Storeless Semantics and Alias Logic

Marius Bozga, Radu Iosif, Yassine Lakhnech

Storeless Semantics and Alias Logic

Marius Bozga, Radu Iosif and Yassine Lakhnech
VERIMAG
2 Avenue de Vignate
38610 Gieres, France
{Marius.Bozga, Radu.Iosif, Yassine.Lakhnech}@imag.fr

Abstract

Pioneering work has been done by Jonkers [11] to define a semantics of pointer manipulating programs that is abstract in the sense of ignoring low-level aspects such as dangling pointers and garbage objects. We explore the principles of such storeless semantics from a logical point of view, first defining a simple logic to completely characterize heap structures up to isomorphism. Second, we extend this language to a full-blown alias logic (AL) that allows to express regular properties of unbounded heap structures. Along the development, we present an operational storeless semantics and give sound and complete total correctness axioms for deterministic programs in the form of Hoare triples, using AL.

1 Introduction

This paper provides a formalism for describing properties of linked data structures such as lists, trees and graphs. It also provides an associated program logic for reasoning about programs that destructively update such data structures. The introduced logic, which is called AL for Alias Logic, allows to describe aliases naturally. Moreover, the truth of formulae of this logic is insensitive to garbage collection. Reynolds argues in [14] that program logics such as Hoare calculus based on low level view of storage are incompatible with garbage collection. On the other hand, garbage collection is an essential feature of runtime environments of languages as Lisp, ML and Java.

We consider a set of instructions for altering stacks and heaps and provide three different semantics for these instructions. We also study the relationships between these semantics (See Fig 1).

We start from a low level concrete semantics, denoted by $[\![\cdot]\!]_1$ in Fig. 1, where the state of a program is described by a stack and heap. We then show that this semantics is fully abstract in the sense that it does not distinguish between states that are equal up-to the identity of locations. In other words, renaming of locations induces an equivalence relation that is a bisimulation (Diagram (1) in Fig 1). Moreover, we provide a logic characterization of this equivalence relation.

This semantics is, however, sensitive to garbage and dangling pointers from which we would like to abstract. Therefore, we present a store-less semantics $[\![\cdot]\!]_3$ where a heap is modeled as a structure that is a set of regular languages satisfying additional conditions. The idea behind this semantics is that each heap element can be represented by a Rabin-Scott automaton with a single accepting state. Our storeless semantics is akin to Jonker's [11] and Deutsch's [4] semantics. In contrast to them, however, our semantics operates on structures representing equivalence classes of heaps. The additional conditions ensures this property. Thus, the structures we introduce can be seen as symbolic representations of equivalence classes of states and the semantics operating on them as a symbolic semantics. We show that it is equivalent to the concrete semantics modulo renaming of locations (Diagram (2) in Fig. 1).

The third semantics we introduce is an axiomatic semantics, i.e., a Hoare logic-like proof system. The main feature of this proof system is that it allows to prove properties of programs that are insensitive to garbage. Thus, our proof system is compatible with garbage collection. The problem of designing program logics enjoying this property has been studied by Hoare and Jifeng in [10]. To obtain the desired result they explicitly introduce in the assertion language an operator which models garbage collection, i.e., transforms a state into a garbage free state. A drawback of this approach, as emphasized by Calcagno, O'Hearn and Bornat in [2], is that one has to explicitly carry around a state parameter. These authors propose in [2] an alternative approach which consists in altering the semantics of the assertion language

in such a way that the program logic becomes compatible with garbage collection. This approach is based on using a possible-world interpretation for existential quantification, where the current heap is the world. A similar approach is followed in [3]. Our approach is based on an assertion language that allows to reason explicitly about aliasing. It is called Alias Logic, AL for short. AL allows to describe sets of symbolic structures as introduced in the symbolic garbage insensitive semantics. Interpreting the logic on symbolic structures gives as the property that it is insensitive to garbage and allows for designing a weakest precondition based program logic that enjoys the same property. We prove soundness and completeness of our program logic (Diagram (3)) and show its applicability on a well known example.

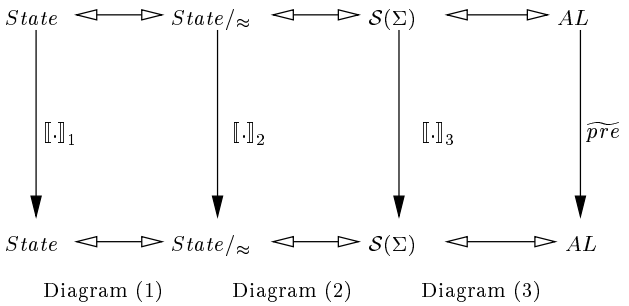

Figure 1: Summary of the contributions

2 A Simple Heap Logic

This section is dedicated to the definition of a simple logic for performing observations on the heap. This language is propositional logic, where the atomic terms describe pointer aliasing. We name this language heap logic (HL). The abstract syntax of HL is given in Figure 2 (left), and its store-based denotational semantics, in Figure 2 (right). We denote by Σ the set of all pointer variables and by Σ^+ the set of all access paths i.e., non-empty sequences of pointer variables. Since we are interested only in describing shapes, we consider that all variables in a program are pointers.

For the semantics, we consider a set Loc of memory locations. This set is infinite but countable and we represent it as $\{l_0, l_1, \ldots\}$. As usual, a *store* is a partial mapping between variables and values. Since in our case all variables are pointers, all values are locations. We express the fact that a mapping f is undefined in a point x by $f(x) = \perp$. For a given set A, the notation $A_\perp$ means $A \cup \{\perp\}$; we always assume $\perp \notin A$. A *heap* is a partial mapping between locations and stores. More precisely, given a heap h and a location l, the expression $h(l)$ denotes a store or $\perp$ if $l = \perp$. For a variable x, the notation $h(l, x)$ stands for $h(l)(x)$. We may refer to the stores in the range of a heap as to *objects*. We assume that heap functions are strict i.e., $h(\perp) = \perp$. The denotation of the terms and expressions of our language is given with respect to *states*. A state st is a pair store-heap (s, h), in which the first component represents the values of *global* variables i.e., variables that are not heap-allocated. Notice that in this setting the denotation of a string in a state $[\![\sigma]\!]_{s,h}$ is a location, or $\perp$ if the access path σ is dangling in the state. Two paths are said to be *aliased* if and only if they lead to the same location. The propositional logic connectives are defined as usual.

