

# On the Verification of SCOOP Programs

Georgiana Caltais<sup>a</sup>, Bertrand Meyer<sup>a,b</sup>

<sup>a</sup>*Department of Computer Science, ETH Zürich, Switzerland*

<sup>b</sup>*Software Engineering Lab, Innopolis University, Russia*

---

## Abstract

In this paper we focus on the development of a toolbox for the verification of programs in the context of SCOOP – an elegant concurrency model, recently formalized based on Rewriting Logic (RL) and Maude. SCOOP is implemented in Eiffel and its applicability is demonstrated also from a practical perspective, in the area of robotics programming. Our contribution consists in devising and integrating an alias analyzer and a Coffman deadlock detector under the roof of the same RL-based semantic framework of SCOOP. This enables using the Maude rewriting engine and its LTL model-checker “for free”, in order to perform the analyses of interest. We discuss the limitations of our approach for model-checking deadlocks and provide solutions to the state explosion problem. The latter is mainly caused by the size of the SCOOP formalization which incorporates all the aspects of a real concurrency model. On the aliasing side, 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 easily implementable on top of the SCOOP formalization. An important property of our extension is that, in non-concurrent settings, the corresponding alias expressions can be over-approximated in terms of a notion of regular expressions. This further enables us to derive 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.

*Keywords:* concurrency, SCOOP, operational semantics, alias analysis, deadlock detection, Maude, rewriting logic

---

## 1. Introduction

In light of the widespread deployment and complexity of concurrent systems, the development of corresponding frameworks for rigorous design and analysis has been a great challenge. Along this research direction, the focus can be two-fold. On the one hand, the interest might be on formalizing and reasoning about a concurrency model, its characteristic concepts and synchronization mechanisms, for instance. On the other hand, of equal importance, the efforts might be directed towards checking the behavior of concurrent applications and their properties. These two research areas are closely related: it is often the case that, for example, analysis tools for concurrent applications depend on the underlying concurrency model. Hence, the development of a *unifying framework for the design and analysis of both the model and its applications* is of interest.

In this paper we are targeting SCOOP [1], 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, and 2) the burden of orchestrating concurrent executions is handled within the model, therefore reducing the risk of correctness issues. The reference implementation is Eiffel [2], but implementations have also been built on top of languages such as Java. The success of SCOOP is

---

*Email addresses:* [georgiana.caltais@inf.ethz.ch](mailto:georgiana.caltais@inf.ethz.ch) (Georgiana Caltais), [bertrand.meyer@inf.ethz.ch](mailto:bertrand.meyer@inf.ethz.ch) (Bertrand Meyer)

demonstrated not only from a research perspective, but also from a practical perspective, with applications appearing, for instance, in the area of robotics programming [3].

The basis of a framework for the design and analysis of the SCOOP model has already been set. In this respect, we refer to the recent formalization of SCOOP in [1] based on Rewriting Logic (RL) [4], which is “executable” and straightforwardly implementable in the programming language Maude [5]. In [5] these capabilities have been successfully exploited in order to reason on the original SCOOP model and to identify a number of design flaws.

Moreover, an executable semantics can be exploited in order to formalize and “run” analysis tools for SCOOP programs as well. This facilitates the extension of the aforementioned SCOOP formalization to the level of a unifying executable semantic framework for the design and analysis of both the model and its concurrent applications. In this paper we focus on the *development of a RL-based toolbox for the analysis of SCOOP programs* on top of the formalization in [1]. We are interested in constructing an alias analyzer and a deadlock detector.

*Alias analysis* has been an interesting research direction for the 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 [6, 7] that handle isolated functions only, and their inter-procedural counterparts [7, 8, 9] that consider the interactions between function calls; (ii) type-based techniques [10]; (iii) flow-based techniques [11, 12] that establish aliases depending on the control-flow information of a procedure; (iv) context-(in)sensitive approaches [13, 14] that depend on whether the calling context of a function is taken into account or not; (v) field-(in)sensitive approaches [15, 16] that depend on whether the individual fields of objects in a program are traced or not.

There is a huge literature on heap analysis for aliasing [17], but hardly any paper that presents a calculus allowing the derivation of alias relations as the result of applying various instructions of a programming language. Hence, 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 [18]. The 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 [18] abstracts the aliasing information in terms of explicit access paths [19] referred to as *alias expressions* straightforwardly computed in an equational fashion, based on the language constructs. As we shall see later on in this paper, the language-based expressions can be exploited in order to reason on “may aliasing” in a finite number of steps in non-concurrent settings and, moreover, can be easily incorporated in the semantic rules defining SCOOP in [1].

*Deadlock* is one of the most serious problems in concurrent systems. It occurs when two or more executing threads are each waiting for the other to finish. Along time, the complexity of the problem determined various approaches to combat deadlocks [20]. Examples include: (i) deadlock prevention [21] which ensures that at least one of the deadlock conditions cannot hold, (ii) deadlock avoidance [22] that provides a priori information so that the system can predict and avoid deadlock situations, (iii) deadlock detection [23, 24] that detects and recovers from a deadlock state.

Our focus is on *deadlock detection* for SCOOP programs. We base our work on the fact that this type of analysis is in strict connection with the underlying model of interest. Consequently, as described in the corresponding subsequent sections, our approach consists in formalizing deadlocks in the context of the SCOOP concurrency model and enriching its semantics in [1] with the equivalent operational-based definition of deadlocks. This enables using the Maude rewriting capabilities “for free” in order to test SCOOP programs for deadlock. Nevertheless, the more ambitious goal of using the Maude LTL model-checker for deadlock detection is not straightforward. As discussed in more detail later on in this paper, verification of deadlocks was possible after reducing the SCOOP semantics in [1] and abstracting it based on aliasing information, and modifying a series of implementation aspects (such as indexed-based parameterizations) that determined state explosion issues.

*Our contribution.* This paper is an extended version of [25] where we proposed:

1. a translation of the (finite) alias calculus in [18] to the setting of unbounded program executions such as infinite loops and recursive calls, together with a sound over-approximation technique based on (finitely representable) “regular alias expressions” capturing unbounded executions in non-concurrent settings;
2. a RL-based specification of the extended calculus suitable for integration within the SCOOP formalization in [1] (for this purpose we chose the  $\mathbb{K}$  semantic framework as a RL-based formalism enabling compact and modular definitions);
3. an algorithm for “may aliasing” (exploiting the finiteness property in 1.) that always terminates in non-concurrent settings.

The current work adds to 1.–3. above:

4. the full RL-based specification in 2. and the complete formal proofs showing the soundness of the over-approximating technique based on “regular alias expressions”;
5. examples of exploiting the algorithm in 3. and its implementation on top of the SCOOP formalization in Maude [1];
6. the formalization and integration of a deadlock detection mechanism on top of the SCOOP operational semantics [1], together with discussions on the limitations of our approach and associated workarounds.

*Paper structure.* The paper is organized as follows. In Section 2 we provide a brief overview of SCOOP. In Section 3 we introduce the extension of the alias calculus in [18] to unbounded executions. In Section 4 we provide the full RL-based executable specification of the calculus. The implementation in SCOOP and further applications are discussed in Section 5. Section 6 is dedicated to deadlocking in SCOOP. In Section 7 we draw the conclusions, discuss some of the related works and provide pointers to future developments.

## 2. Biref introduction to SCOOP

As already stated, the purpose of the current work is the development of a toolbox for the analysis of SCOOP programs by exploiting the semantics proposed in [1]. SCOOP is particularly attractive due to its simplicity and elegance, as it allows the switch from sequential to concurrent programming in a rather straightforward fashion, by means of just one keyword, namely, *separate*. Transparent to the user, the key notion in SCOOP is the processor, or *handler* (that can be a CPU, or it can also be implemented in software, as a process or thread). Handlers are in charge of executing the routines of “separate” objects, in a concurrent fashion.

For an example, assume a processor  $p$  that performs a call  $o.f(a_1, a_2, \dots)$  on an object  $o$ . If  $o$  is declared as “separate”, then  $p$  sends a request for executing  $f(a_1, a_2, \dots)$  to  $q$  – the handler of  $o$  (note that  $p$  and  $q$  can coincide). Meanwhile,  $p$  can continue. Moreover, assume that  $a_1, a_2, \dots$  are of “separate” types. According to the SCOOP semantics, the application of the call  $f(\dots)$  will *wait* until it has been able to *lock* all the separate objects associated to  $a_1, a_2, \dots$ . This mechanism guarantees exclusive access to these objects. Given a processor  $p$ , by  $W(p)$  we denote the set of processors  $p$  *waits* to release the resources  $p$  needs for its asynchronous execution. Orthogonally, by  $H(p)$  we represent the set of resources (more precisely, resource handlers that)  $p$  already acquired.

The semantics of SCOOP in [1] is defined over tuples of shape

$$\langle p_1 :: St_1 \mid \dots \mid p_n :: St_n, \sigma \rangle \tag{1}$$

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

Roughly speaking, one could classify the operational rules formalizing SCOOP in [1] in: a) *language rules* that provide the semantics of language constructs such as “**if ... then ... else ... end**” or “**until ... loop ... end**”, and b) *control rules* implementing mechanisms such as locking or scheduling.

