



HAL
open science

A lightweight deadlock analysis for programs with threads and reentrant locks

Cosimo Laneve

► **To cite this version:**

Cosimo Laneve. A lightweight deadlock analysis for programs with threads and reentrant locks. Science of Computer Programming, 2019, 181, pp.64 - 81. 10.1016/j.scico.2019.06.002 . hal-02392938

HAL Id: hal-02392938

<https://hal.inria.fr/hal-02392938>

Submitted on 4 Dec 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



A lightweight deadlock analysis for programs with threads and reentrant locks



Cosimo Laneve¹

Department of Computer Science and Engineering, University of Bologna – INRIA Focus, Italy

ARTICLE INFO

Article history:

Received 5 October 2018

Received in revised form 11 June 2019

Accepted 14 June 2019

Available online 21 June 2019

Keywords:

Deadlock analysis

Threads and reentrant locks

Lams

Circularities

Static semantics

ABSTRACT

Deadlock analysis of multi-threaded programs with reentrant locks is complex because these programs may have infinitely many states. We define a simple calculus featuring recursion, threads and synchronizations that guarantee exclusive access to objects. We detect deadlocks by associating an abstract model to programs – the *extended lam* model – and we define an algorithm for verifying that a problematic object dependency (e.g. a *circularity*) between threads will not be manifested.

The analysis is lightweight because the deadlock detection problem is fully reduced to the corresponding one in lams (without using other models). In fact, the technique is intended to be an effective tool for the deadlock analysis of programming languages by defining ad-hoc extraction processes. We demonstrate this effectivity by applying our analysis to a core calculus featuring shared objects, threads and Java-like synchronization primitives. We also discuss a prototype verifier, called JaDA, that covers several features of Java and deliver initial assessments of the tool.

© 2019 Elsevier B.V. All rights reserved.

1. Introduction

Threads and locks are a common model of concurrent programming that is nowadays widely used by the mainstream programming languages (Java, C#, C++, Objective C, etc.). In these languages, deadlocks are flaws that occur when two or more threads are blocked because each one is attempting to acquire an object's lock held by another one. As an example, consider the following method

```
setTable(C x, C y, int n) = (new C z) (
  if (n=0) then sync(y){ sync(x){ 0 } }
  else fork{ sync(x){ sync(z){ 0 } } }
  z.setTable(z, y, n-1)
)
```

where `new C z` creates a new object of class `C`, `fork{P} Q` creates a new thread whose body is `P` and runs it in parallel with the continuation `Q`, and `sync(x){ P }` is the operation locking the object `x`, performing `P`, and releasing `x` when `P` terminates. The method `setTable` creates a table of $n + 1$ threads – the *philosophers* – each one sharing an object – the *fork* – with the nearby philosopher. Every philosopher, except one, grabs the fork on his left – the first argument – and on his right – the second argument – in this order and then releases them. The exceptional case is the `then`-branch ($n=0$) where the grabbing strategy is opposite. It is well-known that, when the method is invoked with `x.setTable(x, x, n)`,

E-mail address: cosimo.laneve@unibo.it.

¹ Research partly supported by the H2020-MSCA-RISE project ID 778233 “Behavioural Application Program Interfaces (BEHAPI)”.

no deadlock will ever occur because at least one philosopher has a strategy that is different from the other ones. On the contrary, if we change the `then`-branch into `sync(this){ sync(y){ 0 } }` a deadlock may occur because philosophers' strategies are all symmetric. It is worth to notice that `x.setTable(x,x,0)` is deadlock-free because it just locks twice the object `x`, which is admitted in the model with threads and locks – a thread may acquire a same lock several times (*lock-reentrancy*).

In order to ensure termination, current analyzers [4,7,1,16,27,28] use finite approximate models representing the dependencies between object names. The corresponding algorithms usually return false positives with input `x.setTable(x,x,n)` because they are not powerful enough to manage structures that are not statically bounded.