To simplify reasoning about states, we consider the function *Reach*, which maps a state st and a location l to the set of all access paths reaching l in st. Also, the function *Reachable* gives, for a state the set of reachable locations within it.

$$\begin{aligned} Reach &: State \times Loc \to \mathcal{P}(\Sigma^+) \\ Reach(st, l) &\triangleq \{\sigma \in \Sigma^+ \mid [\![\sigma]\!]_{st} = l\} \end{aligned} \tag{1}$$

$$\begin{aligned} Reachable &: State \to \mathcal{P}(Loc) \\ Reachable(st) &\triangleq \{l \in Loc \mid \exists \sigma \in \Sigma^+ \ [[\![\sigma]\!]_{st} = l]\} \end{aligned} \tag{2}$$

Using these functions, we can now define the notions of *total* and *garbage free* states.

$$
\begin{array}{rcl}
u, v, x & \in & \Sigma \\
\sigma, \tau, \theta & \in & \Sigma^* \\
\\
f & := & \sigma \diamond \tau \mid f_1 \vee f_2 \mid \neg f
\end{array}
\qquad
\begin{array}{rcl}
s & \in & Store \triangleq \Sigma \mapsto Loc_\perp \\
h & \in & Heap \triangleq Loc \mapsto Store_\perp \\
st & \in & State \triangleq Store \times Heap \\
\\
[\![\sigma]\!]_{s,h} & \triangleq & \begin{cases} s(v) & \text{if } \sigma \equiv v \\ h([\![\tau]\!]_{s,h}, v) & \sigma \equiv \tau v \\ \perp & \text{otherwise} \end{cases} \\
\\
[\![\sigma \diamond \tau]\!]_{st} & \triangleq & [\![\sigma]\!]_{st} = [\![\tau]\!]_{st}
\end{array}
$$

Figure 2: The Heap Logic HL

Definition 1 *A state $s, h \in Store \times Heap$ is said to be:*

- finite *if and only if the domain of h is finite.*
- total *if and only if $Reachable(s, h) \subseteq dom(h)$, and,*
- garbage free *if and only if $dom(h) \subseteq Reachable(s, h)$.*

Next, we define *garbage collection* on states as the function gc which restricts the domain of the heap h in a state (s, h) to the set of the reachable locations. Formally:

$$
\begin{array}{rcl}
gc & : & State \rightarrow State \\
gc(s, h) & \triangleq & s, h \downarrow_{Reachable(s,h)}
\end{array}
\tag{3}
$$

2.1 An Imperative Programming Language

We consider a simple language of atomic statements and let programs be sequences of statements. The abstract syntax of statements is shown in Figure 3. The first statement sets the left-hand side variable to null, which may cause the deletion of non-reachable objects by the garbage collector. The second statement allocates a fresh cell for further uses. The third statement is the assignment operation between variables.

$$
\begin{array}{rcll}
Stmn & := & \sigma := null & \\
 & \mid & \sigma := new & \\
 & \mid & \sigma := \sigma' & \text{where } \sigma \in \Sigma^+
\end{array}
$$

Figure 3: Syntax of Statements

The operational semantics is given in Figure 4. For each statement, we distinguish two cases, depending on the length of the left-hand side path σ. If $|\sigma| = 1$ (Figure 4, left), the statement changes the value of a local variable. Otherwise (Figure 4, right) the statement affects a heap-allocated variable.

Note that, due to the lack of an explicit delete operation, all states generated by this semantics starting with a total state are total. Moreover, they are also garbage free, since the gc function is invoked for each transition. Let $\leadsto \subseteq State \times Stmn \times State$ be the transition relation defined by the rules in Figure 4, and $\leadsto^*$ denote its transitive closure. For a sequence of statements $\omega \in Stmn^*$ we write $st \overset{\omega}{\leadsto}^* st'$ if st' is obtained from st by applying the sequence of operations ω.

3 On Heap Isomorphism

The store-based model of computation is redundant. Intuitively, two states that differ only by a renaming of locations are equivalent with respect to formulas written in HL. Moreover, this equivalence

$$\frac{}{s, h \overset{v:=null}{\rightsquigarrow} gc(s[v \mapsto \bot], h)} \quad (4)$$

$$\frac{l_{new} \notin dom(h) \quad h_{new} = h[l_{new} \mapsto \lambda x.\bot]}{s, h \overset{v:=new}{\rightsquigarrow} gc(s[v \mapsto l_{new}], h_{new})} \quad (5)$$

$$\frac{[\![\theta]\!]_{s,h} = l \qquad l \neq \bot}{s, h \overset{v:=\theta}{\rightsquigarrow} gc(s[v \mapsto l], h)} \quad (6)$$

$$\frac{[\![\tau]\!]_{s,h} = l, \;\; l \neq \bot, \;\; s' = h(l)[v \mapsto \bot]}{s, h \overset{\tau.v:=null}{\rightsquigarrow} gc(s, h[l \mapsto s'])} \quad (7)$$

$$\frac{[\![\tau]\!]_{s,h} = l, \;\; l \neq \bot \qquad l_{new} \notin dom(h) \quad h_{new} = h[l_{new} \mapsto \lambda x.\bot] \quad s' = h(l)[v \mapsto l_{new}]}{s, h \overset{\tau.v:=new}{\rightsquigarrow} gc(s, h_{new}[l \mapsto s'])} \quad (8)$$

$$\frac{[\![\tau]\!]_{s,h} = l \quad l \neq \bot \quad [\![\theta]\!]_{s,h} = l' \quad l' \neq \bot \quad s' = h(l)[v \mapsto l']}{s, h \overset{\tau.v:=\theta}{\rightsquigarrow} gc(s, h[l \mapsto s'])} \quad (9)$$

Figure 4: Operational Semantics

is a bisimulation [12] i.e., it has been shown that the direct successors of two equivalent states are also equivalent [8]. We formalize this notion as follows.

Definition 2 (Isomorphism) *Two states are said to be isomorphic, denoted by $s, h \approx s', h'$, if and only if there exists a bijection $\pi : Loc \to Loc$ such that:*

- $s' = \lambda v.\pi(s(v))$, *and*
- $h' = \lambda l.\lambda v.\pi(h(\pi^{-1}(l), v))$.

To establish whether two paths in the heap are aliased, we need to compare their denotations (Figure 2). The outcome of this observation is however independent of the actual values of paths, and can be shown to be invariant under isomorphic transformations of states. Hence we can associate each location in the domain of the heap an invariant set of paths, which is the set of all incoming paths.

The following theorem relates states isomorphism and the heap logic HL: two isomorphic states are indistinguishable by any HL formula, and viceversa, two states that are indistinguishable by HL are isomorphic. Due to space limitations, all proofs are deferred to [1].

Theorem 1 *Let $st, st' \in State$ denote any two total garbage free states, and f denote any HL formula over the alphabet Σ. Then we have:*

$$st \approx st' \iff \forall f \; [[\![f]\!]_{st} = [\![f]\!]_{st'}]$$

This result shows that, despite its simplicity, the HL language is powerful enough to distinguish non-isomorphic states. We also conjecture a stronger result: for each finite total garbage free state $st \in State$, there exists an HL formula f^{st} that characterizes st up to isomorphism. Investigating this issue into further detail is somehow outside the scope of this paper and is considered for an extended version.

