

Expression-based aliasing for OO-languages

Georgiana Caltais¹

Department of Computer Science, ETH Zürich, Switzerland

Abstract. Alias analysis has been an interesting research topic in verification and optimization of programs. The undecidability of determining whether two expressions in a program may reference to the same object is the main source of the challenges raised in alias analysis. In this paper we propose an extension of a previously introduced alias calculus based on program expressions, to the setting of unbounded program executions such as infinite loops and recursive calls. Moreover, we devise a corresponding executable specification in the K-framework. An important property of our extension is that, in a non-concurrent setting, the corresponding alias expressions can be over-approximated in terms of a notion of regular expressions. This further enables us to show that the associated K-machinery implements an algorithm that always stops and provides a sound over-approximation of the “may aliasing” information, where soundness stands for the lack of false negatives. As a case study, we analyze the integration and further applications of the alias calculus in SCOOP. The latter is an object-oriented programming model for concurrency, recently formalized in Maude; K definitions can be compiled into Maude for execution.

1 Introduction

A research direction of interest in Computer Science is the application of *alias analysis* in verification and optimization of programs. One of the challenges along this line of research has been the undecidability of determining whether two expressions in a program *may* reference the same object. A rich suite of approaches aiming at providing a satisfactory balance between scalability and precision has already been developed in this regard. Examples include: (i) intra-procedural frameworks [17, 16] that handle isolated functions only, and their inter-procedural counterparts [16, 23, 12] that consider the interactions between function calls; (ii) type-based techniques [9]; (iii) flow-based techniques [4, 7] that establish aliases depending on the control-flow information of a procedure; (iv) context-(in)sensitive approaches [10, 30] that depend on whether the calling context of a function is taken into account or not; (v) field-(in)sensitive approaches [21, 1] that depend on whether the individual fields of objects in a program are traced or not. More details on such classifications can be found in [26], for instance. For a comprehensive survey on alias analyses for object-oriented programs, corresponding issues and remaining open problems, we refer the interested reader to the works in [29, 11].

Of particular interest for the work in this paper is the untyped, flow-sensitive, field sensitive, inter-procedural and context-sensitive calculus for *may aliasing*, introduced in [15]. The aforementioned calculus covers most of the aspects of a modern object-oriented language, namely: object creation and deletion, conditionals, assignments, loops and (possibly recursive) function calls. The approach in [15] abstracts the aliasing information in terms of explicit access paths [18] referred to as *alias expressions*. Consider, for an example, the code

$$\begin{array}{l} x := y; \\ \mathbf{loop} \ x := x.next \ \mathbf{end} \end{array} \quad (1)$$

The corresponding execution causes x to become aliased to $y.next.next. \dots$, with a possibly infinite number of occurrences of the field $next$. The set of associated alias expressions can be equivalently written as:

$$\{[x, y.next^k] \mid k \geq 0\}. \quad (2)$$

The sources of imprecision introduced by the calculus in [15] are limited to ignoring tests in conditionals, and to “cutting at length L ”, for the case of possibly infinite alias relation as in (2). Intuitively, the cutting technique considers sequences longer than a given length L as aliased to all expressions.

There is a huge literature on heap analysis for aliasing [11], but hardly any paper that presents a calculus as in [15] allowing the derivation of alias relations as the result of applying various instructions of a programming language.

Our focus is two folded. First, we want extend the framework in [15] to the setting of unbounded program executions such as infinite loops and recursive calls. In accordance, the goal is to provide a way to shift from “finite” to “infinite behaviours”. This can be achieved in a rather straightforward manner, by redefining the construct **loop** p **end** in [15] according to the informal semantics: “execute p repeatedly any number of times, including zero”. However, developing a corresponding mechanism for reasoning on “may aliasing” in a finite number of steps is not trivial. The key observation that paves the way to a possible (finite state-based) modeling in a non-concurrent setting is that the alias expressions corresponding to loops and recursive calls grow in a regular fashion. Hence, they are finitely representable, as it is easy to see in (2), for instance. Such regularities cannot be exploited in concurrent contexts, due to the “non-determinism” of process interaction.

A similar technique exploiting regular behaviour of (non-concurrent) programs, in order to reason on “may aliasing”, was previously introduced in [2]. In short, the results in [2] utilize abstract representations of programs in terms of finite pushdown systems, for which infinite execution paths have a regular structure (or are “lasso shaped”) [3]. Then, in the style of abstract interpretation [8], the collecting semantics is applied over the (finite state) pushdown systems to obtain the alias analysis itself. In short, the main difference with the results in [2] consists in how the abstract memory addresses corresponding to pointer variables are represented. In [2] these range over a finite set of natural numbers. In this paper we consider alias expressions build according to the calculus in [15].

The work in [2] also proposes an implementation of pushdown systems in the \mathbb{K} -framework [27]. The latter is an executable semantic framework based on Rewriting Logic (RL) [19], and has successfully been used for defining programming languages and corresponding formal analysis tools. Moreover, \mathbb{K} definitions have a direct implementation in \mathbb{K} -Maude [28].