For an example in category a) above, consider the rules specifying “**if**” instructions:

$$\frac{\text{a is fresh}}{\langle p :: \mathbf{if} \ e \ \mathbf{then} \ St_1 \ \mathbf{else} \ St_2 \ \mathbf{end}; St, \sigma \rangle \rightarrow \langle p :: \text{eval}(a, e); \text{wait}(a); \text{provided } a.\text{data} \ \text{then} \ St_1 \ \text{else} \ St_2; St, \sigma \rangle} \quad (2)$$

$$\frac{\cdot}{\langle p :: \text{provided } \text{true} \ \text{then} \ St_1 \ \text{else} \ St_2; St, \sigma \rangle \rightarrow \langle p :: St_1; St, \sigma \rangle} \quad (3)$$

$$\frac{\cdot}{\langle p :: \text{provided } \text{false} \ \text{then} \ St_1 \ \text{else} \ St_2; St, \sigma \rangle \rightarrow \langle p :: St_2; St, \sigma \rangle} \quad (4)$$

Intuitively, “ $\text{eval}(a, e)$ ” evaluates  $e$  and puts the result on a fresh channel  $a$  and “ $\text{wait}(a)$ ” enables processor  $p$  to use the evaluation result stored in  $a.\text{data}$ . It is straightforward to see that, according to (3), in case the condition  $e$  is evaluated to *true* then the “**if** branch”  $St_1$  is placed on top of the call stack of  $p$ . Otherwise, based on (4), if  $e$  is evaluated to *false*, the “**else** branch” is executed.

As we shall see in Section 5, an operational view on the alias calculus in [18] exploiting the instructions of a programming language will enable a straightforward implementation on top of the “language rules” of SCOOP.

For the case b) above we refer to the locking rule:

$$\frac{\forall q_i \in \{q_1, \dots, q_m\} : \sigma.\text{rq\_locked}(q_i) = \text{false}}{\langle p :: \text{lock}(\{q_1, \dots, q_m\}); St, \sigma \rangle \rightarrow \langle p :: St, \sigma.\text{lock\_rqs}(p, \{q_1, \dots, q_m\}) \rangle} \quad (5)$$

stating that a processor  $p$  can lock a set of handlers  $\{q_1, \dots, q_m\}$  by calling  $\text{lock\_rqs}$  on the state  $\sigma$  whenever none of the handlers  $q_i$  has previously been acquired by other processors, *i.e.*,  $\sigma.\text{rq\_locked}(q_i) = \text{false}$ .

As it will become clear in Section 6, “control rules” pave the way to an immediate implementation of a corresponding “deadlock rule” on top of the Maude formalization of SCOOP in [1].

### 3. The alias calculus

The calculus for *may aliasing* introduced in [18] abstracts the aliasing information in terms of explicit access paths referred to as “alias expressions”. Consider, for an example, the case of a linked list. We write  $x_i$  ( $i \geq 0$ ) to represent node  $i$  in the list, and use a setter to assign the next node of the list:

$$\begin{array}{l} \mathbf{create} \ x_0 \\ \mathbf{loop} \\ \quad i := i + 1 \\ \quad \mathbf{create} \ x_i \\ \quad \quad x_i.\text{set\_next}(x_{i-1}) \\ \mathbf{end} \end{array} \quad (6)$$

The result of the execution of the code above can be intuitively depicted as the infinite sequence:

$$x_0 \xleftarrow{\text{next}} x_1 \xleftarrow{\text{next}} \dots x_{k-1} \xleftarrow{\text{next}} x_k \xleftarrow{\text{next}} x_{k+1} \dots$$

Hence,  $x_0$  becomes aliased to  $x_1.\text{next}$ ,  $x_2.\text{next.next}$ ,  $x_3.\text{next.next.next}$ , so on and so on. In short, the set of associated alias expressions can be equivalently written as:

$$\{[x_i, x_{i+k}.\text{next}^k] \mid i \geq 0 \wedge k \geq 1\}. \quad (7)$$

The sources of imprecision introduced by the calculus in [18] are limited to ignoring tests in conditionals, and to “cutting at length  $L$ ” for the case of possibly infinite alias relation corresponding to unbounded executions as in (6). The cutting technique considers sequences longer than a given length  $L$  as aliased to all expressions.

In this section we define an extension of the calculus in [18], to unbounded program executions. Moreover, based on the idea behind the *pumping lemma for regular languages* [26], 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 4, in the formal setting of the  $\mathbb{K}$  semantic framework [27]. Note that  $\mathbb{K}$  is used more as a notational convention, as its operational flavor enables a straightforward integration within the SCOOP formalization in [1].

*Brief overview of the alias calculus.* We proceed by recalling the notion of *alias relation* and a series of associated notations and basic operations, as introduced in [18].

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  $e.\text{Current} = \text{Current}.e = e$ . 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  $t = x.\text{next}.\text{next}$ . 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  $u$ .

### 3.1. Extension to unbounded executions

We further introduce an extension of the alias calculus in [18] 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 a program affects the aliasing between expressions. As in [18], the calculus ignores tests in conditionals and loops. The *program instructions* are defined as follows:

$$\begin{aligned} p ::= & p; p \mid \mathbf{then} \ p \ \mathbf{else} \ p \ \mathbf{end} \mid \\ & \mathbf{create} \ x \mid \mathbf{forget} \ x \mid t := s \mid \\ & \mathbf{loop} \ p \ \mathbf{end} \mid \mathbf{call} \ f(l) \mid x.\mathbf{call} \ f(l). \end{aligned} \tag{8}$$

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:

$$r \gg (p; q) = (r \gg p) \gg q. \tag{9}$$

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:

$$r \gg (\mathbf{then} \ p \ \mathbf{else} \ q \ \mathbf{end}) = r \gg p \cup r \gg q. \tag{10}$$

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:

$$r \gg (\mathbf{loop} \ p \ \mathbf{end}) = \bigcup_{n \in \mathbb{N}} (r \gg p^n) \quad (11)$$

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 [18] that proposes a “cutting” technique restricting the model to a maximum length  $L$ . In [18], 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 3.2.

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

$$\begin{aligned} r \gg (\mathbf{create} \ x) &= r - x \\ r \gg (\mathbf{forget} \ x) &= r - x. \end{aligned} \quad (12)$$

The (qualified) function calls comply to their initial definitions in [18]:

$$\begin{aligned} r \gg (\mathbf{call} \ f(l)) &= (r[f^\bullet:l]) \gg |f| \\ r \gg (x.\mathbf{call} \ f(l)) &= x.((x'.r) \gg \mathbf{call} \ f(x'.l)). \end{aligned} \quad (13)$$

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 [18]:

$$r \gg (t := s) = \mathbf{given} \ r_1 = r[ot = t] \ \mathbf{then} \ (r_1 - t)[t = (r_1/s - t)] - ot \ \mathbf{end} \quad (14)$$

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 (10) and (14) 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.*

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

$$p ::= p; p \mid \mathbf{then} \ p \ \mathbf{else} \ p \ \mathbf{end} \mid \mathbf{create} \ x \mid \mathbf{forget} \ x \mid t := s. \quad (15)$$

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 (6):

```

create  $x_0$ 
loop
   $i := i + 1$ 
  create  $x_i$ 
   $x_i.set\_next(x_{i-1})$ 
end

```

Its execution generates an alias relation including an infinite number of pairs of shape:

$$[x_i, x_{i+1}.next], [x_i, x_{i+2}.next.next], [x_i, x_{i+3}.next.next.next] \dots \quad (16)$$

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 (16), which are pairs of type  $[x_i, x_{i+k}.next^{k \geq 1}]$ .

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 [28]. We call an alias relation *regular* if it consists of pairs of regular expressions.

**Lemma 2.** *Assume  $p$  a program built according to the rules in (8). Then, in a sequential setting, the relation  $\emptyset \gg p$  is regular.*

In order to prove Lemma 2, we proceed by demonstrating a series of intermediate results.

**Remark 3.** *We first observe that the operations  $r/s$ ,  $r - x$ , dot-completeness and  $r[x = u]$  introduced in Section 3 preserve the regularity of an alias relation  $r$ .*

Then, we define a notion of *finite execution* control blocks:

$$p ::= \mathbf{create} \ x \mid \mathbf{forget} \ x \mid t := s \mid p; p \mid \mathbf{then} \ p \ \mathbf{else} \ p \ \mathbf{end} \mid \mathbf{call} \ f(l) \mid x.\mathbf{call} \ f(l) \quad (17)$$

where  $f$  stands for a non-recursive function.

It is easy to see that the execution of control blocks as in (17) preserve the regularity of alias relations as well.

**Lemma 4.** *For all regular alias relations  $r$  and  $p$  a finite-execution control block, in a sequential setting, it holds that  $r \gg p$  is also regular.*

*Proof.* The proof follows immediately, by induction on the structure of  $p$  and Remark 3. Base cases are: **create**  $x$ , **forget**  $x$  and  $t := s$ . For function calls, the result is a consequence of their corresponding unfolding, based on the definitions in (13).  $\square$

**Remark 5.** *With respect to may aliasing, recursive calls can be handled via loops. Consider, for instance the recursive function*

$$f(x) \{ B_1; f(y); B_2 \}$$

where  $B_1$  and  $B_2$  are instruction blocks built as in (8). It is intuitive to see that computing the may aliases resulted from the execution of  $f(x)$  reduces executing unfoldings of shape:

**loop**  $B_1$  **end**; **loop**  $B_2$  **end**.

Moreover, unbounded program executions also preserve regularity.