In [11,17] we solved this problem for value-passing CCS [21] and pi-calculus [22]. In that case, the technique used two formal models: Petri Nets [25] and *deadLock Analysis Models – lams* [12]. The former ones are a well-known model of concurrent and distributed systems; the latter ones are basic recursive models that collect dependencies and features recursion and dynamic name creation. In the pi-calculus analyzer, Petri Nets were used to verify the consistency of the communication protocol of every channel, while lams were used for guaranteeing the correctness of the dependencies between different channels. The analysis algorithm of [11,17] required a tool for verifying the reachability of Petri Nets that model channels' behaviors (which has exponential computational complexity with respect to the size of the net [15]) and a tool for analyzing lams (which has exponential computational complexity with respect to the number of arguments of functions).

In this paper we demonstrate that it is possible to define a deadlock analyzer for programs with threads and (reentrant) locks by only using an extension of lams. For example, the lam function corresponding to `setTable` is²

$$\text{setTable}(t, x, y) = (v s, z) \left((y, x)_t + ((x, z)_s \& \text{setTable}(t, z, y)) \right).$$

The term $(y, x)_t$, called *dependency*, indicates that the thread t , which owns the lock of y , is going to grab the lock of x . The operation “+” and “&” are disjunction and conjunctions of dependencies, respectively. The index t of $(y, x)_t$ was missing in [12,11,17]; it has been necessary for modeling reentrant locks. In particular, (x, x) is a circularity in the standard lam model, whilst $(x, x)_t$ is not a circularity in the *extended lam model* because t can acquire x twice. Therefore, `setTable(t, x, x)` manifests a circularity in the model of [12,11,17] and it does not in the extended model. A circularity in the extended model is $(y, x)_t \& (x, y)_s$, which denotes that two *different* threads are attempting to acquire two objects in different order. This lam gives $(x, x)_\checkmark$, which represents a circularity.

Because of the foregoing extension, the algorithm for detecting circularities in extended lams is different than the one in [11,17]. In particular, while there is a *decision algorithm* for the presence/absence of circularities in standard lams, in Section 2 we define an algorithm that verifies the absence and is imprecise in some cases (it may return that a lam will manifest a circularity while it will not be the case – a false positive).

We also define a simple language of concurrent programs featuring recursion, threads, shared objects and synchronization primitives inspired to those of Java. (The foregoing method `setTable` is a term of our calculus.) The syntax, semantics, and examples of the calculus are in Section 3. Section 4 reports a type system that associates lams to processes. For example, the type system returns the lam function `setTable` for the method `setTable`. As a byproduct of the type system and the algorithm for detecting circularities in lams, our technique can detect deadlocks of programs like `setTable`.

In order to deliver initial assessments of our technique, we have extended the basic language in Section 3 to cover several features of Java. In Section 5 we focus on inheritance that allows classes to be defined as sub-classes of other ones by deriving the corresponding method (and field) definitions and eventually redefine them (method overriding). In a language with inheritance, locating a method definition is performed at run-time through a mechanism called *dynamic dispatch* that uses the actual value of the called object. Since our technique extracts lams from programs *at static time*, the precise modeling of dynamic dispatch is not feasible. We therefore over-approximate the type by returning the lam functions of every possible overridden method in the subclasses. In Section 6 we also analyze assignments and recursive data-types. The analysis of other features of Java can be found in Garcia's PhD thesis [9]. This analysis does not address coordination primitives (methods `wait`, `notify` and `notifyAll`) and Java concurrent libraries, such as `java.util.concurrent`, – see Section 9 for a discussion.

Our theoretical developments are complemented by an analyzer prototype called JaDA [10,19]. While the type system in this paper simply checks static information, JaDA infers the behavioral types from the code. Inference is important in practice because it lightens the analysis but checking is crucial for type safety. Section 7 reports on initial assessments of the tool.

We discuss a few weaknesses of the techniques in Section 8 and we point to related works and give some concluding remarks in Section 9. The paper also includes an appendix that contains the demonstration of type safety for the language of Section 3 and the type system of Section 4.

This paper collects two conference contributions about deadlock analysis of Java-like languages: [18] and [19]. Here we extend [18] with the proof of type soundness and with the discussion on the extensions of the basic calculus in Section 3.

² Actually, this is a simplified form of `setTable`. The lam function associated to `setTable` by the type system in Section 4 has two additional names that record the last name synchronized by the thread t and the carrier of the invocation.