We agree that it could be worth presenting our analysis as an abstract interpretation (AI) [8]. A modelling exploiting the machinery of AI (based on abstract domains, abstraction and concretization functions, Galois connections, fixed-points, etc.) is an interesting, but different research topic per se.

Our second interest w.r.t. may aliasing is its integration in SCOOP [22] – a simple object oriented programming model for concurrency; thus an operational based approach on handling the alias calculus is more appropriate. The basis of a RL-based framework for the design and analysis of the SCOOP model was recently set in [22]. The reference implementation of SCOOP is Eiffel [20]. The integration of alias analysis belongs to a more ambitious goal, namely, the construction of a RL-based toolbox for the analysis of SCOOP programs (examples include a deadlock detector and a type checker).

Our contribution. By drawing inspiration from, and building on top of the results in [15, 2], in this paper we propose:

- an extension of the (finite) alias calculus in [15] to the setting of unbounded program executions, and a sound over-approximation technique based on “regular alias expressions”, for non-concurrent settings;
- a RL-based specification of the extended calculus;
- an algorithm that always terminates and provides a sound over-approximation of “may aliasing” by exploiting a notion of regular (finitely representable) aliases, for non-concurrent settings.

Moreover, we analyze the integration, implementation and further applications of the alias calculus in SCOOP.

We refer the interested reader to [5] for the extended version of the current paper including: the full specification of the RL-based machinery, two examples emphasizing the naturalness of applying the executable aliasing framework and a case study exploiting the corresponding implementation in SCOOP, respectively, together with the detailed proofs of the formal results.

Paper structure. The paper is organized as follows. In Section 2 we introduce the extension of the alias calculus in [15] to unbounded executions. In Section 3 we provide the RL-based executable specification of the calculus in the \mathbb{K} semantic framework. The implementation in SCOOP, and further applications are discussed in Section 4. In Section 5 we draw the conclusions and provide pointers to future work.

2 The alias calculus

In this section we define an extension of the calculus in [15], to unbounded program executions. Moreover, based on the idea behind the *pumping lemma* for

regular languages [25], we devise a corresponding sound over-approximation of “may aliasing” in terms of regular expressions, applicable in sequential contexts. This paves the way to developing an algorithm for the aliasing problem, as presented in Section 3, in the formal setting of the \mathbb{K} semantic framework [27].

Preliminaries. We proceed by briefly recalling the notion of *alias relation* and a series of associated notations and basic operations, as introduced in [15].

We call an *expression* a (possibly infinite) path of shape $x.y.z. \dots$, where x is a local variable, class attribute or *Current*, and y, z, \dots are attributes. Here, *Current*, also known as *this* or *self*, stands for the current object. For an arbitrary alias expression e , it holds that $e.Current = Current.e = e$. Let E represent the set of all expressions of a program. An *alias relation* is a symmetric and irreflexive binary relation over $E \times E$.

Given an alias relation r and an expression e , we define

$$r/e = \{e\} \cup \{x: E \mid [x, e] \in r\}$$

denoting the set consisting of all elements in r which are aliased to e , plus e itself.

Let x be an expression; we write $r - x$ to represent r without the pairs with one element of shape $x.e$.

We say that an alias relation is *dot complete* whenever for any t, u, v and a it holds that if $[t, u]$ and $[t.a, v]$ are alias pairs, then $[u.a, v]$ is an alias pair and, moreover, if a is in the domain of t , then $[t.a, u.a]$ is an alias pair. By the “domain of t ” we refer to a method or a field in the class corresponding to the object referred by the expression associated to t . For instance, given a class `NODE` with a field `next` of type `NODE`, and a `NODE` object x , we say that `next` is in the domain of $t = x.next.next$. For the sake of brevity, we write *dot-complete*(r) for the closure under dot-completeness of a relation r .

The notation $r[x = u]$ represents the relation r augmented with pairs $[x, y]$ and made dot complete, where y is an element of u .

2.1 Extension to unbounded executions

We further introduce an extension of the alias calculus in [15] to infinite alias relations corresponding to unbounded executions such as infinite loops or recursive calls. The main difference in our approach is reflected by the definition of loops, which now complies to the usual fixed-point denotational semantics.

The alias calculus is defined by a set of axioms “describing” how the execution of a program affects the aliasing between expressions. As in [15], the calculus ignores tests in conditionals and loops. The *program instructions* are defined as follows:

$$\begin{aligned}
 p ::= & p; p \mid \mathbf{then} \ p \ \mathbf{else} \ p \ \mathbf{end} \mid \\
 & \mathbf{create} \ x \mid \mathbf{forget} \ x \mid t := s \mid \\
 & \mathbf{loop} \ p \ \mathbf{end} \mid \mathbf{call} \ f(l) \mid x.\mathbf{call} \ f(l).
 \end{aligned}
 \tag{3}$$

In short, we write $r \gg p$ to represent the alias information obtained by executing p when starting with the initial alias relation r .

