma})$ for reference expressions and also $\Im$, which simplifies the treatment of assignments to expressions of the form $\Im.f$. $\overline{\mathit{eval}}_{\mathsf{rexp}}$ returns (i) the abstract references that a given reference expression (or $\Im$) may represent together with the abstract conditions under which it may represent these abstract references. In these conditions, the target edge-local identifiers must correspond to value fields of the resulting abstract references. And (ii) the abstract heap that results from evaluating the given expression, since allocation has side-effects. Formally:

$$\overline{\mathit{eval}}_{\mathsf{rexp}}(\mathit{re}, \overline{\sigma}) = \begin{cases} \mathit{alloc}_{\overline{\Sigma}}(\overline{\sigma}) & \text{if } \mathit{re} = \mathsf{new\ C} \\ (\{(\mathsf{null}, \mathit{add}(\overline{\mathsf{EId}}_{\mathsf{Src}}, \mathit{general}_{\overline{\mathsf{V}}}(\overline{\sigma})))\}, \overline{\sigma}) & \text{if } \mathit{re} = \mathsf{null} \\ (\overline{\mathit{eval}}_{\mathsf{rAE}}(\mathit{re}, \overline{\sigma}), \overline{\sigma}) & \text{if } \mathit{re} \in \mathsf{rAE} \\ (\{(\Im, \mathit{add}(\overline{\mathsf{EId}}_{\mathsf{Src}}, \mathit{expand}(R, \mathit{general}_{\overline{\mathsf{V}}}(\overline{\sigma}))))\}, \overline{\sigma}) & \text{if } \mathit{re} = \Im \end{cases}$$

where $R = \{((\Im, f_v), (\mathsf{Trg}, f_v)) \mid f_v \in \mathsf{Field}_{\mathsf{Val}}\}$. Since the abstract conditions returned by $\overline{\mathit{eval}}_{\mathsf{rexp}}$ may participate in conditions of edges (see line (6) of the abstract semantics of reference assignment in Sec. 5.1), we need to make sure that they include the edge-local identifiers.

### 4.2 Evaluation of Value Expressions

Here, we formalize the evaluation of value expressions ($\mathsf{vexp}$). A value expression is either a literal ($\mathsf{n}$), a value field access expression ($\mathsf{vAE}$), or a binary expression ($\mathsf{vexp}\ \langle\mathsf{op}\rangle\ \mathsf{vexp}$). First, we formalize the evaluation of $\mathsf{vAE}$. Then we define the evaluation function for value expressions.

**Value Field Access Expressions.** For a value field access expression $p.f_\mathsf{v}$, the receiver expression $p$ is either $\Im$ or a reference field access expression. In the latter case, it may correspond to multiple paths in the abstract heap. Therefore, there may be multiple ways to evaluate $p.f_\mathsf{v}$, and the evaluation function must consider each of them.

We use function $\overline{eval}_{\mathsf{rexp}}$ (Sec. 4.1) to determine the set of abstract references representing $p$, and for each of them an abstract condition under which $p$ represents this reference. The edge-local identifier $(\mathsf{Trg}, f_\mathsf{v})$ in these conditions refers to $p.f_\mathsf{v}$, and, therefore, is important for the evaluation. However, in Sec. 5.2, we will detach these conditions from their edges such that edge-local identifiers lose their meaning. To preserve the information about $p.f_\mathsf{v}$, we retain it using a separate, non-edge-local identifier before removing the edge-local identifiers. For this purpose, we extend the set of abstract identifiers $\overline{\mathsf{ID}}$ of the value analysis with *value field access identifiers*, that is, $\overline{\mathsf{ID}}_{\mathsf{vAE}} = \overline{\mathsf{ID}} \cup \mathsf{vAE}$. These identifiers are used only temporarily for intermediate steps, but do not occur in any value states before or after a value assignment is performed.

Function $\overline{eval}_{\mathsf{vAE}} : (\mathsf{vAE} \times \overline{\Sigma}) \to \mathcal{P}(\overline{\mathsf{V}})$ yields the set of value states corresponding to different ways a given value field access expression may be evaluated. Formally:

$$\overline{eval}_{\mathsf{vAE}}(p.f_\mathsf{v}, \overline{\sigma}) = \left\{ drop(\overline{\mathsf{Eld}}, \overline{s}) \;\middle|\; \begin{array}{l} \exists\ (\overline{n}, \overline{s}') \in \overline{eval}_{\mathsf{rexp}}(p, \overline{\sigma}) \cdot \overline{n} \neq \mathsf{null} \wedge \\ \overline{s} = expand(\{((\mathsf{Trg}, f_\mathsf{v}), p.f_\mathsf{v})\}, \overline{s}') \end{array} \right\}$$

For a given value field access expression $p.f_\mathsf{v}$, the function first computes the abstract references representing $p$, together with the value states under which $p$ represents them. Abstract null references are excluded since they would lead to null-pointer dereferencing when evaluating $p.f_\mathsf{v}$ and, thus, abortion of the execution. Then, for each abstract condition, the function renames the identifier $(\mathsf{Trg}, f_\mathsf{v})$ to the new value field access identifier $p.f_\mathsf{v}$ (via *expand*). Finally, since the resulting value state describes conditions on the entire heap, the function removes all edge-local identifiers.