4 A Storeless Model

As mentioned before, the level of detail in the store-based model for the heap is too high. This allows to distinguish between semantically equivalent states. Moreover, this model retains information related to garbage and dangling pointers, from which we would like to abstract. Hence, we introduce a symbolic representation, based on the theory of regular languages. In the new, *storeless* model, a heap is a collection of languages. The idea behind this representation is that each object is the language accepted by the heap graph, viewed as an automaton with that object as a unique final state.

Definition 3 (Storeless structure) *A storeless structure $\Gamma \in \mathcal{P}(\Sigma^+)$ is either the empty set or a set $\{S_1, S_2, \ldots, S_n\}$ satisfying the following conditions, for all $1 \leq i, j \leq n$ and for some $1 \leq k \leq n$:*

(C1) non-emptiness: $S_i \neq \emptyset$,

(C2) determinism: $i \neq j \Rightarrow S_i \cap S_j = \emptyset$,

(C3) prefix closure and right regularity: $\forall \sigma \in S_i \; [\exists \tau, \theta \in \Sigma^+ [\sigma = \tau\theta] \Rightarrow \tau \in S_k \wedge S_k\theta \subseteq S_i]$.

Let $\mathcal{S}(\Sigma)$ denote the set of all storeless structures over the alphabet Σ. An alternative way of defining a storeless structure is by considering an equivalence relation (alias) on the set of all heap paths. This is the approach taken by Jonkers [11] and Deutsch [5]. By requiring that the equivalence relation be right-regular, they obtain that each language in the heap is recognizable by a finite automaton. Instead, we choose to represent equivalence classes explicitly and impose the right-regularity condition as (C3). Notice that our structures are deterministic (C2) since a path is not allowed to belong to two different sets. Moreover, we exclude empty sets (C1) from our representation; empty sets could serve as an abstract representation of (all) garbage objects which we have chosen to ignore.

By requiring that structures are formed only with non-empty paths ($\Gamma \in \mathcal{P}(\Sigma^+)$) we represent only *rooted* graphs i.e., graphs in which there are no incoming paths towards an initial node. This constraint suits our model of heap well since all paths of lenght one denote store (local) variables and we allow cycles only in the heap.

Let us point now the discussion towards proving the soundness and completness of the three rules characterizing the storeless semantics. We will do so by relating storeless structures to the previous store-based semantics.

Definition 4 (Correspondence) *Let $\Gamma \in \mathcal{S}(\Sigma)$ be a structure and $st = s, h \in State$ be a state. We say that Γ and st correspond, denoted $\Gamma \leftrightharpoons st$, if and only if there exists a bijection $\pi : dom(h) \to \Gamma$ such that:*

1. *for all $u \in \Sigma$, $s(u) = l$ if and only if $u \in \pi(l)$.*
2. *for all $l, l' \in dom(h)$ and $u \in \Sigma$, $h(l, u) = l'$ if and only if $\pi(l)u \subseteq \pi(l')$.*

In principle, we cannot represent a state with garbage or dangling pointers by a corresponding storeless structure, without violating condition (C3). The following lemma shows a method of transforming a store-based state into a storeless structure. Moreover, it states that there is only one way to do so. Our construction is in fact the equivalent of the left quotienting in automata theory [7].

Lemma 1 *Let $st = s, h \in State$ be a total garbage free state. Then the set $\{Reach(st, l) \mid l \in dom(h)\}$ is a storeless structure. Moreover, if $\Gamma \in \mathcal{S}(\Sigma)$ is a storeless structure such that $st \leftrightharpoons \Gamma$, then $\Gamma = \{Reach(st, l) \mid l \in dom(h)\}$.*

This result implies that a storeless structure can be used as a *canonical* symbolic representation of isomorphic states. The following theorem is the first important result of this section. It postulates the correctness of conditions (C1), (C2) and (C3) by finding, for each store-based state a corresponding storeless structure and viceversa, for each storeless structure a corresponding state.

Theorem 2 *For each total garbage free state $st \in State$ there exists a unique structure $\Gamma \in \mathcal{S}(\Sigma)$ such that $st \leftrightharpoons \Gamma$. Dually, for each structure $\Gamma \in \mathcal{S}(\Sigma)$ there exists a garbage free total state $st \in State$ such that $st \leftrightharpoons \Gamma$.*

This theorem expresses the fact that our definition of storeless structures has the same expressive power as the right-regular equivalences used by Jonkers [11] and Deutsch [5] and the trace model described by Hoare and Jifeng [10]: we are now capable to describe rooted directed labeled graphs. However, the explicit use of regular languages, enables us to give an operational semantics on storeless structures that is easier to understand and implement using finite regular automata. This is the discussion point of the next subsection. Next, in Section 5, we develop an alias logic that uses regular expressions too, and whose decision procedure relies on checking emptiness of finite regular automata.

4.1 Storeless Operational Semantics

Having defined storeless structures as a symbolic representation for states, we can now define program actions as operations on regular languages. In order to simplify the presentation, we define first three primitive transformations, and later, present the full semantics using compositions of the primitive actions. Informally, *rem* (10) describes the effect of removing an arc v from a graph node represented by the language S. Notice that eliminating a single arc removes a possibly infinite number of paths from the structure, potentially introducing garbage objects. These objects are automatically represented by an empty set, which is finally eliminated from the structure. Next, *new* (11) is used to model object creation. From an origin node S we create a new node represented by Sv and add it to the structure. The most complex operation is *add* (12) which adds a (possibly new) arc between two nodes S and T. This complexity is an inherent consequence of the fact that cycles might be introduced. Nevertheless, all

the transformations occuring in *add* are easily implemented via automata. Interestingly, our semantics for *add* matches exactly Hoare and Jifeng's semantics for pointer swing [10]. The rest of this section is concerned with proving the correctness of the storeless operational semantics.

$$\begin{aligned} rem &: \mathcal{P}(\Sigma^*) \times \Sigma \to \mathcal{P}(\Sigma^*) \to \mathcal{P}(\Sigma^*) \\ rem(S,v) &\triangleq \lambda\Gamma.\{X \setminus Sv\Sigma^* \mid X \in \Gamma\} \setminus \{\emptyset\} \end{aligned} \tag{10}$$

$$\begin{aligned} new &: \mathcal{P}(\Sigma^*) \times \Sigma \to \mathcal{P}(\Sigma^*) \to \mathcal{P}(\Sigma^*) \\ new(S,v) &\triangleq \lambda\Gamma.\Gamma \cup \{Sv\} \end{aligned} \tag{11}$$

$$\begin{aligned} add &: \mathcal{P}(\Sigma^*) \times \Sigma \times \mathcal{P}(\Sigma^*) \to \mathcal{P}(\Sigma^*) \to \mathcal{P}(\Sigma^*) \\ add(S,v,T) &\triangleq \lambda\Gamma.\{\chi^{S,v,T}(X) \mid X \in \Gamma\} \\ \text{where } \chi^{S,v,T}(X) &\triangleq X \cup Sv((T^{-1}S)v)^*(T^{-1}X) \end{aligned} \tag{12}$$

Notice first that the primitive operations can be applied to any languages $S, T \subseteq \Sigma^*$, any symbol $v \in \Sigma$ and any set of languages $\Gamma \in \mathcal{P}(\Sigma^*)$. The following three lemmas are then the first steps in our correctness proof. Assuming that we start with a storeless structure $\Gamma \in \mathcal{S}(\Sigma)$, and two languages in $S, T \in \Gamma_\epsilon$, we postulate necessary and sufficient conditions for the result of the three operations to be a valid storeless structure. In particular, there are no side conditions for *rem*.