The axiom for sequential composition is defined in the obvious way:

$$r \gg (p; q) = (r \gg p) \gg q. \quad (4)$$

Conditionals are handled by considering the union of the alias pairs resulted from the execution of the instructions corresponding to each of the two branches, when starting with the same initial relation:

$$r \gg (\mathbf{then} \ p \ \mathbf{else} \ q \ \mathbf{end}) = r \gg p \cup r \gg q. \quad (5)$$

As previously mentioned, we define $r \gg \mathbf{loop} \ p \ \mathbf{end}$ according to its informal semantics : “execute p repeatedly any number of times, including zero”. The corresponding rule is:

$$r \gg (\mathbf{loop} \ p \ \mathbf{end}) = \bigcup_{n \in \mathbb{N}} (r \gg p^n) \quad (6)$$

where \cup stands for the union of alias relations, as above. This way, our calculus is extended to infinite alias relations. This is the main difference with the approach in [15] that proposes a “cutting” technique restricting the model to a maximum length L . In [15], sequences longer than L are considered as aliased to all expressions. Orthogonally, for sequential settings, we provide finite representations of infinite alias relations based on over-approximating regular expressions, as we shall see in Section 2.2.

Both the creation and the deletion of an object x eliminate from the current alias relation all the pairs having one element prefixed by x :

$$\begin{aligned} r \gg (\mathbf{create} \ x) &= r - x \\ r \gg (\mathbf{forget} \ x) &= r - x. \end{aligned} \quad (7)$$

The (qualified) function calls comply to their initial definitions in [15]:

$$\begin{aligned} r \gg (\mathbf{call} \ f(l)) &= (r[f^\bullet:l]) \gg |f| \\ r \gg (x.\mathbf{call} \ f(l)) &= x.((x'.r) \gg \mathbf{call} \ f(x'.l)). \end{aligned} \quad (8)$$

Here f^\bullet and $|f|$ stand for the formal argument list and the body of f , respectively, whereas $r[u:v]$ is the relation r in which every element of the list v is replaced by its counterpart in u . Intuitively, the negative variable x' is meant to transpose the context of the qualified call to the context of the caller. Note that “.” (i.e., the constructor for alias expressions) is generalized to distribute over lists and relations: $x.[a, b, \dots] = [x.a, x.b, \dots]$.

For an example, consider a class C in an OO-language, and an associated procedure f that assigns a local variable y , defined as: $f(x) \{ y := x \}$. Then, for instance, the aliasing for $a.\mathbf{call} \ f(a)$ computes as follows:

$$\begin{aligned} \emptyset \gg a.\mathbf{call} \ f(a) &= \\ a.(a'.\emptyset \gg y := a'.a) &= \\ a.(\emptyset \gg y := \mathit{Current}) &= \\ \mathit{dot-complete}(\{[a.y, a]\}) & \end{aligned}$$

Recursive function calls can lead to infinite alias relations. In sequential settings, as for the case of loops, the mechanism exploiting sound regular over-approximations in order to derive finite representations of such relations is presented in the subsequent sections.

The axiom for assignment is as well in accordance with its original counterpart in [15]:

$$r \gg (t := s) = \mathbf{given} \ r_1 = r[ot = t] \quad \mathbf{then} \ (r_1 - t)[t = (r_1/s - t)] - ot \ \mathbf{end} \quad (9)$$

where ot is a fresh variable (that stands for “old t ”). Intuitively, the aliasing information w.r.t. the initial value of t is “saved” by associating t and ot in r and closing the new relation under dot-completeness, in r_1 . Then, the initial t is “forgotten” by computing $r_1 - t$ and the new aliasing information is added in a consistent way. Namely, we add all pairs (t, s') , where s' ranges over $r_1/s - t$ representing all expressions already aliased with s in r_1 , including s itself, but without t . Recall that alias relations are not reflexive, thus by eliminating t we make sure we do not include pairs of shape $[t, t]$. Then, we consider again the closure under dot-completeness and forget the aliasing information w.r.t. the initial value of t , by removing ot .

Remark 1. It is worth discussing the reason behind *not* considering transitive alias relations. Assume the following program:

then $x := y$ **else** $y := z$ **end**

Based on the equations (5) and (9) handling conditionals and assignments, respectively, the calculus correctly identifies the alias set: $\{[x, y], [y, z]\}$. Including $[x, z]$ would be semantically equivalent to the execution of the two branches in the conditional at the same time, which is not what we want.

2.2 A sound over-approximation

In a sequential setting, the challenge of computing the alias information in the context of (infinite) loops and recursive calls reduces to evaluating their corresponding “unfoldings”, captured by expressions of shape

$$r \gg p^\omega,$$

with ω ranging over naturals plus infinity, r an (initial) alias relation ($r = \emptyset$), and p a *basic control block* defined by:

$$\begin{aligned} p ::= & p; p \mid \mathbf{then} \ p \ \mathbf{else} \ p \ \mathbf{end} \mid \\ & \mathbf{create} \ x \mid \mathbf{forget} \ x \mid \\ & t := s. \end{aligned} \quad (10)$$

The value $r \gg p^\omega$ refers to the alias relation obtained by recursively executing the control block p , and it is calculated in the expected way:

$$\begin{aligned} r \gg p^0 &= r \\ r \gg p^{k+1} &= (r \gg p^k) \gg p. \end{aligned}$$

Consider again the code in (1):

```

 $x := y;$ 
loop  $x := x.next$  end.