**Value Expressions.** We define an evaluation function for value expressions $\overline{eval}_{\mathsf{vexp}} : (\mathsf{vexp} \times \overline{\Sigma}) \to \mathcal{P}(\overline{\mathsf{V}})$, which returns the set of value states corresponding to every possible way a given value expression may be evaluated in a given abstract heap. Each resulting value state contains value field access identifiers for

all value field access expressions presents in the given value expression. Formally:

$$\overline{eval}_{\mathsf{vexp}}(ve, \overline{\sigma}) = \begin{cases} \{general_{\overline{\mathsf{V}}}(\overline{\sigma})\} & \text{if } ve \in \mathsf{n} \\ \overline{eval}_{\mathsf{vAE}}(ve, \overline{\sigma}) & \text{if } ve \in \mathsf{vAE} \\ \left\{ meet_{\mathsf{vAE}}(\overline{s}_1, \overline{s}_2) \,\middle|\, \begin{array}{l} \overline{s}_1 \in \overline{eval}_{\mathsf{vexp}}(ve_1, \overline{\sigma}) \wedge \\ \overline{s}_2 \in \overline{eval}_{\mathsf{vexp}}(ve_2, \overline{\sigma}) \end{array} \right\} & \text{if } ve_1 \langle \mathsf{op} \rangle ve_2 = ve \end{cases}$$

For literals $\mathsf{n}$, the function yields the general condition since a literal can be evaluated in any heap. For value field access expressions, the function applies $\overline{eval}_{\mathsf{vAE}}$ from the previous section. The evaluation of binary expressions works as follows.

Recursively evaluating the expressions $ve_1$ and $ve_2$ yields sets of abstract states that represent all ways of evaluating these expressions. The resulting state for the binary expression is essentially the point-wise application of the greatest lower bound to each combination of states from these sets. However, these states may in general contain different value field access identifiers $\overline{I}_1$ and $\overline{I}_2$. To avoid information loss when applying the greatest lower bound, we first enlarge all involved value states to include all identifiers in $\overline{I}_1 \cup \overline{I}_2$. Let $get_{\mathsf{vAE}} : \overline{\mathsf{V}} \to \mathcal{P}(\overline{\mathsf{ID}}_{\mathsf{vAE}})$ yield a set of value field access identifiers of a given value state. Function $meet_{\mathsf{vAE}} : \overline{\mathsf{V}} \times \overline{\mathsf{V}} \to \overline{\mathsf{V}}$ yields the value state that is the greatest lower bound of the given value states and contains all value field access identifiers from these states:

$$meet_{\mathsf{vAE}}(\overline{s}_1, \overline{s}_2) = add(get_{\mathsf{vAE}}(\overline{s}_2) \setminus get_{\mathsf{vAE}}(\overline{s}_1), \overline{s}_1) \sqcap add(get_{\mathsf{vAE}}(\overline{s}_1) \setminus get_{\mathsf{vAE}}(\overline{s}_2), \overline{s}_2)$$

## 5 Abstract Semantics

In this section, we formalize the semantics of reference and value assignments.

### 5.1 Reference Assignment

An abstract store includes disjunctive information. Therefore, for a reference assignment $\mathsf{p.f_r} = \mathsf{rhs}$, there may be several abstract references for the receiver $\mathsf{p}$ and the right-hand side $\mathsf{rhs}$, which may be reached through different paths with different value conditions. The value states on the edges along each path specify the conditions under which $\mathsf{p}$ and $\mathsf{rhs}$ evaluate to a particular abstract reference. The abstract semantics for reference assignments adds an abstract edge for each possible combination of receiver $\mathsf{p}$ and right-hand side $\mathsf{rhs}$, with an abstract condition that reflects when this combination exists.

The rule below formalizes reference assignments of the form $\mathsf{p.f_r} = \mathsf{rhs}$, where $\mathsf{p.f_r} \in \mathsf{rAE}$ and $\mathsf{rhs} \in \mathsf{rexp}$. Since we encode local variables as fields of a special reference $\Im$, the rule also covers assignments to those. It uses an auxiliary function $\overline{eval}_{\mathsf{rexp}}$, which takes a reference expression (or $\Im$) $re$ and an abstract

state $\overline{\sigma}$ and yields (a) a set $NC$ of pairs, each consisting of an abstract reference to which $re$ may evaluate in $\overline{\sigma}$ and the condition under which $re$ may evaluate to this reference, and (b) a resulting abstract state, which is used to encode allocation, that is, when $re$ contains $\mathsf{new}$ expressions (see Sec. 4.1).

$$(NC_{\mathsf{rhs}}, (\overline{St}_{\mathsf{rhs}}, \overline{con}_{\mathsf{rhs}}, isSum_{\mathsf{rhs}})) = \overline{eval}_{\mathsf{rexp}}(\mathsf{rhs}, \overline{\sigma}) \tag{1}$$

$$(NC_{\mathsf{p}}, \_) = \overline{eval}_{\mathsf{rexp}}(\mathsf{p}, (\overline{St}_{\mathsf{rhs}}, \overline{con}_{\mathsf{rhs}}, isSum_{\mathsf{rhs}})) \tag{2}$$

$$strong \iff \exists\, \overline{n} \in \overline{\mathsf{Ref}} \cdot (NC_{\mathsf{p}} = \{(\overline{n}, \_)\} \ \wedge\ \neg isSum_r(\overline{n})) \tag{3}$$

$$strong \Rightarrow (\overline{St} = \{(\overline{n}_1, f, \overline{n}_2) \in \overline{St}_{\mathsf{rhs}} \mid (\overline{n}_1, \_) \notin NC_{\mathsf{p}} \ \vee\ \mathsf{f_r} \neq f\}) \tag{4}$$