Lemma 2 *Let $\Gamma \in \mathcal{S}(\Sigma)$ be a storeless structure and $S \in \Gamma_\epsilon$ be a set. Then, for all symbols $v \in \Sigma$ we have $rem(S, v, \Gamma) \in \mathcal{S}(\Sigma)$.*

The *new* operation yields a correct result if and only if there are no common paths between the new node and some existing node or the new node already existed in the structure, in which case nothing is changed. Notice that, if this condition is violated the result will necessarily be non-deterministic, hence violate the (C2) condition.

Lemma 3 *Let $\Gamma \in \mathcal{S}(\Sigma)$ be a storeless structure and $S \in \Gamma_\epsilon$ be a set. Then, for all symbols $v \in \Sigma$ we have $new(S, v, \Gamma) \in \mathcal{S}(\Sigma)$ if and only if for all $T \in \Gamma$, either $Sv \cap T = \emptyset$ or $T = Sv$.*

Before proving soundness with respect to the store-based model, we give a fixpoint formulation for the $\chi^{S,v,T}$ function, used in the definition of *add* (12); while the ready-made formula (12) is easy to implement, the fixpoint formulation will be of more use in reasoning about the correcntess of *add*. To improve readability, we skip the superscripts of χ. Given three languages S, T and X, we define the following function:

$$\xi_X(x) \triangleq X \cup Sv(T^{-1}x) \tag{13}$$

Now let us show that $\chi(X) = \text{fix}\, \xi_X$. Denote $Y \triangleq T^{-1}\chi(X)$. We have:

$$\begin{aligned} Y &= T^{-1}(X \cup Sv(T^{-1}\chi(X))) \\ &= T^{-1}X \cup T^{-1}(Sv(T^{-1}\chi(X))) \\ &= T^{-1}X \cup (T^{-1}S)vY \end{aligned}$$

We have used that $T^{-1}(Sv) = (T^{-1}S)v$, which is easily checked. Also $\epsilon \notin (T^{-1}S)v$, and, by Arden's lemma, we obtain that $Y = ((T^{-1}S)v)^*(T^{-1}X)$ is the unique solution to the above equation. Since $\chi(X) = X \cup SvY$, we have the result. It is easy to check that the ξ_X function is affine i.e., for any $\mathcal{Y} \in \mathcal{P}(\mathcal{P}(\Sigma^*))$ we have $\xi_X(\bigcup \mathcal{Y}) = \bigcup \xi_X(\mathcal{Y})$. Hence $\chi(X) = \bigcup_{i \geq 0} \xi_X^i(\emptyset)$.

The *add* operation yields a correct result if and only if the newly added arc from S to T does not already exist between S and a node different than T. If this condition would be violated the resulting structure would be non-deterministic.

Lemma 4 *Let $\Gamma \in \mathcal{S}(\Sigma)$ be a storeless structure and $S \in \Gamma_\epsilon$, $T \in \Gamma$ be two sets. Then, for all symbols $v \in \Sigma$ we have $add(S, v, T, \Gamma) \in \mathcal{S}(\Sigma)$ if and only if for all $T' \in \Gamma$, either $Sv \cap T' = \emptyset$ or $T' = T$.*

Figure 5 presents the operational semantics of the three statements defined by the syntax in Figure 3. As in the definition of the store-based semantics, we treat separately the case where the left-hand side of the assignment is a local variable v or a path $\tau.v$ of length two or more. In the first case, the

$$\frac{}{\Gamma \overset{v:=null}{\hookrightarrow} rem(\epsilon, v, \Gamma)} \quad (14) \qquad \frac{\exists S \in \Gamma\ [\tau \in S]}{\Gamma \overset{\tau.v:=null}{\hookrightarrow} rem(S, v, \Gamma)} \quad (17)$$

$$\frac{}{\Gamma \overset{v:=new}{\hookrightarrow} (new(\epsilon, v) \circ rem(\epsilon, v))(\Gamma)} \quad (15) \qquad \frac{\exists S \in \Gamma\ [\tau \in S]}{\Gamma \overset{\tau.v:=new}{\hookrightarrow} (new(S, v) \circ rem(S, v))(\Gamma)} \quad (18)$$

$$\frac{\exists S \in \Gamma\ [\theta \in S]}{\Gamma \overset{v:=\theta}{\hookrightarrow} (add(\epsilon, v, S) \circ rem(\epsilon, v))(\Gamma)} \quad (16) \qquad \frac{\exists S, T \in \Gamma\ [\tau \in S \wedge \theta \in T]}{\Gamma \overset{\tau.v:=\theta}{\hookrightarrow} (add(S, v, T) \circ rem(S, v))(\Gamma)} \quad (19)$$

Figure 5: Storeless Operational Semantics

first argument of *rem*, *new* and *add* is ϵ. Otherwise, we need to identify a node S in the source structure to which the v variable belongs. In order to keep the semantics small, we do not treat null pointer dereferencing errors.

To use the primitive operations previously defined, we need to make sure that the side conditions stated in Lemma 3 and 4 are actually met. This is accomplished using *rem* before *new* or *add* to first clear the "inconsistent" paths from the structure. In the case of pointer assignment (rules (16) and (19)), removing first some paths from the structure leads to a well known problem: if the left-hand side path is a prefix of the right-hand side and there are no other incoming paths to the right-hand side node, then this node will be eliminated before the assignment takes place. A solution proposed in the literature [11, 10] uses *fresh* paths that are explicitly added to the right-hand side node and removed after the assignment. To simplify the semantics, we precompile our program introducing a fresh temporary variable whenever $\tau.v$ is a prefix of θ i.e., transforming the statement $\tau.v := \theta$ into $v_{fresh} := \theta;\ \tau.v := v_{fresh}$. Notice that the prefix condition is just a syntax check easily performed during program parsing. Let $\hookrightarrow \subseteq \mathcal{S}(\Sigma) \times Stmn \times \mathcal{S}(\Sigma)$ be the relation defined by the rules in Figure 5. As in the case of store-based semantics, we write $\Gamma \overset{\omega}{\hookrightarrow}^* \Gamma'$ for its transitive closure.