```

Its execution generates the alias relation

$$((\emptyset \gg (x := y)) \gg (x := x.next)) \gg (x := x.next) \dots$$

including an infinite number of pairs of shape:

$$[x, y.next], [x, y.next.next], [x, y.next.next.next] \dots \quad (11)$$

A similar reasoning does not hold for concurrent applications, where process interaction is not “regular”.

In what follows we provide a way to compute finite representations of infinite alias relations in sequential settings. The key observation is that alias expressions corresponding to unbounded program executions grow in a regular fashion. See, for instance, the aliases in (11), which are pairs of type $[x, y.next^{k \geq 1}]$.

Regular expressions are defined similarly to the regular languages over an alphabet. We say that an expression is *regular* if it is a local variable, class attribute or *Current*. Moreover, the concatenation $e_1 . e_2$ of two regular expressions e_1 and e_2 is also regular. Given a regular alias expression e , the expression e^* is also regular; here $(-)^*$ denotes the Kleene star [14]. We call an alias relation *regular* if it consists of pairs of regular expressions.

Lemma 1. *Assume p a program built according to the rules in (3). Then, in a sequential setting, the relation $\emptyset \gg p$ is regular.*

Proof. The result follows by induction on the structure of p .

Inspired by the idea behind the *pumping lemma for regular languages* [25], we define a *lasso* property for alias relations, which identifies the repetitive patterns within the structure of the corresponding alias expressions. The intuition is that such patterns will occur for an infinite number of times due to the execution of loops or recursive function calls. Then, we supply sound over-approximations of “lasso” relations, based on regular alias expressions.

In the context of alias relations, we say that the lasso property is satisfied by r and r' whenever the following two conditions hold: (1) r behaves like a *lasso base* of r' . Namely, all the pairs $[e_1, e_2] \in r$ are used to generate elements $[e'_1, e'_2] \in r'$, by repeating tails of prefixes of e_1 and e_2 , respectively, and (2) r' is a *lasso extension* of r . Namely, all the pairs in r' are generated from elements of r by repeating tails of their prefixes. For example, if e_1 above is an expression of shape $x.y.z.w$, then e'_1 can be $x.y.y.z.w$ if we consider the tail y of the prefix $x.y$, or $x.y.z.y.z.w$ if we take the tail $y.z$ of the prefix $x.y.z$.

Formally, consider r and r' two alias relations, and x_i, y_i and z_i a set of (possibly empty) expressions, for $i \in \{1, 2\}$. Then:

$$\text{lasso}(r, r') = ([x_1 y_1 z_1, x_2 y_2 z_2] \in r \text{ iff } [x_1 y_1 y_1 z_1, x_2 y_2 y_2 z_2] \in r'). \quad (12)$$

For the simplicity of notation we sometimes omit the dot-separators between expressions. For instance, we write $x y z$ in lieu of $x.y.z$.

Assuming a lasso over r and r' , we compute a relation consisting of regular expressions over-approximating r and r' as:

$$\text{reg}(r, r') = \{ [x_1 y_1^* z_1, x_2 y_2^* z_2] \mid \begin{array}{l} [x_1 y_1 z_1, x_2 y_2 z_2] \in r \wedge \\ [x_1 y_1 y_1 z_1, x_2 y_2 y_2 z_2] \in r' \end{array} \} \quad (13)$$

where x_i, y_i and z_i are possibly empty expressions, for $i \in \{1, 2\}$. As previously indicated, the over-approximation is sound w.r.t. the repeated application of a basic control block as in (10), in the way that it does not introduce any false negatives:

Lemma 2. *Consider r and r' two alias relations, and p a basic control block in a sequential setting. If $r \gg p = r'$ and $\text{lasso}(r, r') = \text{true}$, then the following holds for all $n \geq 1$:*

$$r \gg p^n \in \text{reg}(r, r').$$

Proof. The reasoning is by induction on n . The base case follows immediately, whereas the induction step is proved by “reductio ad absurdum”. \square

3 A \mathbb{K} -machinery for collecting aliases

In this section we provide the specification of a RL-based mechanism collecting the alias information in the \mathbb{K} semantic framework [27]. We choose \mathbb{K} more as a notational convention to enable compact and modular definitions. In reality, the \mathbb{K} -rules in this section are implemented in Maude, as rewriting theories, on top of the formalization of SCOOP [22] (we refer to Section 4 for more details on our approach).

In short, our strategy is to start with a program built on top of the control structures in (3), then to apply the corresponding \mathbb{K} -rules in order to get the “may aliasing” information in a designated \mathbb{K} -cell ($\langle - \rangle_{\text{al}}$). Independently of the setting (sequential or concurrent) one can exploit this approach in order to evaluate the aliases of a given finite length L . We also show that for sequential contexts, the application of the \mathbb{K} -rules is finite and the aliases in the final configuration soundly over-approximate the (infinite) “may alias” relations of the calculus.

Brief overview of \mathbb{K} . \mathbb{K} [27] is an executable semantic framework based on Rewriting Logic [19]. It is suitable for defining (concurrent) languages and corresponding formal analysis tools, with straightforward implementation in \mathbb{K} -Maude [28]. \mathbb{K} -definitions make use of the so-called *cells*, which are labelled and can be nested, and (rewriting) *rules* describing the intended (operational) semantics.