$$(\neg strong) \Rightarrow (\overline{St} = \overline{St}_{\mathsf{rhs}}) \tag{5}$$

$$\overline{con}_{asg} = \left[ (\overline{n}_{\mathsf{p}}, \mathsf{f_r}, \overline{n}_{\mathsf{rhs}}) \mapsto (\mathsf{Trg}\, To \mathsf{Src}(\overline{s}_{\mathsf{p}}) \sqcap \overline{s}_{\mathsf{rhs}}) \;\middle|\; \begin{array}{l} (\overline{n}_{\mathsf{p}}, \overline{s}_{\mathsf{p}}) \in NC_{\mathsf{p}} \wedge \\ (\overline{n}_{\mathsf{rhs}}, \overline{s}_{\mathsf{rhs}}) \in NC_{\mathsf{rhs}} \end{array} \right] \tag{6}$$

$$\overline{St}' = \overline{St} \cup dom(\overline{con}_{asg}) \tag{7}$$

$$\overline{con}' = \overline{con}_{\mathsf{rhs}} \left[ \overline{e} \mapsto \overline{s} \;\middle|\; \begin{array}{l} \overline{e} \in dom(\overline{con}_{asg}) \ \wedge\ (\overline{e} \notin \overline{St} \Rightarrow \overline{s} = \overline{con}_{asg}(\overline{e})) \wedge \\ (\overline{e} \in \overline{St} \Rightarrow \overline{s} = \overline{con}_{asg}(\overline{e}) \sqcup \overline{con}_{\mathsf{rhs}}(\overline{e})) \end{array} \right] \tag{8}$$

$$\langle \mathsf{p.f_r} = \mathsf{rhs}, \overline{\sigma} \rangle \rightarrow_{\overline{\Sigma}} (\overline{St}', \overline{con}', isSum_{\mathsf{rhs}})$$

A reference assignment first evaluates $\mathsf{rhs}$ to obtain the possible abstract references for the right-hand side expression together with the corresponding conditions, as well as a successor state (1). The receiver $\mathsf{p}$ is evaluated in this successor state. Since it is side-effect free (see Fig. 3), we discard the state resulting from its evaluation (2). The analysis performs a strong update iff there is only one abstract reference $\overline{n}$ for the receiver, which is a definite node (3). In that case, the analysis removes all edges whose source is the receiver node and that are labeled with the assigned field $\mathsf{f_r}$ (4); otherwise, it performs a weak update, that is, retains all existing edges (5). To add the edges for all possible combinations of receivers and right-hand sides, we first create a map $\overline{con}_{asg}$ that maps each of the new edges to the abstract condition that describes when the particular combination exists, that is, the greatest lower bound of the conditions for choosing a particular abstract reference for the receiver and a particular abstract reference for the right-hand side, respectively (6). The only twist in this step is how to handle edge-local identifiers. The receiver is denoted by $\mathsf{Trg}$ in conditions on edges pointing to the receiver, but by $\mathsf{Src}$ in the new edges. Function $\mathsf{Trg}\, To \mathsf{Src}$ performs this conversion. Since the map $\overline{con}_{asg}$ contains an entry for each new edge, we obtain the final abstract store by adding the domain of this map to the store constructed in step 4 or 5 (7). Finally, the abstract conditions are updated: For each new edge that is not present in the store before the reference assignment, we add the condition from $\overline{con}_{asg}$. For each edge that is already present (which may happen during a weak update), we join the condition from $\overline{con}_{asg}$ and the existing condition (8).

*Example.* Fig. 9 *without* the bold edges and the highlighted constraints, shows the abstract heap after line 5 in Fig. 1. It is obtained from the abstract heap in Fig. 4 (the loop invariant) by (i) assuming the loop guard ($\mathsf{i} > \mathsf{0}$) in all abstract conditions and (ii) applying the abstract semantics of the statement $\mathsf{Node\ p = new\ Node()}$ (line 5), which introduces the definite node $\overline{n}_3$. (Its $\mathsf{next}$

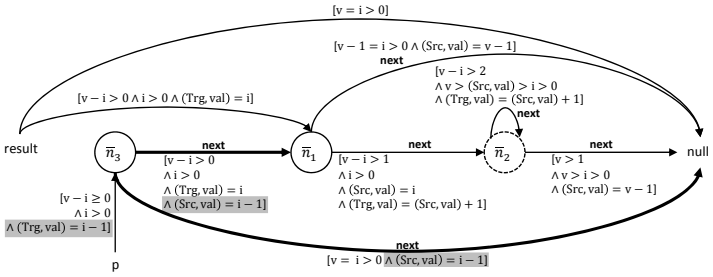


Fig. 9: The abstract heap after line 7 of the program in Fig. 1. The bold edges are added by the reference assignment in line 6; the highlighted constraints come from the value assignment in line 7.

field initially points to null, which is not shown in the figure.) We will now illustrate the abstract semantics of the reference assignment $\mathsf{p.next} = \mathsf{result}$ (line 6).