Theorem 3 *Let $m \in Stmn$ be a statement, $st \in State$ be a total garbage free state and $\Gamma \in \mathcal{S}(\Sigma)$ be a structure such that $st \leftrightharpoons \Gamma$. If $st \overset{m}{\rightsquigarrow} st'$ for some $st' \in State$ then there exists $\Gamma' \in \mathcal{S}(\Sigma)$ such that $\Gamma \overset{m}{\hookrightarrow} \Gamma'$ and $st' \leftrightharpoons \Gamma'$. Dually, if $\Gamma \overset{m}{\hookrightarrow} \Gamma'$ for some $\Gamma' \in \mathcal{S}(\Sigma)$, then there exists $st' \in State$ such that $st \overset{m}{\rightsquigarrow} st'$ and $st' \leftrightharpoons \Gamma'$.*

This proves soundness of the storeless semantics with respect to the classical store-based model.

5 Alias Logic

In this section we describe the full-blown alias logic AL which embeds our initial heap logic HL and which is next developped into a program logic. Having only about aliases is not expressive enough for a precondition calculus. Therefore AL has a modality operator which allows to specify into which node in the graph paths may flow. Since nodes are given as recognizable languages, the decision of modalities boils down to deciding language emptiness of product automata.

Let Var denote a set of free variables, $Reg(Var, \Sigma)$ denote the set of all regular expressions over Σ containing variables from Var, and $Term(Var, \Sigma)$ denote the set of all terms built out of regular expressions, equality operator and a modality operator together with the classic connectives of first-order logic. Figure 6 gives the syntax (left) and the semantics (right) of AL.

In the following, we will refer to the $\langle \rho_1 \rangle \rho_2$ terms as to *modalities*. To ease notation, we introduce some syntactic shortcuts: $\rho_1 \setminus \rho_2 \triangleq \rho_1 \cap \bar{\rho_2}$, $\emptyset \triangleq \rho \setminus \rho$ (for some $\rho \subseteq \Sigma^*$), $\rho_1 \subseteq \rho_2 \triangleq \rho_1 \cap \bar{\rho_2} = \emptyset$, and $\rho_1^{-1}\rho_2 \triangleq \cup\{\rho \mid \rho_1 \cdot \rho \cap \rho_2 \neq \emptyset\}$. Using disjunction, negation and existential quantification, we can define the rest of logical connectives as $true \triangleq \varphi \vee \neg\varphi$, $false \triangleq \neg true$, $\varphi_1 \wedge \varphi_2 \triangleq \neg(\neg\varphi_1 \vee \varphi_2)$, $\varphi_1 \Rightarrow \varphi_2 \triangleq \neg\varphi_1 \vee \varphi_2$, and $\forall X\ [\varphi] \triangleq \neg\exists X\ [\neg\varphi]$.

Given a term, the *free* function returns the set of free variables occuring within its regular expressions. This function is usually defined by induction on the structure of terms i.e., $free(\bullet\varphi) \triangleq free(\varphi)$ for each unary term and $free(\varphi_1 \bullet \varphi_2) \triangleq free(\varphi_1) \cup free(\varphi_2)$ for each binary term. Existential quantification eliminates free variables i.e., $free(\exists X\ [\varphi]) \triangleq free(\varphi) \setminus \{X\}$. A structure Γ is said to be a *model* for a

$$
\begin{array}{rcl}
u & \in & \Sigma \\
X & \in & Var \\
\rho & \in & Reg(Var, \Sigma) \\
\varphi & \in & Term(Var, \Sigma)
\end{array}
$$

$$
\begin{array}{rcl}
\rho & ::= & u \mid \Sigma \mid X \mid \rho_1 \cdot \rho_2 \mid \rho^* \\
 & & \mid \rho_1 \cup \rho_2 \mid \rho_1 \cap \rho_2 \mid \bar{\rho} \\
\varphi & ::= & \rho_1 = \rho_2 \mid \langle \rho_1 \rangle \rho_2 \\
 & & \mid \varphi_1 \vee \varphi_2 \mid \neg\varphi \mid \exists X\ [\varphi]
\end{array}
$$

$$
\begin{array}{rcl}
\Gamma & \in & \mathcal{S}(\Sigma) \\
free & : & Reg(Var, \Sigma) \cup Term(Var, \Sigma) \rightarrow \mathcal{P}(Var) \\
\nu & : & Var \rightarrow \Sigma^*
\end{array}
$$

$$
\begin{array}{rcl}
[\![\rho]\!]_\nu & \triangleq & \rho[\vec{X}/\nu(\vec{X})] \text{ where } \vec{X} = free(\rho) \\
[\![\rho_1 = \rho_2]\!]_\nu & \triangleq & [\![\rho_1]\!]_\nu = [\![\rho_2]\!]_\nu \\
[\![\langle \rho_1 \rangle \rho_2]\!]_{\Gamma,\nu} & \triangleq & [\![\rho_1]\!]_\nu \in \Gamma \wedge [\![\rho_1]\!]_\nu \cap [\![\rho_2]\!]_\nu \neq \emptyset \\
[\![\exists X\ [\varphi]]\!]_\nu & \triangleq & \exists \rho \in \Sigma^*\ [\![\varphi]\!]_{[X \rightarrow \rho]\nu}
\end{array}
$$

Figure 6: The Alias Logic AL

term φ, denoted $\Gamma \models \varphi$ if and only if $[\![\varphi]\!]_{\Gamma,\lambda X.\epsilon}$ is true. We define *pure assertions* to be formulas not containing modalities. Notice that the semantics of a pure assertion is given independently of a heap structure i.e., if φ is pure then either $\forall\ \Gamma \in \mathcal{S}(\Sigma)\ [\Gamma \models \varphi]$ or $\forall\ \Gamma \in \mathcal{S}(\Sigma)\ [\Gamma \not\models \varphi]$.

We can now introduce further handy notation: $\rho_1 \diamond \rho_2 \triangleq \exists X\ [\langle X \rangle \rho_1 \wedge \langle X \rangle \rho_2]$ to express the may-aliasing between two regular paths ρ_1 and ρ_2, and $in(X) \triangleq \langle X \rangle \Sigma^*$ to express the presence of node X in the storeless structure. The fact that a node Y is reachable from another node X can be defined as $reach(X, Y) \triangleq in(X) \wedge in(Y) \wedge X\Sigma^* \cap Y \neq \emptyset$. The fact that X and Y belong to a cycle is expressed as $cycle(X, Y) \triangleq reach(X, Y) \wedge reach(Y, X)$.

It is worthwhile pointing out that, by defining the $\diamond$ predicate, we have embedded the HL logic into AL. Hence AL can be used to describe sets of store-based states as well as sets of storeless configurations, according to Theorem 3. Next, we present examples of AL formulas that describe common place heap structures. The next subsection will discuss the use of AL to define the semantics of programs, in Hoare style.