A *cell* is denoted by $\langle - \rangle_{[\text{name}]}$, where [name] stands for the *name of the cell*. A construction $\langle . \rangle_n$ stands for an *empty cell* named n . We use “pattern

matching” and write $\langle c \dots \rangle_n$ for a cell with content c at the top, followed by an arbitrary content (\dots) . Orthogonally, we can utilize cells of shape $\langle \dots c \rangle_n$ and $\langle \dots c \dots \rangle_n$, defined in the obvious way.

Of particular interest is $\langle - \rangle_k$ – the *continuation cell*, or the *k-cell*, holding the stack of program instructions (associated to one processor), in the context of a programming language formalization. We write

$$\langle i_1 \curvearrowright i_2 \dots \rangle_k$$

for a set of instructions to be “executed”, starting with instruction i_1 , followed by i_2 . The associative operation \curvearrowright is the instruction sequencing.

A \mathbb{K} -rewrite rule

$$\langle c \dots \rangle_{n_1} \langle c' \rangle_{n_2} \Rightarrow \langle c' \dots \rangle_{n_1} \langle \dots c' \rangle_{n_3} \quad (14)$$

reads as: if cell n_1 has c at the top and cell n_2 contains value c' , then c is replaced by c' in n_1 and c' is added at the end of the cell n_3 . The content of n_2 remains unchanged. In short, (14) is written in a \mathbb{K} -like syntax as:

$$\langle \frac{c}{c'} \dots \rangle_{n_1} \langle c' \rangle_{n_2} \langle \dots \frac{\cdot}{c'} \rangle_{n_3}.$$

We further provide the details behind the \mathbb{K} -specification of the alias calculus. As expected, the *k-cell* retains the instruction stack of the object-oriented program. We utilize cells $\langle - \rangle_{\text{al}}$ to enclose the current alias information, and the so-called *back-tracking cells* $\langle - \rangle_{\text{bkt} \dots}$ enabling the sound computation of aliases for the case of **then – else – end** and, in non-concurrent contexts, for loops and (possibly recursive) function calls. As a convention, we mark with (\clubsuit) the rules that are sound only for non-concurrent applications, based on Lemma 2. Due to space limitations, in what follows we introduce only the \mathbb{K} -rules for handling assignments and loops. The entire specification can be found in [5].

As expected, the assignment rule simply restores the current alias relation according to its axiom in (9), and removes the assignment instruction from the top of the *k-cell*:

$$\langle \frac{r}{(r_1 - t)[t = (r_1/s - t)] - ot} \rangle_{\text{al}} \langle \frac{t := s \dots}{\cdot} \rangle_{\text{k}} \quad \text{with } r_1 = r[ot = t] \quad (15)$$

For **loop p end**, we utilize a meta-construction $p \boxed{1} \text{loop } p \text{ end}$ simulating the unfolding corresponding to (6), and a back-tracking stack $\langle - \rangle_{\text{bkt-1}}$ collecting the alias information obtained after each execution of p . Moreover, the \mathbb{K} -implementation exploits the result in Lemma 2. Whenever a “lasso” is reached, the infinite rewriting is prevented by resuming the infinite application of p in terms of a sound over-approximating alias relation. The \mathbb{K} -rules are as follows.

First, the aforementioned unfolding is performed, and the alias relation before p is stored in the back-tracking cell as $\langle r \rangle_{\text{al-o}} \langle p \rangle_1$:

$$\langle r \rangle_{\text{al}} \langle \frac{\text{loop } p \text{ end}}{p \boxed{1} \text{loop } p \text{ end}} \dots \rangle_{\text{k}} \langle \frac{\cdot}{\langle r \rangle_{\text{al-o}} \langle p \rangle_1} \dots \rangle_{\text{bkt-1}} \quad (16)$$

If the alias relation r' obtained after the successful execution of p (marked by $\boxed{1}$ at the top of the continuation) is not a lasso of the aliasing r before p (previously stored in $\langle - \rangle_{\text{bkt-1}}$) then p is constrained to a new execution by becoming the top of the k -cell, and r' is memorized for back-tracking:

$$\langle r' \rangle_{\text{al}} \langle \frac{\boxed{1}}{p} \mathbf{loop} \ p \ \mathbf{end} \ \dots \rangle_k \langle \frac{\langle r \rangle_{\text{al-o}} \langle p \rangle_1}{\langle r' \rangle_{\text{al-o}} \langle p \rangle_1} \dots \rangle_{\text{bkt-1}} \text{ if not } \text{lasso}(r, r') \quad (\clubsuit) \quad (17)$$

Last, if a lasso is reached after the execution of p , then the current aliasing is soundly replaced by a “regular” over-approximation $\text{reg}(r, r')$, the corresponding back-tracking information is removed from $\langle - \rangle_{\text{bkt-1}}$ and the **loop** instruction is eliminated from the k -cell:

$$\langle \frac{r'}{\text{reg}(r, r')} \rangle_{\text{al}} \langle \boxed{1} \mathbf{loop} \ p \ \mathbf{end} \ \dots \rangle_k \langle \frac{\langle r \rangle_{\text{al-o}} \langle p \rangle_1}{\dots} \dots \rangle_{\text{bkt-1}} \text{ if } \text{lasso}(r, r') \quad (\clubsuit) \quad (18)$$

In a non-concurrent setting, the machinery orchestrating the K-rules introduced in this section implements an algorithm that always terminates and provides a sound over-approximation of “may aliasing”.