**Lemma 6.** *For all regular alias relations  $r$  and  $p$  a control block that can execute unboundedly, in a sequential setting, it holds that  $r \gg p$  is also regular.*

*Proof.* The proof follows by induction on the number of nested loops in  $p$  and Remark 5.  $\square$

At this point, the result in Lemma 2 follows immediately by Lemma 4 and Lemma 6.

Inspired by the idea behind the *pumping lemma for regular languages* [26], 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:

$$\text{lasso}(r, r') = ([x_1y_1z_1, x_2y_2z_2] \in r \text{ iff } [x_1y_1y_1z_1, x_2y_2y_2z_2] \in r'). \quad (18)$$

For the simplicity of notation we sometimes omit the dot-separators between expressions. For instance, we write  $xyz$  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:

$$\text{reg}(r, r') = \{ [x_1y_1^*z_1, x_2y_2^*z_2] \mid [x_1y_1z_1, x_2y_2z_2] \in r \wedge [x_1y_1y_1z_1, x_2y_2y_2z_2] \in r' \} \quad (19)$$

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 (15), in the way that it does not introduce any false negatives:

**Lemma 7.** *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.* We proceed by induction on  $n$ .

- *Base case:*  $n = 1$ . By hypothesis it holds that  $\text{lasso}(r, r') = \text{true}$ . Hence, according to the definition of  $\text{lasso}(-, -)$  in (18), there exists a one-to-one correspondence of the shape

$$[x_1y_1z_1, x_2y_2z_2] \in r \text{ iff } [x_1y_1y_1z_1, x_2y_2y_2z_2] \in r'$$

between the elements of  $r$  and  $r'$ , respectively.

Consequently, by the definition of  $\text{reg}(-, -)$  in (19), it is easy to see that

$$r' \in \text{reg}(r, r').$$

- *Induction step.* Fix a natural number  $n$  and suppose that

$$r \gg p^k \in \text{reg}(r, r') \tag{20}$$

for all  $k \in \{1, \dots, n\}$ . We want to prove that (20) holds also for  $k = n + 1$ .

We continue by “*reductio ad absurdum*”. Consider

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

and assume that

$$\bar{r} \gg p \notin \text{reg}(r, r') \tag{21}$$

Clearly, the execution of  $p$  when starting with  $\bar{r}$  identifies an alias pair which is not in  $\text{reg}(r, r')$ . Given that  $p$  is a basic control block as in (15), and based on the corresponding definitions in (9)–(14), it is not difficult to observe that the regular structure of the alias information can only be broken via a new added pair  $(t, s')$  associated to an assignment  $t := s$  within  $p$ .

Let  $p = C[t := s]$ , where  $C$  is a context built according to (15), and  $t := s$  is the upper-most assignment instruction in the syntactic tree associated to  $p$ , that introduces a pair  $[t, s']$  which is not in  $\text{reg}(r, r')$ . Assume that  $\tilde{r}$  is the intermediate alias relation obtained by reducing  $\bar{r} \gg C[t := s]$  according to the equations (9)–(14), before the application of the assignment axiom corresponding to  $t := s$ .

Note that  $t := s$  was executed at least once before, as  $n \geq 1$ , and observe that  $\tilde{r} \in \text{reg}(r, r')$ . Hence, we identify two situations in the context of the aforementioned execution: (a) either all the newly added pairs corresponding to the assignment  $t := s$  complied to the regular structure, or (b) each new pair  $[t', s']$  that did not fit the regular pattern was later removed via a subsequent instruction “**create**  $u$ ” or “**forget**  $u$ ” within  $p$ , with  $u$  a prefix of  $t'$  or  $s'$ .

If the case (a) above was satisfied, then, based on the definition of dot-completeness, a pair

$$(t, s') \in (\tilde{r}_1 - t)[t = \tilde{r}_1/s - t] - ot,$$

where

$$\tilde{r}_1 = \tilde{r}[ot = t]$$

cannot break the regular pattern of the alias expressions either. For the case (b) above, all the “non-well-behaved” new pairs will be again removed via a subsequent “**create**  $u$ ” or “**forget**  $u$ ” within  $p$ .

Therefore, the assumption in (21) is false, so it holds that:

$$\bar{r} \gg p = r \gg p^{n+1} \in \text{reg}(r, r').$$

□

#### 4. 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 [1] (we refer to Section 5 for more details on our approach).

In short, our strategy is to start with a program built on top of the control structures in (8), 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 [4]. It is suitable for defining (concurrent) languages and corresponding formal analysis tools, with straightforward implementation in  $\mathbb{K}$ -Maude [29].  $\mathbb{K}$ -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  $[\text{name}]$  stands for the *name of the cell*. A construction  $\langle \cdot \rangle_n$  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

$$\langle c \dots \rangle_{n_1} \langle c' \rangle_{n_2} \Rightarrow \langle c' \dots \rangle_{n_1} \langle \dots c' \rangle_{n_3} \quad (22)$$

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, (22) is written in a  $\mathbb{K}$ -like syntax as:

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

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 7.

The following  $\mathbb{K}$ -rules are straightforward, based on the axioms (9)–(14) in Section 3.1. Namely, the rule implementing an instruction  $p; q$  simply forces the sequential execution of  $p$  and  $q$  by positioning  $p \curvearrowright q$  at the top of the continuation cell:

$$\langle \frac{p; q}{p \curvearrowright q} \dots \rangle_k \quad (23)$$

Handling **create**  $x$  and **forget**  $x$  complies to the associated definitions. Namely, it updates the current alias relation by removing all the pairs having (at least) one element with  $x$  as prefix. In addition, it also pops the corresponding instruction from the continuation stack:

$$\langle \frac{r}{r - x} \rangle_{\text{al}} \langle \frac{\text{create } x \dots}{\cdot} \rangle_k \quad \langle \frac{r}{r - x} \rangle_{\text{al}} \langle \frac{\text{forget } x \dots}{\cdot} \rangle_k \quad (24)$$

The assignment rule restores the current alias relation according to its axiom in (14), and removes the assignment instruction from the top of the  $k$ -cell:

$$\langle \frac{r}{(r_1 - t)[t = (r_1/s - t)] - ot} \rangle_{\text{al}} \quad \langle \frac{t := s \dots}{.} \rangle_k \quad \text{with } r_1 = r[ot = t] \quad (25)$$

The  $\mathbb{K}$ -implementation of a **then  $p$  else  $q$  end** statement is more sophisticated, as it instruments a stack-based mechanism enabling the computation of the union of alias relations  $r \gg p \cup r \gg q$  in three steps. First, we define the  $\mathbb{K}$ -rule:

$$\langle r \rangle_{\text{al}} \quad \langle \frac{\mathbf{then } p \mathbf{ else } q \mathbf{ end} \dots}{p \boxed{\text{et}} q \boxed{\text{ee}}} \rangle_k \quad \langle \frac{.}{\langle r, p \rangle_t \langle r, q \rangle_e} \dots \rangle_{\text{bkt-te}} \quad (26)$$

saving at the top of the back-tracking stack  $\langle - \rangle_{\text{bkt-te}}$  the initial alias relation  $r$  to be modified by both  $p$  and  $q$ , via two cells  $\langle r, p \rangle_t$  and  $\langle r, q \rangle_e$ , respectively. Note that the original instruction in the  $k$ -cell is replaced by a meta-construction marking the end of the executions corresponding to the **then** and **else** branches with  $\boxed{\text{et}}$  and  $\boxed{\text{ee}}$ , respectively.

Second, whenever the successful execution of  $p$  (signaled by  $\boxed{\text{et}}$ ) at the top of the  $k$ -cell) builds an alias relation  $r'$ , the execution of  $q$  starting with the original relation  $r$  is forced by replacing  $r'$  with  $r$  in  $\langle - \rangle_{\text{al}}$ , and by positioning  $q \boxed{\text{ee}}$  at the top of the  $k$ -cell. The new alias information after  $p$ , denoted by  $\langle r', p \rangle_t$ , is updated in the back-tracking cell:

$$\langle \frac{r'}{r} \rangle_{\text{al}} \quad \langle \frac{\boxed{\text{et}} q \boxed{\text{ee}} \dots}{q \boxed{\text{ee}}} \rangle_k \quad \langle \frac{\langle r, p \rangle_t \langle r, q \rangle_e \dots}{\langle r', p \rangle_t} \rangle_{\text{bkt-te}} \quad (27)$$

Eventually, if the successful execution of  $q$  (marked by  $\boxed{\text{ee}}$  at the top of  $\langle - \rangle_k$ ) produces an alias relation  $r''$ , then the final alias information becomes  $r' \cup r''$ , where  $r'$  is the aliasing after  $p$ , stored as showed in (27). The corresponding back-tracking information is removed from  $\langle - \rangle_{\text{bkt-te}}$ , and the next program instruction is enabled in the  $k$ -cell:

$$\langle \frac{r''}{r' \cup r''} \rangle_{\text{al}} \quad \langle \frac{\boxed{\text{ee}} \dots}{.} \rangle_k \quad \langle \frac{\langle r', p \rangle_t \langle r, q \rangle_e \dots}{.} \rangle_{\text{bkt-te}} \quad (28)$$