With respect to [19], we adopt a different presentation style. We discuss typing rules for the Java language, while in the conference paper the focus was on Java bytecode.

2. Lams and the algorithm for detecting circularities

This section extends the theory developed in [11,17] to cover thread reentrancy. In particular, the new definitions are those of transitive closure and Definition 2.3. Theorem 2.8 is new.

Preliminaries. We use an infinite set \mathcal{A} of names, ranged over by x, y, t, s, \dots . A relation on \mathcal{A} , denoted $\mathbb{R}, \mathbb{R}', \dots$, is an element of $\mathcal{P}(\mathcal{A} \times \mathcal{A} \times (\mathcal{A} \cup \{\checkmark, \bullet\}))$, where $\mathcal{P}(\cdot)$ is the standard powerset operator, $\cdot \times \cdot$ is the cartesian product, and $\checkmark, \bullet \notin \mathcal{A}$ are two special names. The elements of \mathbb{R} , called *dependencies*, are denoted by $(x, y)_t$, where t is called *thread*. The name \checkmark indicates that the dependency is due to the contributions of two or more threads; \bullet indicates that the dependency is due to a thread whose name is unknown.

Let

- \mathbb{R}^+ be the least relation containing \mathbb{R} and closed under the operations:
 1. if $(x, y)_t, (y, z)_{t'} \in \mathbb{R}^+$ and $t \neq t'$ then $(x, z)_{\checkmark} \in \mathbb{R}^+$;
 2. if $(x, y)_t, (y, z)_t \in \mathbb{R}^+$, $t \in \mathcal{A} \cup \{\checkmark\}$, then $(x, z)_t \in \mathbb{R}^+$;
 3. if $(x, y)_{\bullet}, (y, z)_{\bullet} \in \mathbb{R}^+$, $(x, y)_{\bullet} \neq (y, z)_{\bullet}$, then $(x, z)_{\checkmark} \in \mathbb{R}^+$.
- $\{\mathbb{R}_1, \dots, \mathbb{R}_m\} \subseteq \{\mathbb{R}'_1, \dots, \mathbb{R}'_n\}$ if and only if, for all \mathbb{R}_i , there is \mathbb{R}'_j such that
 1. if $(x, y)_t \in \mathbb{R}_i$, $t \in \mathcal{A}$, then $(x, y)_t \in \mathbb{R}'_j$;
 2. if $(x, y)_{\checkmark} \in \mathbb{R}_i$ then either $(x, y)_{\checkmark} \in \mathbb{R}'_j$ or there is $t \in \mathcal{A}$ such that $(x, y)_t \in \mathbb{R}'_j$;
 3. if $(x, y)_{\bullet} \in \mathbb{R}_i$ then either $(x, y)_{\bullet} \in \mathbb{R}'_j$ or there is $t \in \mathcal{A}$ such that $(x, y)_t \in \mathbb{R}'_j$.
- $\{\mathbb{R}_1, \dots, \mathbb{R}_m\} \& \{\mathbb{R}'_1, \dots, \mathbb{R}'_n\} \stackrel{\text{def}}{=} \{\mathbb{R}_i \cup \mathbb{R}'_j \mid 1 \leq i \leq m \text{ and } 1 \leq j \leq n\}$.

We use $\mathcal{R}, \mathcal{R}', \dots$ to range over $\{\mathbb{R}_1, \dots, \mathbb{R}_m\}$, which are elements of $\mathcal{P}(\mathcal{P}(\mathcal{A} \times \mathcal{A} \times (\mathcal{A} \cup \{\checkmark, \bullet\})))$.