Theorem 1. *Consider p a program built on top of the control structures in (3), that executes in a sequential setting. Then, the application of the corresponding K-rules when starting with p and an empty alias relation, is a finite rewriting of shape*

$$\langle \emptyset \rangle_{\text{al}} \langle p \rangle_k \xrightarrow{(*)} \langle r \rangle_{\text{al}} \langle \cdot \rangle_k,$$

with r a sound over-approximation of the aliasing information corresponding to the execution of p .

Proof. The key observation is that, due to the execution of loops and/or recursive calls, expressions can infinitely grow in a *regular* fashion. Hence, a lasso is always reached. Consequently, the control structure generating the infinite behaviour is removed from the k -cell, according to the associated K-specification for loops and/or recursive calls. This guarantees termination. Moreover, recall that the regular expressions replacing the current alias information are a sound over-approximation, according to Lemma 2. \square

Observe that the *RL*-based machinery can simulate precisely the “cutting at length L ” technique in [15]. It suffices to disable the rules (\clubsuit) and stop the rewriting after L steps.

4 Integration in SCOOP

In this section we provide a brief overview on the integration and applicability of the alias calculus in SCOOP [22] – a simple object-oriented programming model for concurrency. Two main characteristics make SCOOP simple: 1) just

one keyword programmers have to learn and use in order to enable concurrent executions, namely, *separate* and 2) the burden of orchestrating concurrent executions is handled within the model, therefore reducing the risk of correctness issues.

In short, the key idea of SCOOP is to associate to each object a processor, or *handler* (that can be a CPU, or it can also be implemented in software, as a process or thread). Assume a processor p that performs a call $o.f()$ on an object o . If o is declared as “separate”, then p sends a request for executing $f()$ to q – the handler of o (note that p and q can coincide). Meanwhile, p can continue. Processors communicate via *channels*.

The Maude semantics of SCOOP in [22] is defined over tuples of shape

$$\langle p_1 :: St_1 \mid \dots \mid p_n :: St_n, \sigma \rangle$$

where, p_i denotes a processor (for $i \in \{1, \dots, n\}$), St_i is the call stack of p_i and σ is the *state* of the system. States hold the information about the *heap* (which is a mapping of references to objects) and the *store* (which includes formal arguments, local variables, *etc.*).

The assignment instruction, for instance, is formally specified as the transition rule:

$$\frac{\text{a is fresh}}{\Gamma \vdash \langle p :: t := s; St, \sigma \rangle \rightarrow \langle p :: \text{eval}(a, s); \text{wait}(a); \text{write}(t, a.data); St, \sigma \rangle} \quad (19)$$

where, intuitively, “ $\text{eval}(a, s)$ ” evaluates s and puts the result on channel a , “ $\text{wait}(a)$ ” enables processor p to use the evaluation result, “ $\text{write}(t, a.data)$ ” sets the value of t to $a.data$, St is a call stack, and Γ is a typing environment [24] containing the class hierarchy of a program and all the type definitions.

At this point it is easy to understand that the \mathbb{K} -rule for assignments

$$\langle \frac{r}{(r_1 - t)[t = (r_1/s - t)] - ot} \rangle_{\text{al}} \langle t := s \dots \rangle_{\text{k}} \quad \text{with } r_1 = r[ot = t] \quad (15)$$

can be straightforwardly integrated in (19) by enriching the state structure with a new field encapsulating the alias information, and considering instead the transition $\Gamma \vdash \langle p :: t := s; St, \sigma \rangle \rightarrow \langle p :: \text{eval}(a, s); \text{wait}(a); \text{write}(t, a.data); St, \sigma' \rangle$ where

$$\sigma.\text{aliases} = r \quad \sigma'.\text{aliases} = (r_1 - t)[t = (r_1/s - t)] - ot$$

with r and r_1 as in (15). The integration of all the \mathbb{K} -rules of the alias calculus on top of the Maude formalization of SCOOP can be achieved by following a similar approach.

For a case study, one can download the SCOOP formalization at:

<https://dl.dropboxusercontent.com/u/1356725/SCOOP.zip>

and run the command

```
> maude SCOOP.maude ..\examples\aliasing-linked_list.maude
```

corresponding to the code in (1):

```
x := y; loop x := x.next end.
```

The console outputs the aliased expressions for a rewriting of depth 100 which include, as expected, pairs of shape $[x, y.next^k]$. (The over-approximating mechanism for sequential settings is still to be implemented.)

As can be observed based on the code in `aliasing-linked_list.maude`, in order to implement our applications in Maude, we use intermediate (still intuitive) representations. For instance, the class structure defining a node in a simple linked list, with field `next` is declared as:

```
class 'NODE
  create {'make}
  ( attribute { 'ANY } 'next : [?, . , 'NODE] ; )
  [...]
end ;
```

where `'next : [?, . , 'NODE]` stands for an object of type `NODE`, that is handled by the current processor (`.`) and that can be `Void (?)`, and `'make` plays the role of a constructor. The intermediate representation of the instruction block in (1) is:

```
assign ('x, 'y);
until False loop ( assign ('x, 'x . 'next(nil)) ; ) end ;
```

For a detailed description of SCOOP and its Maude formalization we refer the interested reader to the work in [22].