For **loop  $p$  end**, we utilize a meta-construction  $p \boxed{1} \mathbf{loop } p \mathbf{ end}$  simulating the set union in (11), and a back-tracking stack  $\langle - \rangle_{\text{bkt-l}}$  collecting the alias information obtained after each execution of  $p$ . Moreover, the  $\mathbb{K}$ -implementation exploits the result in Lemma 7. 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  $\langle r \rangle_{\text{al-o}} \langle p \rangle_1$ :

$$\langle r \rangle_{\text{al}} \quad \langle \frac{\mathbf{loop } p \mathbf{ end} \dots}{p \boxed{1} \mathbf{loop } p \mathbf{ end}} \rangle_k \quad \langle \frac{.}{\langle r \rangle_{\text{al-o}} \langle p \rangle_1} \dots \rangle_{\text{bkt-l}} \quad (29)$$

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  $\langle - \rangle_{\text{bkt-l}}$ ) then  $p$  is constrained to a new execution by becoming the top of the  $k$ -cell, and  $r'$  is memorized for back-tracking:

$$\langle r' \rangle_{\text{al}} \quad \langle \frac{\boxed{1} \mathbf{loop } p \mathbf{ end} \dots}{p \boxed{1} \mathbf{loop } p \mathbf{ end}} \rangle_k \quad \langle \frac{\langle r \rangle_{\text{al-o}} \langle p \rangle_1 \dots}{\langle r' \rangle_{\text{al-o}} \langle p \rangle_1} \rangle_{\text{bkt-l}} \quad \text{if not lasso}(r, r') \quad (\clubsuit) \quad (30)$$

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  $\langle - \rangle_{\text{bkt-l}}$  and the **loop** instruction is eliminated from the  $k$ -cell:

$$\langle \frac{r'}{\text{reg}(r, r')} \rangle_{\text{al}} \quad \langle \frac{\boxed{1} \mathbf{loop } p \mathbf{ end} \dots}{.} \rangle_k \quad \langle \frac{\langle r \rangle_{\text{al-o}} \langle p \rangle_1 \dots}{.} \rangle_{\text{bkt-l}} \quad \text{if lasso}(r, r') \quad (\clubsuit) \quad (31)$$

For handling function calls such as  $\mathbf{call} f(l)$  we use a meta-construction  $| f | \boxed{f}$ . Here  $| f |$  stands for the body of  $f$  and  $\boxed{f}$  marks the end of the corresponding execution. Moreover, a stack  $\langle - \rangle_{\text{bkt-cf}}$  is utilized in order to store the alias information before each (possibly recursive) call of  $f$ , with the purpose of identifying the lassos generated by the (possibly repeated) execution of  $f$ . In order to guarantee a sound implementation of (mutually) recursive calls, both  $\boxed{f}$  and  $\langle - \rangle_{\text{bkt-cf}}$  are parameterized by  $f$  – the name of the function. An example illustrating this reasoning mechanism is provided in Section 4.1.

The first  $\mathbb{K}$  rule for handling function calls matches the associated axiom in (13): the alias information is set to  $r[f^\bullet:l]$ , whereas the next instructions to be executed are given by  $| f |$ . Note that the original aliasing is retained in the (initially empty) back-tracking cell via  $\langle r \rangle_{\text{al-o}}$ .

$$\langle \frac{r}{r[f^\bullet:l]} \rangle_{\text{al}} \quad \langle \frac{\mathbf{call} f(l) \dots}{| f | \boxed{f}} \rangle_{\text{k}} \quad \langle \frac{\cdot}{\langle r \rangle_{\text{al-o}}} \rangle_{\text{bkt-cf}} \quad (32)$$

**Remark 8.** *Observe that the back-tracking cell does not need to be parameterized by the actual argument list  $l$  of  $f$ . Each such argument is anyways replaced in the current alias relation  $r$  by its counterpart in the formal argument list of  $f$ . In short:  $r$  becomes  $r[f^\bullet:l]$ .*

A successful execution of  $\mathbf{call} f(l)$  is distinguished by the occurrence of  $\boxed{f}$  at the top of the continuation stack. If this is the case, then the corresponding back-tracking alias information is removed from  $\langle - \rangle_{\text{bkt-cf}}$  and the next program instruction (if any) is enabled at the top of the  $k$ -cell:

$$\langle r' \rangle_{\text{al}} \quad \langle \boxed{f} \dots \rangle_{\text{k}} \quad \langle \langle r \rangle_{\text{al-o}} \dots \rangle_{\text{bkt-cf}} \quad (33)$$

Recursive calls are treated by means of two  $\mathbb{K}$ -rules. Note that a recursive context is identified whenever the current program instruction is of shape  $\mathbf{call} f(l)$  and the associated back-tracking structure is not empty, *i.e.*, rule (32) was previously applied. Then, if the recursive call of  $f$  when starting with  $r$  produces a lasso  $r'$ , the execution of  $f(l)$  is stopped by soundly over-approximating the alias information with  $\text{reg}(r, r')$ , according to Lemma 7, and by removing  $\mathbf{call} f(l)$  from the  $k$ -cell:

$$\langle \frac{r'}{\text{reg}(r, r')} \rangle_{\text{al}} \quad \langle \mathbf{call} f(l) \dots \rangle_{\text{k}} \quad \langle \langle r \rangle_{\text{al-o}} \dots \rangle_{\text{bkt-cf}} \quad \text{if } \text{lasso}(r, r') \quad (\clubsuit) \quad (34)$$

If a lasso is not reached, then the body of  $f$  is executed once more, and the current aliasing is pushed to the back-tracking cell:

$$\langle r' \rangle_{\text{al}} \quad \langle \frac{\mathbf{call} f(l) \dots}{| f | \boxed{f}} \rangle_{\text{k}} \quad \langle \frac{\cdot}{\langle r' \rangle_{\text{al-o}}} \langle r \rangle_{\text{al-o}} \dots \rangle_{\text{bkt-cf}} \quad \text{if not } \text{lasso}(r, r') \quad (\clubsuit) \quad (35)$$

Qualified calls  $x.\mathbf{call} f(l)$  are handled by two  $\mathbb{K}$ -rules as follows. First, based on the definition in (13), the “negative variable”  $x'$  transposing the context of the call to to the context of the caller is distributed to the elements of the initial alias relation  $r$ , and to  $l$  – the argument list of  $f$ . Moreover, a meta-construction  $\boxed{\text{qf}}$  is utilized in order to mark the end of the qualified call in the continuation cell, similarly to the rule (32). The caller is stored in a back-tracking stack  $\langle \cdot \rangle_{\text{bkt-qf}}$  also parameterized by  $f$  – the name of the function. The current instruction in the  $k$ -cell becomes  $\mathbf{call} f(x'.l)$ , as expected:

$$\langle \frac{r}{x'.r} \rangle_{\text{al}} \quad \langle \frac{x.\mathbf{call} f(l)}{\mathbf{call} f(x'.l) \boxed{\text{qf}}} \dots \rangle_{\text{k}} \quad \langle \frac{\cdot}{\langle x \rangle_{\text{f}}} \rangle_{\text{bkt-qf}} \quad (36)$$

Second, when the successful termination of the qualified call is signaled by  $\boxed{\text{qf}}$  at the top of the  $k$ -cell, the corresponding stored caller is distributed to the current alias relation and removed from the back-tracking cell. The next instruction in the continuation cell is released by eliminating the top  $\boxed{\text{qf}}$ :

$$\langle \frac{r}{x'.r} \rangle_{\text{al}} \quad \langle \boxed{\text{qf}} \dots \rangle_{\text{k}} \quad \langle \langle x \rangle_{\text{f}} \dots \rangle_{\text{bkt-qf}} \quad (37)$$

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

**Theorem 9.** *Consider  $p$  a program built on top of the control structures in (8), that executes in a sequential setting. Then, the application of the corresponding  $\mathbb{K}$ -rules when starting with  $p$  and an empty alias relation, is a finite rewriting of shape*

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

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  $\mathbb{K}$ -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 7.  $\square$

Observe that the *RL*-based machinery can simulate precisely the “cutting at length  $L$ ” technique in [18]. It suffices to disable the rules ( $\clubsuit$ ) and stop the rewriting after  $L$  steps.

The naturalness of applying the resulted aliasing framework is illustrated in the example in Section 4.1, for the case of two mutually recursive functions.

#### 4.1. The $\mathbb{K}$ -machinery by example

For an example, in this section we show how the  $\mathbb{K}$ -machinery developed in Section 4 can be used in order to extract the alias information for the case of two mutually recursive functions defined as:

$$f(x) \{ x := x.a; \text{ call } g(x) \} \qquad g(x) \{ x := x.b; \text{ call } f(x) \}$$

We assume that  $x$  is an object of a class with two fields  $a$  and  $b$ , respectively. We consider a sequential setting.