Examples To show the use of AL as a language for describing the shape of pointer structures, we consider the following predicates:

$$
\begin{array}{rcl}
nclist(h, n) & \triangleq & \exists X\ [\langle X \rangle h] \wedge \forall X, Y\ [X \cup Y \subseteq h.n^* \wedge X \diamond Y \Rightarrow X = Y] \\
nshared(h_1, h_2, n) & \triangleq & \neg(h_1.n^* \diamond h_2.n^*) \\
tree(root) & \triangleq & \exists X\ [\langle X \rangle root] \wedge \forall X, Y\ [root.X \diamond root.Y \Rightarrow X = Y] \\
dag(root) & \triangleq & \exists X\ [\langle X \rangle root \wedge \forall Y, Z\ [reach(X, Y) \wedge reach(X, Z) \Rightarrow \neg cycle(Y, Z)]
\end{array}
$$

The *nclist* predicate is true in all states in which there exists a possibly empty non-circular list pointed to by the variable h. The non-circularity requirement is captured by the fact that if two paths are aliased, then they must be the equal. The *nshared* predicate is true when there is no sharing between a list starting with h_1 and a list starting with h_2, if they both use the same selector n. A *tree* structure is described by the lack of sharing within all the nodes reachable from the node pointed to by the *root* variable. To describe a *dag* we only require that there are no cycles between the nodes reachable from the top node.

AL has obvious limitations due to the expressive power of regular expressions. For instance, it would be impossible to specify a doubly-linked list or a balanced tree, since both are *counting* properties. To cope with these problems, extensions of AL towards context free and tree languages are considered as future work.

5.1 Axiomatic Semantics

Having introduced a logic to represent sets of storeless configurations, we tackle now the problem of using this logic to compute weakest preconditions. We define hereby sound and complete inference rules

to characterize the execution of the three statements we have considered throughout the paper. In this setting we deal with total correctness i.e., our assertions distinguish statements that go "wrong" from the ones that execute correctly.

Definition 5 (Weakest Precondition) *Given an AL term φ and a sequence of statements $\omega \in Stmn^*$, define $wp(\omega, \varphi) \subseteq \mathcal{S}(\Sigma)$ to be the least set such that if $\Gamma \in wp(\omega, \varphi)$ then there exists $\Gamma' \in \mathcal{S}(\Sigma)$ such that $\Gamma \overset{\omega}{\hookrightarrow}^* \Gamma'$ and $\Gamma' \models \varphi$.*

We recall a number of classical results [6] on weakest preconditions seen as predicate transformers i.e., the set $wp(\omega, \varphi)$ being characterized by a first-order predicate $\widetilde{pre}(\omega, \varphi)$. For any transition relation over a sequence of statements ω, $\widetilde{pre}$ distributes over conjunction and universal quantification i.e., $\widetilde{pre}(\omega, \varphi_1 \wedge \varphi_2) = \widetilde{pre}(\omega, \varphi_1) \wedge \widetilde{pre}(\omega, \varphi_2)$ and $\widetilde{pre}(\omega, \forall X\ [\varphi]) = \forall X\ [\widetilde{pre}(\omega, \varphi)]$. For total transition relations we have $\widetilde{pre}(\omega, \varphi) \Rightarrow \neg\widetilde{pre}(\omega, \neg\varphi)$. If, moreover, the transition relation is total and deterministic, we have that $\widetilde{pre}$ is its own dual i.e., $\widetilde{pre}(\omega, \varphi) \Leftrightarrow \neg\widetilde{pre}(\omega, \neg\varphi)$. In the latter case $\widetilde{pre}$ distributes over disjunction and existential quantification too.

These properties of $\widetilde{pre}$ for total deterministic programs allow us to define general inference rules for the precondition inductively on the structure of the postcondition. Therefore, we can first give sound and complete characterizations of *wp* for the primitive storeless operations *rem* (10) and *add* (12) in cases where the postconditions are modalities only. Then we can generalize to arbitrary postconditions using the distributivity of $\widetilde{pre}$ operators over first-order connectives in case of deterministic programs. Next, we will generalize the axioms to describe $\widetilde{pre}$ for all statements in Figure 3. In conclusion, we discuss the treatment of non-deterministic programs in AL.

Remove The following rule defines the weakest precondition of a modality formula with respect to a removal operation.

$$\{\exists X\ [X \setminus Sv\Sigma^* = T \wedge \langle X\rangle\sigma \setminus Sv\Sigma^*]\}\ \mathbf{rem}(\mathbf{S}, \mathbf{v})\ \{\langle T\rangle\sigma\} \tag{20}$$

We show now that the remove rule is sound and complete, by proving the following lemma. Although *rem* is a primitive transformation and not a statement, we still denote by $wp(rem(S, v), \varphi)$ the largest set of structures $\Gamma \in \mathcal{S}(\Sigma)$ which, under the transformation $rem(S, v)$, lead to a structure satisfying φ. According to Lemma 2, we need to assume that $S \in \Gamma$, as a side condition, otherwise the structure obtained from Γ by applying $rem(S, v)$ might not be consistent with Definition 3.

Lemma 5 *Given $\Gamma, \Gamma' \in \mathcal{S}(\Sigma)$ two structures such that $\Gamma' \models \langle T\rangle\sigma$, and $S \in \Gamma$ a set, then $\Gamma \in wp(rem(S, v), \langle T\rangle\sigma)$ if and only if $\Gamma \models \exists X\ [X \setminus Sv\Sigma^* = T \wedge \langle X\rangle\sigma \setminus Sv\Sigma^*]$.*

New This rule defines the weakest precondition of a modality with respect to the new object creation operation.

$$\{(T = Sv \wedge \sigma \cap Sv \neq \emptyset) \vee \langle T\rangle\sigma\}\ \mathbf{new}(\mathbf{S}, \mathbf{v})\ \{\langle T\rangle\sigma\} \tag{21}$$

It can be shown that the rule above is sound and complete with respect to the storeless operational semantics. Using the same abuse of notation, we denote by $wp(new(S, v), \varphi)$ the largest set of structures $\Gamma \in \mathcal{S}(\Sigma)$ which, under the transformation $new(S, v)$, lead to a structure satisfying φ. We need to also assume the necessary and sufficient condition from Lemma 3 in order to ensure that the result of $new(S, v)$ is a consistent storeless structure.