4.1 Further applications of the alias calculus

Apart from providing an alias analysis tool, the alias calculus can be exploited in order to build an abstract semantics of SCOOP. For example, an abstraction of the assignment rule (15) would omit the evaluation of the right-hand side of the assignment $t := s$ and the associated message passing between channels:

$$\frac{}{\Gamma \vdash \langle p :: t := s; St, \sigma \rangle \rightarrow \langle p :: St, \sigma' \rangle}$$

where

$$\sigma.aliases = r \quad \sigma'.aliases = (r_1 - t)[t = (r_1/s - t)] - ot$$

with r and r_1 as in (15). This way one derives a simplified, reduced semantics of SCOOP, more appropriate for model checking, for instance; the current SCOOP formalization in Maude is often too large for this purpose. A survey on abstracting techniques on top of Maude executable semantics is provided in [19].

Furthermore, the aliasing information could be used for the so-called “deadlocking” problem, where two or more executing threads are each waiting for the other to finish. In the context of SCOOP, this is equivalent to identifying whether a set of processors reserve each other circularly (*i.e.*, there is a Coffman deadlock). This situation might occur, for instance, in a Dining Philosophers scenario, where both philosophers and forks are objects residing on their own processors. The difficulty of identifying such deadlocks stems from the fact that SCOOP processors are known from object references, which *may be aliased*.

5 Conclusions

In this paper we provide an extension of the alias calculus in [15] from finite alias relations to infinite ones corresponding to loops and recursive calls. Moreover, we devise an associated executable specification in the \mathbb{K} semantic framework [27]. In Theorem 1 we show that the RL-based machinery implements an algorithm that always terminates with a sound over-approximation of “may aliasing”, in non-concurrent settings. This is achieved based on the sound (finitely representable) over-approximation of (“lasso shaped”) alias expressions in terms of regular expressions, as in Lemma 2. We also discuss the integration and applicability of the alias calculus on top of the Maude formalization of SCOOP [22].

An immediate direction for future work is to identify interesting (industrial) case studies to be analyzed using the framework developed in this paper. We are also interested in devising heuristics comparing the efficiency and the precision (*e.g.*, the number of false positives introduced by the alias approximations) between our approach and other aliasing techniques. In this respect, we anticipate that the rewriting modulo associativity, together with the pattern matching capabilities of Maude will accelerate the identification of the “lasso” properties and the corresponding over-approximating regular alias expressions. This could eventually provide an effective reasoning apparatus for the “may aliasing” problem.

Another research direction is to derive alias-based abstractions for analyzing concurrent programs. We foresee possible connections with the work in [13] on *concurrent Kleene algebra* formalizing choice, iteration, sequential and concurrent composition of programs. The corresponding definitions exploit abstractions of programs in terms of traces of events that can depend on each other. Thus, obvious challenges in this respect include: (i) defining notions of dependence for all the program constructs in this paper, (ii) relating the concurrent Kleene operators to the semantics of the SCOOP concurrency model and (iii) checking whether fixed-points approximating the aliasing information can be identified via fixed-point theorems.

Furthermore, it would be worth investigating whether the graph-based model of alias relations introduced in [15] can be exploited in order to derive finite \mathbb{K} specifications of the extended alias calculus. In case of a positive answer, the general aim is to study whether this type of representation increases the speed of the reasoning mechanism, and why not – its accuracy. With the same purpose, we refer to a possible integration with the technique in [6] that handles point-to graphs via a stack-based algorithm for fixed-point computations.

We are also interested to what extent an abstract semantics based on aliases for SCOOP can be exploited for building more efficient analysis tools such as deadlock detectors, for instance. A survey on similar techniques that abstract away from possibly irrelevant information w.r.t. the problem under consideration is provided in [19].

Acknowledgements We are grateful for valuable comments to the anonymous reviewers, Măriuca Asăvoae, Alexander Kogtenkov, José Meseguer, Bertrand Meyer, Benjamin Morandi and Sergey Velder. The research leading to these results has received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007-2013) / ERC Grant agreement no. 291389.