The right-hand side of the assignment, result, evaluates to null or to $\overline{n}_1$ (point (1) of of the rule above). Since the receiver p evaluates to a single definite node $\overline{n}_3$ (2), we perform a strong update (3). The strong update removes all out-edges of $\overline{n}_3$ labeled with next (4) and introduces new edges for all combinations of abstract references for the receiver and for the right-hand side, that is, edges from $\overline{n}_3$ to null and from $\overline{n}_3$ to $\overline{n}_1$ (7). These edges are shown in bold in Fig. 9. The former edge exists if the right-hand side (result) evaluates to null, that is, if $\mathsf{v} = \mathsf{i} > 0$; the latter exists if result evaluates to $\overline{n}_1$, that is, if $\mathsf{v} - \mathsf{i} > 0 \wedge \mathsf{i} > 0 \wedge (\mathsf{Trg}, \mathsf{val}) = \mathsf{i}$ (6). These constraints are the abstract conditions of the new edges (8), as shown in Fig. 9 (still ignoring the highlighted constraints, which will be discussed later).

### 5.2 Value Assignment

Like the semantics of reference assignments, the abstract semantics of a value assignment $\mathsf{p.f_v} = \mathsf{rhs}$ (where $\mathsf{p.f_v} \in \mathsf{vAE}$ and $\mathsf{rhs} \in \mathsf{vexp}$) needs to consider each possible combination of evaluations for the receiver p and for the right-hand side rhs. For each combination, it updates all abstract conditions in the abstract store to reflect the assignment and the conditions under which the combination exists. The following rule formalizes this intuition.

$$\frac{\begin{array}{rl}\overline{S} = \overline{eval}_{\mathsf{vexp}}(\mathsf{rhs}, (\overline{St}, \overline{con}, isSum)) & (1)\\ (NC_{\mathsf{p}}, \_) = \overline{eval}_{\mathsf{rexp}}(\mathsf{p}, (\overline{St}, \overline{con}, isSum)) & (2)\\ \overline{con}' = update_{\overline{\mathsf{Cond}}}(\mathsf{rhs}, \mathsf{f_v}, NC_{\mathsf{p}}, \overline{S}, \overline{con}) & (3)\end{array}}{\langle \mathsf{p.f_v} = \mathsf{rhs}, (\overline{St}, \overline{con}, isSum)\rangle \rightarrow_{\overline{\Sigma}} (\overline{St}, \overline{con}', isSum)}$$

Each way of evaluating the right-hand side expression rhs, if it contains a field access, chooses a path through the abstract store. The abstract conditions of the edges along a path describe when this path may be chosen. Function $\overline{eval}_{\mathsf{vexp}}$ (defined in Sec. 4.2) yields the set $\overline{S}$ of conditions (value states) that describe

each way of evaluating rhs (1). Analogously to step 2 of reference assignment, we evaluate the receiver expression p to obtain the possible receiver references, each with a condition under which p may evaluate to this reference (2). We use the function $update_{\overline{\mathsf{Cond}}}$ (see below) to reflect the value assignment in the value states of all edges in the abstract store (3). This function considers all possible combinations of receiver reference (obtained from $NC_{\mathsf{p}}$) and value state for a particular way of evaluating the right-hand side expression (from $\overline{S}$). For each of them, it propagates the value information that has to hold when this combination is chosen to the conditions of each edge in the abstract store and applies the assignment operation of the value domain. The condition of each edge in the abstract store is then defined to be the join of the conditions obtained for all ways of executing the value assignment.

*Example.* Fig. 9 *without* the highlighted constraints shows the abstract heap after line 6 in Fig. 1. The highlighted constraints are introduced by the abstract semantics of the statement p.val = i - 1 (line 7). There is only one way to evaluate the right-hand side expression. Therefore, $\overline{eval}_{\mathsf{vexp}}$ yields a singleton set (point (1) of the rule above). This set contains the condition $\mathsf{v} - \mathsf{i} \geq 0 \wedge \mathsf{i} > 0$, which holds in each concrete heap (otherwise there would be no value for result). Similarly, the receiver expression p evaluates to a single node, $\overline{n}_3$, under the same condition (2). This condition must be satisfied in order to be able to perform the assignment. Therefore, we conjoin it to each abstract condition in the store (which has no effect in this example), and then assign $\mathsf{i} - 1$ to $(\overline{n}_3, \mathsf{val})$ since p evaluates to $\overline{n}_3$ (3). Moreover, since $\overline{n}_3$ is the target of the edge from p to $\overline{n}_3$, we also add the constraint $(\mathsf{Trg}, \mathsf{val}) = \mathsf{i} - 1$ for the edge-local identifier to the condition on this edge, and analogously for $(\mathsf{Src}, \mathsf{val})$ on both out-edges of $\overline{n}_3$ (3).

**Updating Conditions.** Here, we formalize an update function used in above in the definition of the abstract semantics of value assignment.

We assume that the value analysis has the notion of summary identifiers and provides an operator $assign : \mathcal{P}(\overline{\mathsf{ID}}) \times \mathsf{vexp} \times \overline{\mathsf{V}} \to \overline{\mathsf{V}}$ that performs a strong assignment of a given expression to a given set of identifiers. We define a weak assignment $assign_w : \mathcal{P}(\overline{\mathsf{ID}}) \times \mathsf{vexp} \times \overline{\mathsf{V}} \to \overline{\mathsf{V}}$ as $assign_w\left(\overline{I}, ve, \overline{s}\right) = assign(\overline{I}, ve, \overline{s}) \sqcup \overline{s}$.