At first glance it is easy to see that the execution of  $\text{call } f(x)$ , when starting with an empty alias relation  $r$ , produces the alias expressions:

$$[x, x.(a.b)^*] \quad [x.a, x.(a.b)^*.a] \quad [x.b, x.(a.b)^*.b] \tag{38}$$

The associated reasoning in  $\mathbb{K}$  is depicted in Figure 4.1. The whole procedure starts with an empty alias relation  $r = \emptyset$ , and  $\text{call } f(x)$  in the continuation stack. Then, the corresponding  $\mathbb{K}$  rules (for handling assignments and function calls) are applied in the natural way.

A lasso is reached after two calls of  $f(x)$  that, consequently, determine two calls of  $g(x)$  – identified by  $\boxed{g} \boxed{f} \boxed{g} \boxed{f}$  in the  $k$ -cell. This triggers the application of rule (34) enabling the “regular” over-approximation as in Lemma 7.

Our example also illustrates the importance of isolating the back-traced alias information in cells of shape  $\langle \cdot \rangle_{\text{bkt-cf}}$  parameterized by the (possibly recursive) function  $f$ . More explicitly, rule (34) is soundly applied by identifying the aforementioned lasso based on: the current alias relation  $r_4$ , the recursive call  $f(l)$  at the top of the continuation, and the back-traced aliasing  $\langle \langle r_2 \rangle_{\text{al-o}} \dots \rangle_{\text{bkt-cf}}$  associated to the previous executions of  $f(l)$ .

As introduced in (18), an alias relation  $r'$  is a lasso of a relation  $r$  whenever there is a one-to-one correspondence between their elements as follows:

$$[x_1 y_1 z_1, x_2 y_2 z_2] \in r \quad \text{iff} \quad [x_1 y_1 y_1 z_1, x_2 y_2 y_2 z_2] \in r'$$

The current alias relation

$$r_4 = \{ [x, x.a.b.a.b], [x.a, x.a.b.a.b.a], [x.b, x.a.b.a.b.b] \},$$

before applying rule (34), is a lasso of

$$r_2 = \{[x, x.a.b], [x.a, x.a.b.a], [x.b, x.a.b.b]\}.$$

The aforementioned one-to-one correspondence is summarized in the following table:

$[x_1y_1z_1, x_2y_2z_2] \in r_2$ iff $[x_1y_1y_1z_1, x_2y_2y_2z_2] \in r_4$	$x_1$	$y_1$	$z_1$	$x_2$	$y_2$	$z_2$
$[x, x.a.b] \in r_2$ iff $[x, x.a.b.a.b] \in r_4$	$x$	$\varepsilon$	$\varepsilon$	$x$	$a.b$	$\varepsilon$
$[x.a, x.a.b.a] \in r_2$ iff $[x.a, x.a.b.a.b.a] \in r_4$	$x$	$\varepsilon$	$a$	$x$	$a.b$	$a$
$[x.b, x.a.b.b] \in r_2$ iff $[x.b, x.a.b.a.b.b] \in r_4$	$x$	$\varepsilon$	$b$	$x$	$a.b$	$b$

Here  $\varepsilon$  stands for the *empty alias expression*.

Moreover, according to rule (34), the lasso shaped by  $r_2$  and  $r_4$  also causes the (otherwise infinite) recursive calls to stop, as **call**  $f(l)$  is eliminated from the top of the  $k$ -cell. Hence, the rewriting process finishes with a sound over-approximation  $\text{reg}(r_2, r_4)$  replacing the current alias relation (cf. Lemma 7), defined precisely as in (38).

## 5. Aliasing in SCOOP

In this section we provide a brief overview on the integration and applicability of the alias calculus in SCOOP. First, recall from Section 2 that the Maude semantics of SCOOP in [1] is defined over tuples of shape

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

where,  $p_i$  and  $St_i$  stand for processors and their call stacks, respectively.  $\sigma$  is the state of the system and it holds information about the heap and the store.

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

$$\frac{\text{a is fresh}}{\Gamma \vdash \langle p :: t := s; St, \sigma \rangle \rightarrow \langle p :: \text{eval}(a, s); \text{wait}(a); \text{write}(t, a.data); St, \sigma \rangle} \quad (39)$$

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  $\Gamma$  is a typing environment [30] containing the class hierarchy of a program and all the type definitions.

At this point it is easy to understand that the K-rule for assignments

$$\left\langle \frac{r}{(r_1 - t)[t = (r_1/s - t)] - ot} \right\rangle^{\text{al}} \langle \underline{t := s} \dots \rangle^{\text{k}} \quad \text{with } r_1 = r[ot = t] \quad (25)$$

can be straightforwardly integrated in (39) by enriching the SCOOP configuration with a new component  $\text{alias\_}$  encapsulating the alias information, and considering instead the transition:

$$\Gamma \vdash \langle p :: t := s; St, \sigma, \text{alias}_{old} \rangle \rightarrow \langle p :: \text{eval}(a, s); \text{wait}(a); \text{write}(t, a.data); St, \sigma, \text{alias}_{new} \rangle$$

where

$$\text{alias}_{old} = r \quad \text{alias}_{new} = (r_1 - t)[t = (r_1/s - t)] - ot$$

with  $r$  and  $r_1$  as in (25). The integration of all the K-rules of the alias calculus on top of the Maude formalization of SCOOP is achieved by following a similar approach.

For a case study, one can download the SCOOP formalization at: <https://dl.dropboxusercontent.com/u/1356725/SCOOP-SCP.zip> and run the command

$$\begin{aligned}
& \langle r \rangle_{\text{al}} \langle \mathbf{call} f(x) \rangle_{\text{k}} \\
& \langle \cdot \rangle_{\text{bkt-cf}} \langle \cdot \rangle_{\text{bkt-cg}} \\
& \Downarrow (32) \\
& \langle r \rangle_{\text{al}} \langle x := x.a; \mathbf{call} g(x) \boxed{f} \rangle_{\text{k}} \\
& \langle \langle r \rangle_{\text{al-o}} \rangle_{\text{bkt-cf}} \langle \cdot \rangle_{\text{bkt-cg}} \\
& \Downarrow (25) \\
& \langle r_1 \rangle_{\text{al}} \langle \mathbf{call} g(x) \boxed{f} \rangle_{\text{k}} \\
& \langle \langle r \rangle_{\text{al-o}} \rangle_{\text{bkt-cf}} \langle \cdot \rangle_{\text{bkt-cg}} \\
& \text{where } r_1 = \{[x, x.a], [x.a, x.a.a], [x.b, x.a.b]\} \\
& \Downarrow (35) \\
& \langle r_1 \rangle_{\text{al}} \langle x := x.b; \mathbf{call} f(x) \boxed{g} \boxed{f} \rangle_{\text{k}} \\
& \langle \langle r \rangle_{\text{al-o}} \rangle_{\text{bkt-cf}} \langle \langle r_1 \rangle_{\text{al-o}} \rangle_{\text{bkt-cg}} \\
& \Downarrow (25) \\
& \langle r_2 \rangle_{\text{al}} \langle \mathbf{call} f(x) \boxed{g} \boxed{f} \rangle_{\text{k}} \\
& \langle \langle r \rangle_{\text{al-o}} \rangle_{\text{bkt-cf}} \langle \langle r_1 \rangle_{\text{al-o}} \rangle_{\text{bkt-cg}} \\
& \text{where } r_2 = \{[x, x.a.b], [x.a, x.a.b.a], [x.b, x.a.b.b]\} \\
& \Downarrow (35) \\
& \langle r_2 \rangle_{\text{al}} \langle x := x.a; \mathbf{call} g(x) \boxed{f} \boxed{g} \boxed{f} \rangle_{\text{k}} \\
& \langle \langle r_2 \rangle_{\text{al-o}} \langle r \rangle_{\text{al-o}} \rangle_{\text{bkt-cf}} \langle \langle r_1 \rangle_{\text{al-o}} \rangle_{\text{bkt-cg}} \\
& \Downarrow (25) \\
& \langle r_3 \rangle_{\text{al}} \langle \mathbf{call} g(x) \boxed{f} \boxed{g} \boxed{f} \rangle_{\text{k}} \\
& \langle \langle r_2 \rangle_{\text{al-o}} \langle r \rangle_{\text{al-o}} \rangle_{\text{bkt-cf}} \langle \langle r_1 \rangle_{\text{al-o}} \rangle_{\text{bkt-cg}} \\
& \text{where } r_3 = \{[x, x.a.b.a], [x.a, x.a.b.a.a], [x.b, x.a.b.a.b]\} \\
& \Downarrow (35) \\
& \langle r_3 \rangle_{\text{al}} \langle x := x.b; \mathbf{call} f(x) \boxed{g} \boxed{f} \boxed{g} \boxed{f} \rangle_{\text{k}} \\
& \langle \langle r_2 \rangle_{\text{al-o}} \langle r \rangle_{\text{al-o}} \rangle_{\text{bkt-cf}} \langle \langle r_3 \rangle_{\text{al-o}} \langle r_1 \rangle_{\text{al-o}} \rangle_{\text{bkt-cg}} \\
& \Downarrow (25) \\
& \langle r_4 \rangle_{\text{al}} \langle \mathbf{call} f(x) \boxed{g} \boxed{f} \boxed{g} \boxed{f} \rangle_{\text{k}} \\
& \langle \langle r_2 \rangle_{\text{al-o}} \langle r \rangle_{\text{al-o}} \rangle_{\text{bkt-cf}} \langle \langle r_3 \rangle_{\text{al-o}} \langle r_1 \rangle_{\text{al-o}} \rangle_{\text{bkt-cg}} \\
& \text{where } r_4 = \{[x, x.a.b.a.b], [x.a, x.a.b.a.b.a], [x.b, x.a.b.a.b.b]\} \\
& \Downarrow (34) \\
& \langle \mathbf{reg}(r_2, r_4) \rangle_{\text{al}} \langle \boxed{g} \boxed{f} \boxed{g} \boxed{f} \rangle_{\text{k}} \\
& \langle \langle r_2 \rangle_{\text{al-o}} \langle r \rangle_{\text{al-o}} \rangle_{\text{bkt-cf}} \langle \langle r_3 \rangle_{\text{al-o}} \langle r_1 \rangle_{\text{al-o}} \rangle_{\text{bkt-cg}} \\
& \Downarrow (*) (33) \\
& \langle \{[x, x.(a.b)^*], [x.a, x.(a.b)^*.a], [x.b, x.(a.b)^*.b]\} \rangle_{\text{al}} \langle \cdot \rangle_{\text{k}} \langle \cdot \rangle_{\text{bkt-cf}} \langle \cdot \rangle_{\text{bkt-cg}}
\end{aligned}$$

Figure 1: Aliasing and mutual recursion in  $\mathbb{K}$ .