Lemma 6 *Given $\Gamma, \Gamma' \in \mathcal{S}(\Sigma)$ two structures such that $\Gamma' \models \langle T\rangle\sigma$, and $S \in \Gamma$ a set such that for all $Y \in \Gamma$, either $Sv \cap Y = \emptyset$ or $Y = Sv$, then $\Gamma \in wp(new(S, v), \langle T\rangle\sigma)$ if and only if $\Gamma \models (T = Sv \wedge \sigma \cap Sv \neq \emptyset) \vee \langle T\rangle\sigma$.*

Add The last rule defines the weakest precondition for modalities under the edge *add* operation. The complexity of the precondition formula occurs as an inherent consequence of *add*'s rather complex storeless operational semantics (12).

$$\{\exists X\ [\chi(X) = U \wedge \bigvee_{i=1,2} \psi_i(X)]\}\ \mathbf{add}(\mathbf{S}, \mathbf{v}, \mathbf{T})\ \{\langle U\rangle\sigma\} \tag{22}$$

where

$$\psi_1(X) \triangleq Sv((T^{-1}S)v)^*(T^{-1}X) \cap \sigma = \emptyset \wedge \langle X \rangle \sigma$$
$$\psi_2(X) \triangleq Sv((T^{-1}S)v)^*(T^{-1}X) \cap \sigma \neq \emptyset \wedge \langle X \rangle \Sigma^*$$

The soundness and completness proof is done in a similar way, with $wp(add(S, v, T), \varphi)$ denoting the weakest precondition with respect to *add* and the side condition from Lemma 4, added to ensure consistency of the result.

Lemma 7 *Given $\Gamma, \Gamma' \in \mathcal{S}(\Sigma)$ two structures such that $\Gamma' \models \langle U \rangle \sigma$, and $S, T \in \Gamma$ two sets such that, for all $Y \in \Gamma$, either $Sv \cap Y = \emptyset$ or $Y = T$, then $\Gamma \in wp(add(S, v, T), \langle U \rangle \sigma)$ if and only if $\Gamma \models \exists X\ [\chi(X) = U \wedge \bigvee_{i=1,2} \psi_i(X)]$.*

It is to be noticed that, in the above claims, we have implicitly used the fact that all primitive operations on the storeless heap are total functions i.e., Γ' always exists. This observation leads to the fact that the transition relation $\hookrightarrow$ defined by the rules in Figure 5 is both total and deterministic. According to the previous discussion, the weakest precondition predicate transformer $\widetilde{pre}$ distributes over all first order logical connectives. Under this assumption, we can express the precondition of an arbitrary AL formula φ recursively on the structure of φ. Let *op* be a primitive operation i.e., one of *rem*, *new* and *add* provided with some sound parameters S, v and T, then we denote by $\widetilde{pre}(op, \varphi)$ the formula obtained by recursively applying rules (20), (21) and (22) to φ. Notice that, if φ is a pure assertion, we have $\widetilde{pre}(op, \varphi) = \varphi$. With these considerations, Figure 7 shows the weakest preconditions for the three types of statements considered in this paper. In order to deal with total correctness, for statements that use dereferencing, we must add conditions to match the preconditions of the operational semantic rules in Figure 5. Differently stated, this ensures that all transitions can actually execute. For a sequence of statements $\omega \in Stmn^*$ let, $\widetilde{pre}(\omega, \varphi)$ be the precondition formula defined by the rules in Figure 7 using the classical composition rule $\{\widetilde{pre}(n, \widetilde{pre}(m, \varphi))\}\ \mathbf{m}; \mathbf{n}\ \{\varphi\}$.

$$\frac{\{\widetilde{pre}(rem(\epsilon, v), \varphi)\}}{v := null}\ (23)\qquad \frac{\{\exists S\ [\langle S \rangle \tau \wedge \widetilde{pre}(rem(S, v), \varphi)]\}}{\tau.v := null}\ (26)$$
$$\{\varphi\} \qquad\qquad \{\varphi\}$$

$$\frac{\{\widetilde{pre}(rem(\epsilon, v), \widetilde{pre}(new(\epsilon, v), \varphi))\}}{v := new}\ (24)\qquad \frac{\{\exists S\ [\langle S \rangle \tau \wedge \widetilde{pre}(rem(\epsilon, v), \widetilde{pre}(new(\epsilon, v), \varphi))]\}}{\tau.v := new}\ (27)$$
$$\{\varphi\} \qquad\qquad \{\varphi\}$$

$$\frac{\{\exists S\ [\langle S \rangle \theta \wedge \widetilde{pre}(rem(\epsilon, v), \widetilde{pre}(add(\epsilon, v, S), \varphi))]\}}{v := \theta}\ (25)\qquad \frac{\{\exists S, T\ [\langle S \rangle \tau \wedge \langle T \rangle \theta \wedge \widetilde{pre}(rem(S, v), \widetilde{pre}(add(S, v, T), \varphi))\}]}{\tau.v := \theta}\ (28)$$
$$\{\varphi\} \qquad\qquad \{\varphi\}$$

Figure 7: Weakest Preconditons for Statements

Theorem 4 *Given $\Gamma, \Gamma' \in \mathcal{S}(\Sigma)$ two structures, $\omega \in Stmn^*$ a statement and φ an AL formula such that $\Gamma' \models \varphi$, then $\Gamma \in wp(\omega, \varphi)$ if and only if $\Gamma \models \widetilde{pre}(\omega, \varphi)$.*

In conclusion, we briefly discuss the use of AL to describe the semantics of non-deterministic programs. In practice, non-determinism can be the result of parallel composition i.e., one can immagine a parallel version of the language in Figure 3, or abstraction i.e., non-deterministic choices can be introduced by loss of precision. In terms of weakest preconditions, non-determinism means that the implication $\widetilde{pre}(\omega, \varphi) \Leftarrow \neg\widetilde{pre}(\omega, \neg\varphi)$ does not hold any longer. Consequently we also lose the fact that $\widetilde{pre}(\omega, \varphi_1 \vee \varphi_2) \Rightarrow \widetilde{pre}(\omega, \varphi_1) \vee \widetilde{pre}(\omega, \varphi_2)$ and the same for existential quantifier. However, $\widetilde{pre}(\omega, \varphi_1 \vee \varphi_2) \Leftarrow \widetilde{pre}(\omega, \varphi_1) \vee \widetilde{pre}(\omega, \varphi_2)$ still holds, and similar for the existential quantifier. This results in a loss of completness of weakest preconditions. Notice that, if we still replace $\widetilde{pre}(\omega, \varphi_1 \vee \varphi_2)$ by $\widetilde{pre}(\omega, \varphi_1) \vee \widetilde{pre}(\omega, \varphi_2)$ we obtain a stronger precondition i.e., a sound but incomplete rule. In practice this might be useful still, since any result that we can infer is correct.

However, the reasoning we applied to disjunction and existential quantification cannot be applied to negation, since replacing $\widetilde{pre}(\omega, \neg\varphi)$ by $\neg\widetilde{pre}(\omega, \varphi)$ results in weakening the precondition. Instead, we

write the postcondition in positive normal form (with only atomic terms in the scope of a negation) and give a set of sound axioms for the negated modalities. This construction is further developped in the Appendix. It is easy to verify that using positive normal forms with sound axioms for negation, ensures soundness of the precondition axioms.

5.2 An Example

We assessed our calculus on a classical example in the literature [14]: the in-place list reversal program from Figure 8.