We define an update function $update_{\overline{\mathsf{Cond}}} : (\mathsf{vexp} \times \mathsf{Field}_{\mathsf{Val}} \times \mathcal{P}(\overline{\mathsf{Ref}} \times \overline{\mathsf{V}}) \times \mathcal{P}(\overline{\mathsf{V}}) \times \overline{\mathsf{Cond}}) \to \overline{\mathsf{Cond}}$, which updates abstract conditions of all edges to reflect an assignment to a value field of a receiver with respect to the value states that correspond to the evaluations of the right hand side of the assignment. Formally:

$$update_{\overline{\mathsf{Cond}}}(ve, f_{\mathsf{v}}, NC, \overline{S}, \overline{con}) = \left[ \overline{e} \mapsto \bigsqcup \left\{ \overline{s}_3 \;\middle|\; \begin{array}{l} \exists\ \overline{s} \in \overline{S}, (\overline{n}, \overline{s}') \in NC \cdot \\ \exists\ \overline{s}_1 = propagate_{\overline{\mathsf{V}}}(\overline{s}, \overline{s}', \overline{con}(\overline{e})) \cdot \\ \exists\ \overline{s}_2 = asg(affected_{\overline{\mathsf{ID}}}(\overline{e}, f_{\mathsf{v}}, \overline{n}), ve, \overline{s}_1)\ \cdot \\ \quad \overline{s}_3 = drop(\mathsf{vAE}, \overline{s}_2) \end{array} \right\} \;\middle|\; \overline{e} \in dom(\overline{con}) \right]$$

where $asg = \begin{cases} assign & \text{if } \neg isSum(\overline{n}) \\ assign_w & \text{otherwise} \end{cases}$.

For every edge $\overline{e}$, the function computes a new abstract condition as a join of conditions $\overline{s}_3$ computed as follows:

The set $\overline{S}$ of value states corresponds to all possible evaluations of the value expression $ve$. $NC$ is a set of abstract references representing a receiver together with an abstract condition under which this reference represents the receiver. A possible combination of $\overline{s}$ in $\overline{S}$ and $(\overline{n}, \overline{s}')$ in $NC$ is chosen. The value information from $\overline{s}$ and $\overline{s}'$ is propagated (via $propagate_{\overline{\mathsf{V}}}$, defined below) into the abstract condition of $\overline{e}$, obtaining the value state $\overline{s}_1$. $\overline{s}_1$ contains all value field access identifiers from $ve$. Since several different identifiers in $\overline{\mathsf{ID}}$ may represent the same memory location, all the identifiers that may represent the value field $f_\mathsf{v}$ of the receiver $\overline{n}$ must be updated. We define a function $affected_{\overline{\mathsf{ID}}} : ((\overline{\mathsf{Ref}} \times \mathsf{Field}_{\mathsf{Ref}} \times \overline{\mathsf{Ref}}) \times \mathsf{Field}_{\mathsf{Val}} \times \overline{\mathsf{Ref}}) \to \mathcal{P}(\overline{\mathsf{ID}})$ which yields the set of identifiers for a given edge that may represent the same memory location as the identifier $(\overline{n}, f_\mathsf{v})$. Formally:

$$affected_{\overline{\mathsf{ID}}}((\overline{n}_1, f_r, \overline{n}_2), f_\mathsf{v}, \overline{n}) = \{(\mathsf{Src}, f_\mathsf{v}) \mid \overline{n}_1 = \overline{n}\} \cup \{(\mathsf{Trg}, f_\mathsf{v}) \mid \overline{n}_2 = \overline{n}\} \cup \{(\overline{n}, f_\mathsf{v})\}$$

Then, we apply a strong (if $\overline{n}$ is definite) or a weak (if $\overline{n}$ is summary) assignment of $ve$ to all identifiers that may represent the same memory location as $(\overline{n}, f_\mathsf{v})$, obtaining the value state $\overline{s}_2$. However, $\overline{s}_2$ still contains value field access identifiers of $ve$ which cannot appear in abstract conditions of edges. We obtain $\overline{s}_3$ by removing all value field access identifiers from $\overline{s}_2$ (via $drop$).

Function $propagate_{\overline{\mathsf{V}}} : \overline{\mathsf{V}} \times \overline{\mathsf{V}} \times \overline{\mathsf{V}} \to \overline{\mathsf{V}}$ propagates the value information from a given value state $\overline{s}$ corresponding to an evaluation of a value expression and a given abstract condition $\overline{s}'$ for a receiver to a given abstract condition of an edge $\overline{s}_{\overline{e}}$. Formally:

$$propagate_{\overline{\mathsf{V}}}(\overline{s}, \overline{s}', \overline{s}_{\overline{e}}) = meet_{\mathsf{vAE}}(add(\overline{\mathsf{EId}}, meet_{\mathsf{vAE}}(\overline{s}, \overline{s}')), \overline{s}_{\overline{e}})$$

The above function first gathers the value information from $\overline{s}$ and $\overline{s}'$ into an intermediate state by taking the greatest lower bound of $\overline{s}$ and $\overline{s}'$ and preserving the value field access identifiers of the both states. Then, in order to preserve the edge local identifiers of $\overline{s}_{\overline{e}}$, the edge-local identifiers are introduced into the intermediate state. Finally, we propagate the value information from the intermediate state to $\overline{s}_{\overline{e}}$ by taking the greatest lower bound of the intermediate state and $\overline{s}_{\overline{e}}$ while preserving the value field access identifiers of the intermediate state.