References

1. E. Albert, P. Arenas, S. Genaim, and G. Puebla. Field-sensitive value analysis by field-insensitive analysis. In *Proceedings of the 2Nd World Congress on Formal Methods, FM '09*, pages 370–386, Berlin, Heidelberg, 2009. Springer-Verlag.
2. I. M. Asavae. Abstract semantics for alias analysis in K. *Electr. Notes Theor. Comput. Sci.*, 304:97–110, 2014.
3. A. Bouajjani, J. Esparza, and O. Maler. Reachability analysis of pushdown automata: Application to model-checking. In *CONCUR*, pages 135–150, 1997.
4. M. Burke, P. Carini, J.-D. Choi, and M. Hind. Flow-insensitive interprocedural alias analysis in the presence of pointers. In K. Pingali, U. Banerjee, D. Gelernter, A. Nicolau, and D. Padua, editors, *Languages and Compilers for Parallel Computing*, volume 892 of *Lecture Notes in Computer Science*, pages 234–250. Springer Berlin Heidelberg, 1995.
5. G. Caltais. Expression-based aliasing for OO-languages. *CoRR*, abs/1409.7509, 2014.
6. D. R. Chase, M. N. Wegman, and F. K. Zadeck. Analysis of pointers and structures. In *PLDI*, pages 296–310, 1990.
7. J.-D. Choi, M. Burke, and P. Carini. Efficient flow-sensitive interprocedural computation of pointer-induced aliases and side effects. In *Proceedings of the 20th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL '93*, pages 232–245, New York, NY, USA, 1993. ACM.
8. P. Cousot and R. Cousot. Abstract interpretation and application to logic programs. *J. Log. Program.*, 13(2&3):103–179, 1992.
9. A. Diwan, K. S. McKinley, and J. E. B. Moss. Type-based alias analysis. *SIGPLAN Not.*, 33(5):106–117, May 1998.
10. M. Emami, R. Ghiya, and L. J. Hendren. Context-sensitive interprocedural points-to analysis in the presence of function pointers. In *Proceedings of the ACM SIGPLAN 1994 Conference on Programming Language Design and Implementation, PLDI '94*, pages 242–256, New York, NY, USA, 1994. ACM.
11. M. Hind. Pointer analysis: haven’t we solved this problem yet? In *PASTE*, pages 54–61, 2001.
12. M. Hind, M. Burke, P. Carini, and J.-D. Choi. Interprocedural pointer alias analysis. *ACM Trans. Program. Lang. Syst.*, 21(4):848–894, July 1999.
13. C. A. R. Hoare, B. Möller, G. Struth, and I. Wehrman. Concurrent Kleene algebra. In *CONCUR 2009 - Concurrency Theory, 20th International Conference, CONCUR 2009, Bologna, Italy, September 1-4, 2009. Proceedings*, pages 399–414, 2009.
14. S. C. Kleene. Representation of events in nerve nets and finite automata. In C. Shannon and J. McCarthy, editors, *Automata Studies*, pages 3–41. Princeton University Press, Princeton, NJ, 1956.

15. A. Kogtenkov, B. Meyer, and S. Velder. Alias and change calculi, applied to frame inference. *CoRR*, abs/1307.3189, 2013.
16. W. Landi. Undecidability of static analysis. *ACM Lett. Program. Lang. Syst.*, 1(4):323–337, Dec. 1992.
17. W. Landi and B. G. Ryder. Pointer-induced aliasing: A problem classification. In *Proceedings of the 18th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '91, pages 93–103, New York, NY, USA, 1991. ACM.
18. J. R. Larus and P. N. Hilfinger. Detecting conflicts between structure accesses. In *PLDI*, pages 21–34, 1988.
19. J. Meseguer and G. Rosu. The rewriting logic semantics project: A progress report. In *Fundamentals of Computation Theory - 18th International Symposium, FCT 2011, Oslo, Norway, August 22-25, 2011. Proceedings*, pages 1–37, 2011.
20. B. Meyer. *Eiffel: The Language*. Prentice-Hall, 1991.
21. A. Miné. Field-sensitive value analysis of embedded c programs with union types and pointer arithmetics. In *Proceedings of the 2006 ACM SIGPLAN/SIGBED Conference on Language, Compilers, and Tool Support for Embedded Systems*, LCTES '06, pages 54–63, New York, NY, USA, 2006. ACM.
22. B. Morandi, M. Schill, S. Nanz, and B. Meyer. Prototyping a concurrency model. In *ACSD*, pages 170–179, 2013.
23. E. M. Myers. A precise inter-procedural data flow algorithm. In *Proceedings of the 8th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '81, pages 219–230, New York, NY, USA, 1981. ACM.
24. P. Nienaltowski. *Practical Framework for Contract-based Concurrent Object-oriented Programming*. ETH, 2007.
25. M. O. Rabin and D. Scott. Finite automata and their decision problems. *IBM J. Res. Dev.*, 3(2):114–125, Apr. 1959.
26. V. Robert and X. Leroy. A formally-verified alias analysis. In *CPP*, pages 11–26, 2012.
27. G. Rosu and T. F. Serbanuta. K overview and SIMPLE case study. In *Proceedings of International K Workshop (K'11)*, ENTCS. Elsevier, 2013. To appear.
28. T.-F. Serbanuta and G. Rosu. K-Maude: A rewriting based tool for semantics of programming languages. In *WRLA*, pages 104–122, 2010.
29. M. Sridharan, S. Chandra, J. Dolby, S. J. Fink, and E. Yahav. Alias analysis for object-oriented programs. In *Aliasing in Object-Oriented Programming*, pages 196–232. 2013.
30. R. P. Wilson and M. S. Lam. Efficient context-sensitive pointer analysis for C programs. In *Proceedings of the ACM SIGPLAN 1995 Conference on Programming Language Design and Implementation*, PLDI '95, pages 1–12, New York, NY, USA, 1995. ACM.