> maude SCOOP.maude ..\examples\linked\_list-test.maude  
 corresponding to the execution of the code in (6)

```

create x0
loop
  i := i + 1
  create x_i
  x_i.set_next(x_{i-1})
end

```

for five iterations of the loop. 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` and setter `set_next` is declared as:

```

class 'NODE
  create {'make} (
    attribute { 'ANY } 'next : [?, . , 'NODE] ;
    procedure { 'ANY } 'set_next ( 'a_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 “entry point” of the program corresponds to the function `'make` in the (main) class `'LINKED_LIST_TEST` and is set via:

```
settings('LINKED_LIST_TEST, 'make, aliasing-on) .
```

Observe that the flag for performing the alias analysis is switched to “on”.

The relevant parts of the corresponding Maude output after executing the aforementioned command are as follows:

```

\|||||/
--- Welcome to Maude ---
/|||||
Maude 2.6 built: Mar 31 2011 23:36:02

```

```

rewrite in SYSTEM :
[...] settings('LINKED_LIST_TEST, 'make, aliasing-on))

|-
{0}proc(0) :: nil | {0}proc(1) :: nil,
{['x0 ; 'x0]} U {['x0 ; 'x1.'next]} U
{['x0 ; 'x2.'next.'next]} U {['x0 ; 'x3.'next.'next.'next]} U
{['x0 ; 'x4.'next.'next.'next.'next]} U
[...]
{['x3 ; 'x3]} U {['x3 ; 'x4.'next]}

state
  heap [...]
  store [...]
end

```

In short, one can see that two processors that were created finished executing the instructions on their corresponding call stacks: `{0}proc(0) :: nil` and `{0}proc(1) :: nil`. The aliased expressions include, as expected based on (7), pairs of shape  $[x_i ; x_{i+k}.next^k]$ . Moreover, the output displays the contents of the current system state, by providing information on the `heap` and `store`, as formalized in [1].

## 6. Deadlocking in SCOOP

In what follows we provide the details behind the formalization and the implementation of a deadlock detection mechanism for SCOOP. We discuss how Maude can be exploited in order to test and, respectively, model-check deadlocks in the context of SCOOP programs, we analyze the associated challenges and propose a series of corresponding solutions.

### 6.1. Formalizing deadlocks in SCOOP

Recall that the key idea of SCOOP is to associate to each object declared as *separate* a processor, or handler in charge of executing the routines of that object. Assume a processor  $p$  that performs a call  $o.f(a_1, a_2, \dots)$  on a separate object  $o$ , with separate arguments  $a_i$  ( $i \geq 1$ ). As previously stated in Section 2, in the SCOOP semantics, the application of the call  $f(\dots)$  will *wait* until it has been able to *lock* all the separate objects associated to  $a_1, a_2, \dots$ . This reservation mechanism enables deadlocks to occur whenever a set of processors reserve each other circularly. This situation might happen, for instance, in a Dining Philosophers scenario, where both philosophers and forks are objects residing on their own processors.

**Definition 10** (Deadlock). *For a processor  $p$ , let  $W(p)$  be the set of handlers  $p$  waits for its asynchronous execution, and  $H(p)$  represent the set of resources  $p$  already acquired. A deadlock exists if for some set  $D$  of processors the following holds:*

$$(\forall p \in D).(\exists p' \in D).(p \neq p') \wedge (W(p) \cap H(p') = \emptyset). \quad (40)$$

The integration of a deadlock detection mechanism based on Definition 10 on top of the SCOOP formalization in [1] is immediate. As already presented in Section 2, the operational semantics of SCOOP is defined over tuples of shape:

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

where,  $p_i$  and  $St_i$  stand for processors and their call stacks, respectively, and  $\sigma$  is the state of the system. Given a processor  $p'$  as in (40), the set  $H(p')$  corresponds, based on [1], to  $\sigma.rq\_locks(p')$ . Whenever the top of the instruction stack of a processor  $p$  is of shape  $lock(\{q_i, \dots, q_n\})$ , we say that the wait set  $W(p)$  is the set of processors  $\{q_1, \dots, q_n\}$ . Hence, assuming a predefined system configuration  $\langle deadlock \rangle$ , the SCOOP transition rule in Maude corresponding to (40) can be written as:

$$\frac{(\exists D \subseteq \sigma.procs).(\forall p \in D).(\exists p' \in D).(p \neq p') \wedge (aqs := \dots \mid p :: lock(\{q_i, \dots\}); St \mid \dots) \wedge (\sigma.rq\_locks(p').has(q_i))}{\langle aqs, \sigma \rangle \rightarrow \langle deadlock \rangle} \quad (41)$$

It is intuitive to guess that  $\sigma.procs$  in (41) returns the set of processors in the system, whereas  $aqs$  stands for the list of these processors and their associated instruction stacks (separated by the associative and commutative operator “ $\mid$ ”). We use “ $\dots$ ” to represent an arbitrary sequence of processors and processor stacks.

### 6.2. Testing deadlocks

We implemented (41) and tested the deadlock detection mechanism on top of the formalization in [1] for the Dining Philosophers problem. A case study considering two philosophers sharing two forks can be run by downloading the SCOOP formalization at:

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

and executing the command

```
> maude SCOOP.maude ..\examples\dining-philosophers-rewrite.maude
```

The class implementing the *philosopher* concept is briefly given below:

```

(class 'PHILOSOPHER
  create { 'make } (
    attribute {'ANY} 'left : [!,T,'FORK] ;
    attribute {'ANY} 'right : [!,T,'FORK] ;

    procedure { 'ANY } 'make ( 'f1 : [!,T,'FORK] ;
                               'fr : [!,T,'FORK] ; )
      do ( assign ('left, 'f1) ; assign ('right, 'fr) ; )
      [...]
    end ;

    procedure { 'ANY } 'eat_wrong (nil)
      do ( command ('Current . 'pick_in_turn('left ;)) ; )
      [...]
    end ;

    procedure { 'ANY } 'pick_in_turn ('f : [!,T,'FORK] ; )
      do ( command ('Current . 'pick_two('f ; 'right ;)) ; )
      [...]
    end ;

    procedure { 'ANY } 'pick_two ('fa : [!,T,'FORK] ;
                                   'fb : [!,T,'FORK] ; )
      do
        (
          command ('fa . 'use(nil)) ;
          command ('fb . 'use(nil)) ;
        )
      [...]
    end ;
  [...] end)

```

It declares two forks – 'left and 'right of type 'FORK, that can be handled by any processor (T) and that cannot be Void (!). Assume two philosophers p1 and p2 (of *separate* type PHILOSOPHER) and two forks f1 and f2 (of *separate* type FORK). Moreover, assume that 'left and 'right for p1 correspond to 'f1 and 'f2. For the case of p2 they correspond to 'f2 and 'f1, respectively. Asynchronously, p1 and p2 can execute `eat_wrong`, which calls `pick_in_turn(left)`. In the context of p1, the actual value of `left` is f1, whereas for p2 is f2. Consequently, both resources f1 and f2, respectively, may be locked “at the same time” by p1 and p2, respectively. Note that `pick_in_turn` subsequently calls `pick_two` that, intuitively, should enable the philosophers to use both forks. Thus, if f1 and f2, respectively, are locked by p1 and p2, respectively, the calls `pick_two(f2, f1)` and `pick_two(f1, f2)` corresponding to p1 and p2 will (circularly) wait for each other to finish. According to the SCOOP semantics, `pick_two(f1, f2)` is waiting for p2 to release f2, whereas `pick_two(f2, f1)` is waiting for p1 to release f1, as the forks are passed to `pick_two(...)` as *separate* types. In the context of SCOOP, this corresponds to a Coffman deadlock [31].

The entry point of the program implementing the Dining Philosophers example is the function 'make in the class APPLICATION, which asks the two philosophers p1 and p2 to adopt a wrong eating strategy as above, possibly leading to a deadlock situation. The flag enabling the deadlock analysis as in (41) is set to “on”. This information is specified using the instruction `settings('APPLICATION, 'make, deadlock-on)`.

Unfortunately, none of the executions of the Dining Philosophers scenario by simply invoking the Maude `rewrite` command lead to a deadlock situation. Each of our tests displayed the output:

```

\|/
--- Welcome to Maude ---
/|/
Maude 2.6 built: Mar 31 2011 23:36:02
Copyright 1997-2010 SRI International

```

```

rewrite in SYSTEM :
[...] settings('APPLICATION, 'make, deadlock-on)

```

```

|- {0}proc(0) :: nil | {0}proc(1) :: nil | {0}proc(3) :: nil
   {0}proc(5) :: nil | {0}proc(7) :: nil
   {0}proc(9) :: nil | {0}proc(11) :: nil

```

consisting of a list of processors (including the handlers of both the philosophers and the forks) with empty call stacks (`:: nil`). This indicates that every time, the two philosophers proceeded by lifting their forks simultaneously, hence no deadlock was possible.

One possible workaround is to use predefined strategies [32] in order to guide the rewriting of the Maude rules formalizing SCOOP towards a *deadlock* system configuration. An example of applying such a strategy for the Dining Philosophers case can be tested by running the command:

```
> maude SCOOP.maude ..\examples\dining-philosophers-strategy.maude
```

The command `srew [...]` using `init ; parallelism{lock} ; [...]` ; `deadlock-on` forces the execution of a `pick_in_turn` approach as above. This determines Maude to first trigger the rule `[init]` in the SCOOP formalization in [1]. This makes all the required initializations of the *bootstrap* processor. Then, one of the processors that managed to *lock* the necessary resources is (“randomly”) enabled to proceed to the asynchronous execution of its instruction stack, according to the strategy `parallelism{lock}` . The last step of the strategy calls the rule `[deadlock-on]` implementing the Coffman deadlock detection as in (41).

This time the guided rewriting leads, indeed, to one solution identifying a deadlock. The relevant parts of the corresponding Maude output are as follows:

```

\|/
--- Welcome to Maude ---
/|/
Maude 2.6 built: Mar 31 2011 23:36:02
Copyright 1997-2010 SRI International

```

```

srewrite in SYSTEM :
[...] settings('APPLICATION, 'make, deadlock-on)
using init ; parallelism{lock} ; [...]
```

 ; deadlock-on .

```

Solution 1
rewrites: 479677 in 2674887330ms cpu (7379ms real) (0 rewrites/second)
result Configuration: deadlock

```