## 6 Experimental Results

We implemented our analysis in the static analyzer Sample and applied it to Scala implementations of typical list and tree operations (some of which we took from the literature [4,6,13]), operations on nested recursive data structures (such as lists of lists), and a simple aggregate structure [11]. We performed the experiments on an Intel Core i7-Q820 CPU (1.73GHz, 8GB) running the 64-bit version of Ubuntu 14.04. We instantiated our analysis with the Octagon [22]

| Data Structure | Operations | Octa. | Poly. |
|---|---|---|---|
| SortedSLL | constructor<br>insertKey<br>deleteKey<br>findKey<br>deepCopy | 1.24 | 1.82 |
| SortedDLL | constructor<br>insertKey<br>deleteKey<br>findKey<br>deepCopy | 1.91 | 2.83 |
| BST | constructor<br>insertKey<br>findKey | 1.96 | 2.43 |
| NodeCachingSLL | constructor<br>add<br>remove<br>findKey | 0.87 | 1.04 |
| PersonAndAccount | withdraw<br>deposit<br>changeInterest | 0.38 | 0.43 |

Table 1: Analysis times (in seconds) of classes implementing different data structures when instantiating the analysis with the Octagon and Polyhedra value domains. For each class, we inferred an object invariant by computing a fixed point over all its methods.

and Polyhedra [8] value domains implemented in Apron [16]. We used the default widening parameter, that is, $\mathcal{W}$ is the set of local reference variables. There were no manual annotations for any of the benchmarks.

**Inference of Object Invariants.** Tab. 1 reports the analysis times (the average of 10 runs) for implementations of five different data structures. We instantiated Logozzo's framework [18] with our analysis to infer object invariants for each data structure by computing a fixed point over all its operations.

SortedSLL implements a sorted singly-linked list (SLL). The inferred object invariant expresses that the values stored in the list nodes are non-decreasing.

SortedDLL implements a sorted doubly-linked list (DLL). Our analysis infers sortedness in both directions, that is, via the next and prev fields. However, the analysis cannot infer the structural invariant of doubly-linked lists $n.\mathsf{next} \neq \mathsf{null} \Rightarrow n = n.\mathsf{next}.\mathsf{prev}$ because it has no way of relating the concrete references of the two edges $\overline{n} \xrightarrow{\mathsf{next}} \overline{n}$ and $\overline{n} \xrightarrow{\mathsf{prev}} \overline{n}$ for the summary node $\overline{n}$.

BST implements a binary search tree. The analysis infers both the value and the shape information of a BST data structure. Our implementation stores the infimum and supremum of all keys of a sub-tree in its root. This information allows our analysis to distinguish the left and right sub-tree of a node and, thus, to infer that the shape is not a DAG. We omitted method deleteKey because our analysis is not able to infer that replacing the deleted key with the next smallest key preserves sortedness; it does, however, infer that the tree shape is preserved.

NodeCachingSLL implements an acyclic SLL that maintains a cache of node objects to reduce object creation and garbage collection. The inferred object invariant expresses that the list and the cache are disjoint and that the size of the cache is between 0 and maximumCacheSize. Moreover, we inferred that the addKey method creates a new object only if the cache is empty. Every node of the list stores the length of the list rooted at the node. This information lets our analysis infer that the list and its cache are acyclic, which is needed to infer disjointness of the list and the cache. The latter step required materialization, that is, splitting a definite node off a summary node, which is supported by our implementation, but not explained in this paper.

| Operation | Octa. | Poly. |
|---|---|---|
| insertionSort | 0.43 - SLL<br>0.72 - DLL | 0.48 - SLL<br>0.85 - DLL |
| partitionWithKey | 0.32 - SLL<br>0.48 - DLL | 0.34 - SLL<br>0.55 - DLL |
| createListOfZerosAndSum | 0.22 - SLL<br>0.39 - DLL | 0.23 - SLL<br>0.43 - DLL |
| increasingList | 0.28 - SLL<br>0.41 - DLL | 0.31 - SLL<br>0.50 - DLL |
| sortListOfListsOfValues | 1.45 | 1.88 |
| listToBST | 1.03 | 1.21 |

Table 2: Analysis times (in seconds) for single operations on different data structures when instantiating the analysis with the Octagon and Polyhedra value domains. The first 4 operations work on singly-linked lists (SLL) and doubly-linked lists (DLL). The fifth operation works on lists of singly-linked lists that store values. The last operation transforms an SLL to a binary search tree.

Besides the object invariants for these four classes, our analysis infers that the result of method findKey is either null or has the value of the given key. This postcondition is inferred even if the result is represented by a summary node.

PersonAndAccount implements an aggregate data structure similar to the example from a paper on the verification of object invariants [11]. The analysis infers combined shape and value invariants, for instance, that Account and Person objects reference each other, the sum of the account balance and person's salary is positive, and the interest rate of the account is always non-negative.

**Inference of Method Postconditions.** Tab. 2 reports the analysis times of individual operations on different data structures. The initial abstract states and the abstract heaps that represent the arguments to the operations contain only information that is provided by the static types; no annotations were used. The first four operations manipulate singly and doubly-linked lists. insertionSort takes an unsorted list of values and sorts it. The analysis infers that the result is a sorted list. partitionWithKey takes a list of values and a key, and creates two new lists such that the keys in one are less than or equal to the given key, and the keys in the other are greater. The analysis infers this value property and that the resulting lists are disjoint. createListOfZerosAndSum creates a list of zeros and subsequently traverses the list and sums up the values. The analysis infers that the result is a list of zeros, and the sum of the values is zero. increasingList is the method from Fig. 1, with an analogous implementation for DLLs. The analysis infers the heap in Fig. 2 (and an analogous heap for DLLs).