The goal is to prove that $nclist(i,n) \wedge nclist(j,n) \wedge nshared(i,j)$ is an invariant of the while loop. Since we deal with total correctness, we shall assume a side condition of the form $\exists I\ [\langle I \rangle i]$ throughout the computation. In other words, we assume that the i-list has at least one element, which allows us to iterate at least once.

```
j := null;
while i ≠ null do
        k := i.n;
        i.n := j;
        j := i;
        i := k;
od
```

Figure 8: List Reversal Program

Below we give the bottom-up derivation for the loop body. For presentation purposes we have skipped the simplifications applied to each step. We use the consequence rule i.e., if $P \Rightarrow Q$ and $\{Q\}\ \mathbf{C}\ \{R\}$ then $\{P\}\ \mathbf{C}\ \{R\}$.

$$\{nclist(i) \wedge nclist(j) \wedge nshared(i,j)\}$$
$$\Rightarrow$$
$$\{nclist(i) \wedge nclist(j) \wedge nshared(i,j) \wedge \forall I, J\ [\langle I \rangle i \wedge \langle J \rangle j \Rightarrow J^{-1} I \cap n^* = \emptyset]\}$$
$$k := i.n$$
$$\{nclist(k) \wedge nclist(j) \wedge nshared(k,j) \wedge \forall I, J\ [\langle I \rangle i \wedge \langle J \rangle j \Rightarrow J^{-1} I \cap n^* = \emptyset]\}$$
$$i.n := j$$
$$\{nclist(k) \wedge nclist(i) \wedge nshared(k,i)\}$$
$$j := i$$
$$\{nclist(k) \wedge nclist(j) \wedge nshared(k,j)\}$$
$$i := k$$
$$\{nclist(i) \wedge nclist(j) \wedge nshared(i,j)\}$$

The first two steps can be resumed by substitution, since we are assigning between local variables. In the third step we obtain an extra precondition saying that there should be no n-paths between the node containing j and the node containing i, or else the assignment $i.n := j$ would generate a cycle, invalidating $nclist(i)$. This precondition is however implied by $nshared(i,j)$ and disapears when we apply the consequence rule after the last step.

6 Conclusions and Future Work

This paper presents an assertion language and an associated program logic to reason about linked data structures. The presented program logic is shown to be sound and complete for program instructions that destructively update such data structures. The program logic is based on a storeless semantics that is shown to be equivalent to a low level store-based semantics.

The expressive power as well as the decidability of the assertion language AL or of some relevant fragments need to be studied. An other promising direction is the development of an abstraction interpretation based approach using AL asertions as abtract domain in the spirit of shape analysis [13].

References

[1] Marius Bozga, Radu Iosif and Yassine Lakhnech: Storeless Semantics and Alias Logic. Technical report, http://www-verimag.imag.fr/~bozga/ssal.ps

[2] Cristiano Calcagno, Peter W. O'Hearn and Richard Bornat: Program Logic and Equivalence in the Presence of Garbage Collection. Proc. FOSSACS 2001.

[3] Frank S. de boer. Reasoning about dynamically evolving process structures. Ph.D. Thesis, Vrije Universiteit te Amsterdam (University of Amsterdam), 1991.

[4] Alain Deutsch. A storeless model of aliasing and its abstractions using finite representations of right-regular equivalence relations. In Proceedings of the IEEE 1992 Conference on Computer Languages (April 1992) pp. 2-13

[5] Alain Deutsch. Interprocedural May-Alias Analysis for Pointers: Beyond k-limiting. In Proceedings of the ACM SIGPLAN'94 Conference on Programming Language Design and Implementation (PLDI), Orlando, Florida, June 20-24, 1994. SIGPLAN Notices 29(6) (June 1994) pp. 230-241

[6] Edsger W. Dijkstra and Carel S. Scholten. Predicate Calculus and Program Semantics. Springer-Verlag, New York 1990.

[7] Samuel Eilenberg. Automata, Languages, and Machines. Academic Press (1976)

[8] Radu Iosif: Symmetry Reductions for Model Checking of Concurrent Dynamic Software. Submitted, http://www-verimag.imag.fr/~iosif/papers.html

[9] Samin Ishtiaq and Peter O'Hearn: BI as an Assertion Language for Mutable Data Structures. Proc. of 28th ACM-SIGPLAN Symposium on Principles of Programming Languages (2001)

[10] C.A.R Hoare and He Jifeng: A Trace Model for Pointers and Objects. In Proc. ECOOP'99, Lecture Notes in Computer Science, Vol. 1628, pp. 1-18 (1999)

[11] H. B. M. Jonkers. Abstract Storage Structures. Algorithmic Languages, North-Holland (1981)

[12] M. Hennessy and R. Milner: Algebraic Laws for Nondeterminism and Concurrency. Journal of the ACM Vol. 32 (1985) pp. 137-161

[13] M. Sagiv, M., T. Reps and R. Wilhelm. Parametric Shape Analysis via 3-Valued Logic. Conference Record of the Twenty-Third ACM Symposium on Principles of Programming Languages, (San Antonio, TX, Jan. 20-22, 1999), ACM, New York, NY, 1999.

[14] John C. Reynolds. Separation Logic: A Logic for Shared Mutable Data Structures. Proc 17th IEEE Symposium on Logic in Computer Science (2002)

APPENDIX

The following axioms express sound preconditions for the negated modalities. This is useful for developing a sound (but not complete) semantics of non-deterministic programs. Notice first that $\neg\langle U\rangle\sigma = \neg in(U) \vee \sigma \cap V = \emptyset$. Since $\sigma \cap U = \emptyset$ is a pure assertion, it is sufficient to define preconditions only for $\neg in(U)$. Note that the axioms below can be easily shown to be also complete.

Negate Remove Intuitively, a set U will not belong to a structure after a remove has been performed iff by performing the remove operation from an existing set we will not obtain U.

$$\{\forall X\ [in(X) \Rightarrow X \setminus Sv\Sigma^* \neq U]\}\ \mathbf{rem}(\mathbf{S}, \mathbf{v})\ \{\neg in(U)\}$$

Negate New A set U will not belong to a structure after a new operation has been performed iff it does not belong to the structure before and it is not equal to the set that will be added by *new*.

$$\{\neg in(U) \wedge U \neq Sv\}\ \mathbf{new}(\mathbf{S}, \mathbf{v})\ \{\neg in(U)\}$$

Negate Add A set U will not belong to a structure after an add operation has been performed iff the pre-image of its transformation does not belong to the original structure.

$$\{\forall X\ [in(X) \Rightarrow \chi(X) \neq U]\}\ \mathbf{add}(\mathbf{S}, \mathbf{v}, \mathbf{T})\ \{\neg in(U)\}$$

By adding the $\sigma \cap U = \emptyset$ disjunct to the above preconditions, we will obtain the weakest precondition of $\neg\langle U\rangle\sigma$.