```

No more solutions.
rewrites: 479677 in 2674887330ms cpu (7453ms real) (0 rewrites/second)

```

Nevertheless, such an approach requires lots of ingeniousness (our strategy has more than 300 rules!) and, moreover, is not automated.

### 6.3. Model-checking deadlocks

In this section we provide an overview on our approach to model-checking deadlocks for SCOOP, using the LTL Maude model-checker [33]. As mentioned in the beginning of the current paper, the idea behind our work is to exploit the unifying flavor of the Maude executable semantics of SCOOP [1]. The latter

integrates both the formalization of the language and its concurrency mechanisms, thus enabling using the semantic framework for program analysis purposes, “for free”.

One possible way to proceed is by simply running the Maude LTL model-checker via a search command. For model-checking deadlocks, one could invoke `search [b, d] subject =>* deadlock`. In a Dining Philosophers setting, for instance, this corresponds to performing a breadth-first search starting with `subject` – the corresponding (intermediate Maude representation of the) SCOOP program – to a `deadlock` state, in zero or more steps of proof (`=>*`). The optional arguments `b` and `d` provide an upper bound in the number of solutions to be found and, respectively, the maximum depth of the search.

Unfortunately, running the aforementioned search command led to the state explosion problem. At a first look, the issue was caused by the size of the SCOOP formalization in [1] which includes all the aspects of a real concurrency model.

As a first step, we reduced this formalization by eliminating the parts that are not relevant in the context of deadlocking; examples include the garbage collection and the exception handling mechanisms.

In addition, we provided a simplified, abstract semantics of SCOOP based on aliasing. This idea stems from the fact that SCOOP processors are known from object references, that may be aliased. Therefore, the SCOOP semantics can be simplified by retaining within the corresponding transition rules only the information important for aliasing. Consider, for instance, the rules (2)–(4) specifying “**if**” instructions in Section 2. The abstract transition rule omits the evaluation of the conditional and computes the aliasing information similarly to the semantics of `then . . . else . . . end` in (10), in Section 3. The abstraction collects the aliases resulted after the execution of both “**if**” and “**else**” branches:

$$\frac{\langle p :: \mathbf{if} \ e \ \mathbf{then} \ St_1 \ \mathbf{else} \ St_2 \ \mathbf{end}; St, \sigma, alias_{old} \rangle}{\langle p :: St, \sigma, alias_{new} \rangle} \rightarrow \quad (42)$$

Observe that the SCOOP system configurations in (1) are enriched with a new component `alias_` consisting of a set of alias expressions. Above, `aliasold` is the aliasing before the execution of the “**if**” instruction, whereas, intuitively, `aliasnew` stands for `aliasold » St1 ∪ aliasold » St2`.

As a second step, we analyzed the implementation in [1] from a more engineering perspective, and identified a series of design decisions that either slowed down considerably the rewriting or made the search space grow unnecessarily large.

After running some experiments, we understood that the parallelism rule

$$\frac{\langle p_1 :: St_1, \sigma \rangle \rightarrow \langle p'_1 :: St'_1, \sigma' \rangle}{\langle p_1 :: St_1 \mid p_2 :: St_2, \sigma \rangle \rightarrow \langle p'_1 :: St'_1 \mid p_2 :: St_2, \sigma' \rangle} \quad (43)$$

in [1] was increasing the rewriting time. Though elegant from the formalization perspective, the usage of this rule was not efficient. Therefore, we eliminated it from the SCOOP semantics and made the remaining rules apply directly, by matching at top. For instance, the abstract rule (42) formalizing “**if**” instructions in the context of one processor `p` becomes:

$$\frac{\langle p :: \mathbf{if} \ e \ \mathbf{then} \ St_1 \ \mathbf{else} \ St_2 \ \mathbf{end}; St \mid aqs, \sigma, alias_{old} \rangle}{\langle p :: St \mid aqs, \sigma, alias_{new} \rangle} \rightarrow \quad (44)$$

for an arbitrary list `aqs` of processors and their instruction stacks. For Dining Philosophers, for example, this modification reduced the rewriting time from around 10s to less than 1s.

Recall that SCOOP processors communicate via channels. The implementation in [1] creates *fresh* channels (as in (2), for instance) parameterized by natural indexes. This was inefficient for model-checking purposes, as the state space contained many identical states up-to channel naming.

The implementation of the above observations enabled us to successfully identify a deadlock situation in a Dining Philosophers scenario, by using the Maude LTL model-checker. The new (reduced) formalization of SCOOP can be found at:

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



efficiency and the precision (*e.g.*, the number of false positives introduced by the alias approximations) between our approach and other aliasing techniques.

Another research direction is to derive alias-based abstractions for analyzing concurrent programs. We foresee possible connections with the work in [37] 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 [18] can be exploited in order to derive finite 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 [38] that handles point-to graphs via a stack-based algorithm for fixed-point computations.

Related to deadlock detection, the second contribution of this paper, we provided a formalization based on the sets of acquired resources and the sets of handlers processors wait to lock in their attempt to execute asynchronously. This definition corresponds to Coffman deadlocks [31] in the context of SCOOP, occurring whenever there is a set of processors reserving each other circularly. We presented the equivalent SCOOP semantic rule and discussed the results of using Maude in order to analyze deadlocks in the context of a Dining Philosophers scenario. In addition, we presented how an abstract semantics based on aliases, together with a series of implementation improvements enabled the Maude LTL model-checker to correctly identify deadlocks. A survey on abstracting techniques on top of Maude executable semantics is provided in [4].

The literature on using static analysis [7] and abstracting techniques for (related) concurrency models is considerable. We refer, for instance, to the recent work in [39] that introduces a framework for detecting deadlocks by identifying circular dependencies in the (finite state) model of so-called contracts that abstract methods in an OO-language. Nevertheless, the integration of a deadlock analyzer in SCOOP on top of Maude is an orthogonal approach that aims at constructing a RL-based toolbox for SCOOP programs laying over the same semantic framework.

In [40] SCOOP programs are verified for deadlocks and other behavioral properties using GROOVE [41]. The work in [40] proposes a redefinition of the most common features of the SCOOP semantics based on graph transformation systems (GPSs). This is a bottom-up approach, as it aims at redefining the SCOOP semantics from scratch via GPSs, orthogonal to our rather top-down strategy of narrowing the original semantics proposed in [1].

As a clear direction for future work we consider designing and analyzing deadlock situations for more SCOOP applications. Based on the experience so far, this would help better understand and observe the SCOOP state space, thus providing hints for further improvements in the context of model-checking, for instance. As we could already see, major advances in this regard are obtained when semantically equivalent states are identified and collapsed within the same equivalence classes. This was the case of the indexed-based communication channels in Section 6.3. Similar redundant states are introduced by the so-called “regions” in SCOOP “administrating” objects handled by the same processor. We foresee their elimination would speed-up the LTL model-checker.

*Acknowledgements.* We are grateful for valuable comments to the anonymous reviewers of FTSCS’14, Măriuca Asăvoae, Alexander Kogtenkov, José Meseguer, Benjamin Morandi and Sergej 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.

- [1] B. Morandi, M. Schill, S. Nanz, B. Meyer, Prototyping a concurrency model, in: ACSD, 2013, pp. 170–179.
- [2] B. Meyer, Eiffel: The Language, Prentice-Hall, 1991.
- [3] A. Rusakov, J. Shin, B. Meyer, Simple concurrency for robotics with the Roboscoop framework, in: 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, Chicago, IL, USA, September 14-18, 2014, 2014, pp. 1563–1569. doi:10.1109/IROS.2014.6942763.  
URL <http://dx.doi.org/10.1109/IROS.2014.6942763>
- [4] J. Meseguer, 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, 2011, pp. 1–37.
- [5] M. Clavel, F. Durán, S. Eker, P. Lincoln, N. Martí-Oliet, J. Meseguer, C. L. Talcott (Eds.), All About Maude - A High-Performance Logical Framework, How to Specify, Program and Verify Systems in Rewriting Logic, Vol. 4350 of Lecture Notes in Computer Science, Springer, 2007.
- [6] W. Landi, 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, ACM, New York, NY, USA, 1991, pp. 93–103. doi:10.1145/99583.99599.  
URL <http://doi.acm.org/10.1145/99583.99599>
- [7] W. Landi, Undecidability of static analysis, ACM Lett. Program. Lang. Syst. 1 (4) (1992) 323–337. doi:10.1145/161494.161501.  
URL <http://doi.acm.org/10.1145/161494.161501>
- [8] 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, ACM, New York, NY, USA, 1981, pp. 219–230. doi:10.1145/567532.567556.  
URL <http://doi.acm.org/10.1145/567532.567556>
- [9] M. Hind, M. Burke, P. Carini, J.-D. Choi, Interprocedural pointer alias analysis, ACM Trans. Program. Lang. Syst. 21 (4) (1999) 848–894. doi:10.1145/325478.325519.  
URL <http://doi.acm.org/10.1145/325478.325519>
- [10] A. Diwan, K. S. McKinley, J. E. B. Moss, Type-based alias analysis, SIGPLAN Not. 33 (5) (1998) 106–117. doi:10.1145/277652.277670.  
URL <http://doi.acm.org/10.1145/277652.277670>
- [11] M. Burke, P. Carini, J.-D. Choi, M. Hind, Flow-insensitive interprocedural alias analysis in the presence of pointers, in: K. Pingali, U. Banerjee, D. Gelernter, A. Nicolau, D. Padua (Eds.), Languages and Compilers for Parallel Computing, Vol. 892 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 1995, pp. 234–250. doi:10.1007/BFb0025882.  
URL <http://dx.doi.org/10.1007/BFb0025882>
- [12] J.-D. Choi, M. Burke, 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, ACM, New York, NY, USA, 1993, pp. 232–245. doi:10.1145/158511.158639.  
URL <http://doi.acm.org/10.1145/158511.158639>
- [13] M. Emami, R. Ghiya, 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, ACM, New York, NY, USA, 1994, pp. 242–256. doi:10.1145/178243.178264.  
URL <http://doi.acm.org/10.1145/178243.178264>
- [14] R. P. Wilson, 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, ACM, New York, NY, USA, 1995, pp. 1–12. doi:10.1145/207110.207111.  
URL <http://doi.acm.org/10.1145/207110.207111>
- [15] 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, ACM, New York, NY, USA, 2006, pp. 54–63. doi:10.1145/1134650.1134659.  
URL <http://doi.acm.org/10.1145/1134650.1134659>
- [16] E. Albert, P. Arenas, S. Genaim, G. Puebla, Field-sensitive value analysis by field-insensitive analysis, in: Proceedings of the 2Nd World Congress on Formal Methods, FM '09, Springer-Verlag, Berlin, Heidelberg, 2009, pp. 370–386.
- [17] M. Hind, Pointer analysis: haven't we solved this problem yet?, in: PASTE, 2001, pp. 54–61.
- [18] A. Kogtenkov, B. Meyer, S. Velder, Alias and change calculi, applied to frame inference, CoRR abs/1307.3189.  
URL <http://arxiv.org/abs/1307.3189>
- [19] J. R. Larus, P. N. Hilfinger, Detecting conflicts between structure accesses, in: PLDI, 1988, pp. 21–34.
- [20] C. Shih, J. A. Stankovic, Survey of deadlock detection in distributed concurrent programming environments and its application to real-time systems, Tech. rep., Amherst, MA, USA (1990).
- [21] G. R. Andrews, G. M. Levin, On-the-fly deadlock prevention, in: Proceedings of the First ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing, PODC '82, ACM, New York, NY, USA, 1982, pp. 165–172.
- [22] T. Minoura, Deadlock avoidance revisited, J. ACM 29 (4) (1982) 1023–1048.
- [23] K. Chandy, J. Misra, L. Haas, T. U. A. A. D. O. C. SCIENCES., A Distributed Deadlock Detection Algorithm and Its Correctness Proof, Defense Technical Information Center, 1982.  
URL <http://books.google.ch/books?id=AH-3NwAACAAJ>
- [24] D. Badal, M. Gehl, M. Gehl, N. P. S. M. CA., N. P. S. (U.S.), On Deadlock Detection in Distributed Computing Systems, Defense Technical Information Center, 1983.  
URL <http://books.google.ch/books?id=JjrQ0AAACAAJ>

- [25] G. Caltais, Expression-based aliasing for OO-languages, CoRR abs/1409.7509, accepted in FTSCS'14.  
URL <http://arxiv.org/abs/1409.7509>
- [26] M. O. Rabin, D. Scott, Finite automata and their decision problems, IBM J. Res. Dev. 3 (2) (1959) 114–125.  
doi:10.1147/rd.32.0114.  
URL <http://dx.doi.org/10.1147/rd.32.0114>
- [27] G. Rosu, T. F. Serbanuta, K overview and SIMPLE case study, in: Proceedings of International K Workshop (K'11), ENTCS, Elsevier, 2013, To appear.
- [28] S. C. Kleene, Representation of events in nerve nets and finite automata, in: C. Shannon, J. McCarthy (Eds.), Automata Studies, Princeton University Press, Princeton, NJ, 1956, pp. 3–41.
- [29] T.-F. Serbanuta, G. Rosu, K-Maude: A rewriting based tool for semantics of programming languages, in: WRLA, 2010, pp. 104–122.
- [30] P. Nienaltowski, Practical Framework for Contract-based Concurrent Object-oriented Programming, ETH, 2007.  
URL <http://books.google.ch/books?id=ZDcEkgAACAAJ>
- [31] E. G. Coffman, M. Elphick, A. Shoshani, System deadlocks, ACM Comput. Surv. 3 (2) (1971) 67–78.
- [32] N. Martí-Oliet, J. Meseguer, A. Verdejo, Towards a Strategy Language for Maude, Electronic Notes in Theoretical Computer Science 117 (2005) 417–441.
- [33] S. Eker, J. Meseguer, A. Sridharanarayanan, The maude LTL model checker, Electr. Notes Theor. Comput. Sci. 71 (2002) 162–187.
- [34] I. M. Asavoaie, Abstract semantics for alias analysis in K, Electr. Notes Theor. Comput. Sci. 304 (2014) 97–110.  
doi:10.1016/j.entcs.2014.05.005.  
URL <http://dx.doi.org/10.1016/j.entcs.2014.05.005>
- [35] A. Bouajjani, J. Esparza, O. Maler, Reachability analysis of pushdown automata: Application to model-checking, in: CONCUR, 1997, pp. 135–150.
- [36] P. Cousot, R. Cousot, Abstract interpretation and application to logic programs, J. Log. Program. 13 (2&3) (1992) 103–179.
- [37] C. A. R. Hoare, B. Möller, G. Struth, I. Wehrman, Concurrent Kleene algebra, in: CONCUR 2009 - Concurrency Theory, 20th International Conference, CONCUR 2009, Bologna, Italy, September 1-4, 2009. Proceedings, 2009, pp. 399–414.
- [38] D. R. Chase, M. N. Wegman, F. K. Zadeck, Analysis of pointers and structures, in: PLDI, 1990, pp. 296–310.
- [39] E. Giachino, C. A. Grazia, C. Laneve, M. Lienhardt, P. Y. H. Wong, Deadlock analysis of concurrent objects: Theory and practice, in: Integrated Formal Methods, 10th International Conference, IFM 2013, Turku, Finland, June 10-14, 2013. Proceedings, 2013, pp. 394–411.
- [40] A. Heußner, C. M. Poskitt, C. Corrodi, B. Morandi, Towards practical graph-based verification for an object-oriented concurrency model, in: Proc. Graphs as Models (GaM 2015), Vol. 181 of Electronic Proceedings in Theoretical Computer Science, 2015, pp. 32–47.
- [41] A. H. Ghamarian, M. de Mol, A. Rensink, E. Zambon, M. Zimakova, Modelling and analysis using groove, Int. J. Softw. Tools Technol. Transf. 14 (1) (2012) 15–40.