The last two operations of Tab. 2 demonstrate that the analysis is able to infer non-trivial shape and value properties for programs manipulating nested recursive data structures or a combination of different data structures. sortListOfListsOfValues takes a singly-linked list of SLLs that store values, and sorts each of the lists. The analysis infers that the result is a list of sorted SLLs. listToBST takes an SLL of values and creates a binary search tree out of it, without using the methods of the BST class discussed above. The analysis infers that the result is a binary search tree.

**Discussion.** The analysis times in Tab. 1 and Tab. 2 demonstrate the efficiency of our analysis. For all our benchmark classes, the fixed point over all their methods was computed within 3 seconds when using the Polyhedra domain. When instantiated with a more efficient but less precise value domain, the efficiency of the analysis increases, as illustrated by the usage of the Octagon domain.

Our experiments demonstrate that our analysis can infer invariants that combine shape and value information in interesting ways, for instance, sortedness of lists and trees, or invariants that relate the states of different objects in an aggregate structure. Our analysis leverages data stored in value fields, such as the infimum and supremum in the BST class discussed above, to obtain more precise shape information. As future work, we plan to rely less on such fields by tracking additional abstract conditions (such as injectivity of references) on edges and by generalizing edge-local identifiers to reference fields.

## 7 Related Work

Dillig et al. [9,10] present a precise content analysis for arrays and containers, in which heap edges are qualified by logical constraints over indexes into a container. This idea inspired our approach of tracking disjunctive information via the value states associated with edges in the heap. Our analysis uses generic value domains instead of logical constraints and can therefore be instantiated with different levels of precision and efficiency. Moreover, it uses edge-local identifiers instead of indexes, which allows us to express constraints on arbitrary nodes (especially summary nodes) in the heap, not only on indexed structures such as arrays and containers. Whereas Dillig et al. concentrate on clients of arrays and containers, our analysis targets arbitrary heap-manipulating programs including implementations of containers.

Similarly to our work, Bouajjani et al. [1,2] introduce a static analysis that automatically infers combined shape and numerical invariants and is parametric in the underlying value domain. The main difference is that their technique is specific to programs that manipulate singly-linked lists of values. For such data structures Bouajjani et al.'s approach is more powerful since it can relate an arbitrary number of successive positions in a list. In contrast, the aim of our analysis is to be applicable to general heap-manipulating programs.

Sagiv et al. [23] introduce a shape analysis in which invariants are expressed in 3-valued first order logic with transitive closure (FOLTC). These invariants may combine shape and value constraints. The analysis requires user-supplied predicates, whereas our analysis does not need manual annotations; it represents a state by a set of logical structures, whereas our analysis maintains a single abstract heap, reducing the number of nodes and edges, and therefore the complexity of the overall analysis. The merging of nodes in our widening operator can be viewed as a special case of canonical abstractions.

McCloskey et al. [21] propose a framework for combining shape and numerical domains (encoded as predicates in FOLTC) in a generic way. However, users have to supply shared predicates via which the domains communicate and which

usually resemble the properties one wants to prove. In contrast, our analysis can be parameterized by arbitrary value domains without any manual overhead.

Ferrara et al. [12,13] and Fu [14] combine different heap and value analyses. Whereas their work represents a state as a heap abstraction and a single value state, our analysis attaches a value state to each edge in the heap abstraction, allowing for a precise tracking of disjunctive information. Moreover, in the value states of Ferrara et al.'s and Fu's work, different heap identifiers represent disjoint portions of the heap. This is not the case for our edge-local identifiers, which refer to memory locations already represented by abstract identifiers and which enable a precise treatment of summary nodes.

Chang et al. [6] introduce a shape analysis based on user-supplied invariants that describe data structures such as lists and trees. These invariants are used to abstract over a potentially unbounded number of concrete references. Chang and Rival [4,5] extend this work and present a framework for combining shape and numeric abstractions into a single domain. Their approach enables the precise and modular analysis of heap and numeric invariants, but relies on user-supplied properties, whereas our analysis does not require manual annotations.

Marron et al. [20] introduce heap abstractions that are similar to the graphs representing abstract heaps in our work. In fact, the formalization of the concretization function in Sec. 3.2 is inspired by their work. However, there are important technical differences. In particular, Marron et al.'s analysis maintains a normal form, which makes their lattice finite, but loses information when merging two heap graphs. In contrast, we deal with an infinite lattice, but preserve some of this information. Moreover, Marron et al.'s heap graphs track specific aliasing predicates (such as injectivity of fields or tree shapes), but no value information. Finally, the purpose of their work is to provide a high-level abstraction of *concrete* runtime heaps, whereas we propose an abstract domain and an abstract semantics for a static code analysis.

## 8 Conclusion

In this paper, we have presented a static analysis that infers complex invariants combining shape and value information. The analysis is parametric in the underlying value domain, allowing for different trade-offs between precision and efficiency. A key innovation of our analysis is the introduction of edge-local identifiers to track value information about the source and target of a pointer, which allows it to infer inductive invariants such as sortedness of a linked list. The analysis has been implemented in the static analyzer Sample. Our experiments demonstrate its effectiveness.

*Acknowledgments.* We are grateful to Uri Juhasz and Alexander Summers for numerous discussions, to John Boyland for helpful comments on a draft of this paper, and to Severin Heiniger for his contributions to the implementation.

## References


1. A. Bouajjani, C. Dragoi, C. Enea, and M. Sighireanu. On inter-procedural analysis of programs with lists and data. In *PLDI*. ACM, 2011.
2. A. Bouajjani, C. Dragoi, C. Enea, and M. Sighireanu. Abstract domains for automated reasoning about list-manipulating programs with infinite data. In *VMCAI*. Springer, 2012.
3. H. Bunke, X. Jiang, and A. Kandel. On the minimum common supergraph of two graphs. *Computing*, 65(1):13–25, Aug. 2000.
4. B.-Y. E. Chang and X. Rival. Relational inductive shape analysis. In *POPL*. ACM, 2008.
5. B.-Y. E. Chang and X. Rival. Modular construction of shape-numeric analyzers. In *David A. Schmidt's 60th Birthday Festschrift*, EPTCS, 2013.
6. B.-Y. E. Chang, X. Rival, and G. C. Necula. Shape analysis with structural invariant checkers. In *SAS*. Springer, 2007.
7. P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*. ACM, 1977.
8. P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *POPL*. ACM, 1978.
9. I. Dillig, T. Dillig, and A. Aiken. Fluid updates: Beyond strong vs. weak updates. In *ESOP*. Springer, 2010.
10. I. Dillig, T. Dillig, and A. Aiken. Precise reasoning for programs using containers. In *POPL*. ACM, 2011.
11. S. Drossopoulou, A. Francalanza, P. Müller, and A. J. Summers. A unified framework for verification techniques for object invariants. In *ECOOP*. Springer, 2008.
12. P. Ferrara. Generic combination of heap and value analyses in abstract interpretation. In *VMCAI*. Springer, 2014.
13. P. Ferrara, R. Fuchs, and U. Juhasz. TVAL+ : TVLA and value analyses together. In *SEFM*. Springer, 2012.
14. Z. Fu. Modularly combining numeric abstract domains with points-to analysis, and a scalable static numeric analyzer for Java. In *VMCAI*. Springer, 2014.
15. D. Gopan, F. DiMaio, N. Dor, T. W. Reps, and S. Sagiv. Numeric domains with summarized dimensions. In *TACAS*. Springer, 2004.
16. B. Jeannet and A. Miné. Apron: A library of numerical abstract domains for static analysis. In *CAV*. Springer, 2009.
17. T. Lev-Ami and M. Sagiv. TVLA: A system for implementing static analyses. In *SAS*. Springer, 2000.
18. F. Logozzo. Automatic inference of class invariants. In *VMCAI*. Springer, 2004.
19. R. Manevich, M. Sagiv, G. Ramalingam, and J. Field. Partially disjunctive heap abstraction. In *SAS*. Springer, 2004.
20. M. Marron, C. Sánchez, Z. Su, and M. Fähndrich. Abstracting runtime heaps for program understanding. *IEEE Trans. Software Eng.*, 39(6):774–786, 2013.
21. B. McCloskey, T. W. Reps, and S. Sagiv. Statically inferring complex heap, array, and numeric invariants. In *SAS*. Springer, 2010.
22. A. Miné. The octagon abstract domain. *Higher Order Symbol. Comput.*, 2006.
23. M. Sagiv, T. Reps, and R. Wilhelm. Parametric shape analysis via 3–valued logic. In *POPL*. ACM, 1999.

## A Graph Theory

**Definition 1 (Graph).** *Let $L$ be a finite set of labels of nodes and edges. A labeled* ***graph*** *is a triple $g = (V, E, \eta, \mu)$, where*

- *$V$ is a set of vertices*
- *$E \subseteq V \times L \times V$*
- *$\eta : V \to L$ is a function assigning labels to vertices*

**Definition 2 (Subgraph - $\subseteq$).** *Let $g = (V, E, \eta)$ be a graph. We say that $g' = (V', E', \eta')$ is a* ***subgraph*** *of $g$, $g' \subseteq g$, iff*

- *$V' \subseteq V$*
- *$E' \subseteq E$*
- *$\eta'(v') = \eta(v')$ for all $v' \in V'$*

**Definition 3 (Graph Isomorphism).** *Let $g$ and $g'$ be graphs. A* ***graph isomorphism*** *between $g$ and $g'$ is a bijective mapping $i : V \to V'$ such that*

- *$\eta(v) = \eta'(i(v))$ for all $v \in V$*
- *for any edge $e = (u, l, v) \in E$ there exists an edge $e' = (i(u), l, i(v)) \in E'$, and for any edge $e' = (u', l, v') \in E'$ there exists an edge $e = (i^{-1}(u'), l, i^{-1}(v')) \in E$.*

**Definition 4 (Subgraph Isomorphism).** *If there exists a graph $g' \subseteq g$ and $g'$ is isomorphic to $g_1$, then we say that $g_1$ is* subgraph isomorphic *to $g$.*

**Definition 5 (Maximum Common Subgraph - *mcs*).** *Let $g_1 = (V_1, E_1, \eta_1)$ and $g_2 = (V_2, E_2, \eta_2)$ be graphs. We say that $g = (V, E, \eta)$ is a* ***common subgraph*** *of $g_1$ and $g_2$ iff there exist $g'_1 \subseteq g_1$ and $g'_2 \subseteq g_2$ such that $g$ is graph isomorphic to $g'_1$ and $g'_2$. We call $g$ a* ***maximum common subgraph*** *of $g_1$ and $g_2$, $mcs(g_1, g_2)$, if there exists no other common subgraph of $g_1$ and $g_2$ that has, firstly, more vertices and, secondly, more edges than $g$.*