The names \checkmark and \bullet are managed in an ad-hoc way in the transitive closure \mathbb{R}^+ and in the relation \subseteq . In particular, if $(x, y)_t$ and $(y, z)_{t'}$ belong to a relation and $t \neq t'$, the dependency obtained by transitivity, e.g. $(x, z)_{\checkmark}$, records that it has been produced by a contribution of two different threads – this is important for separating circularities, e.g. $(x, x)_{\checkmark}$, from lock reentrancy, e.g. $(x, x)_t$. The name \bullet copes with another issue: it allows us to abstract away thread names that are created inside methods. For this reason the transitive dependency of $(x, y)_{\bullet}$ and $(y, z)_{\bullet}$ is $(x, z)_{\checkmark}$ because the threads producing $(x, y)_{\bullet}$ and $(y, z)_{\bullet}$ might be different. The meaning of $\mathcal{R} \subseteq \mathcal{R}'$ is that \mathcal{R}' is “more precise” with respect to pairs (x, y) : if this pair is indexed with either \checkmark or \bullet in some $\mathbb{R} \in \mathcal{R}$ then it may be indexed by a t ($t \neq \checkmark$) in the corresponding (transitive closure) relation of \mathcal{R}' . For example $\{(x, y)_{\bullet}, (y, z)_{\bullet}, (x, z)_{\checkmark}\} \subseteq \{(x, y)_t, (y, z)_t\}$ and $\{(x, x)_{\bullet}\} \subseteq \{(x, x)_t, (x, x)_{t'}\}$.

Definition 2.1. A relation \mathbb{R} has a circularity if $(x, x)_{\checkmark} \in \mathbb{R}^+$ for some x . A set of relations \mathcal{R} has a circularity if there is $\mathbb{R} \in \mathcal{R}$ that has a circularity.

Lams. In our technique, dependencies are expressed by means of *lams* [12], noted ℓ , whose syntax is

$$\ell ::= 0 \mid (x, y)_t \mid (\nu x) \ell \mid \ell \& \ell' \mid \ell + \ell' \mid \mathfrak{f}(\bar{x})$$

The term 0 is the empty type; $(x, y)_t$ specifies a dependency between the name x and the name y that has been created by (the thread) t . The operation $(\nu x) \ell$ creates a new name x whose scope is the type ℓ ; the operations $\ell \& \ell'$ and $\ell + \ell'$ define the conjunction and disjunction of the dependencies in ℓ and ℓ' , respectively. The operators $+$ and $\&$ are associative and commutative; the operator $\&$ has precedence over $+$. The term $\mathfrak{f}(\bar{x})$ defines the invocation of \mathfrak{f} with arguments \bar{x} . The argument sequence \bar{x} has always at least two elements in our case: the first element is the thread that performed the invocation, the second element is the last object whose lock has been acquired by it.

A *lam program* is a pair (\mathcal{L}, ℓ) , where \mathcal{L} is a finite set of function definitions

$$\mathfrak{f}(\bar{x}) = \ell_{\mathfrak{f}}$$

with $\ell_{\mathfrak{f}}$ being the *body* of \mathfrak{f} , and ℓ is the *main lam*. We always assume that (i) \bar{x} is a sequence of pairwise different names and (ii) $\ell_{\mathfrak{f}} = (\nu \bar{z}) \ell'_{\mathfrak{f}}$ where $\ell'_{\mathfrak{f}}$ has no ν -binder. The constraint (ii) also applies to ℓ and, whenever we write $(\mathcal{L}, (\nu \bar{x}) \ell')$ we always mean that ℓ' has no ν -binder. The function `setTable` in the Introduction is an example of a lam function.

Let a *lam context*, noted $\mathbb{L}[\]$, be a term derived by the following syntax:

$$\mathbb{L}[\] ::= [\] \mid \ell \& \mathbb{L}[\] \mid \ell + \mathbb{L}[\]$$

As usual $\mathbb{L}[\ell]$ is the lam where the hole of $\mathbb{L}[\]$ is replaced by ℓ . We remark that, according to the syntax, lam contexts have no ν -binder.

We notice that $I_{\mathcal{L}}^{(0)}$ is the function I^\perp . Let us analyze the definition of $I_{\mathcal{L}}^{(k+1)}(\mathfrak{f}_i)$ and, in particular, the function proj :

- first of all, notice that proj applies to the transitive closures of relations, which may have names in $A, \bar{z}_i, \checkmark, \bullet$ and \varkappa ;
- the transitive closure operation is crucial because a circularity may follow with the key contribution of fresh names. For instance the model of $\mathfrak{f}(x) = (\nu t, t', z) (x, z)_t \& (z, x)_{t'}$ is $\{\{(x, x)_{\checkmark}\}\}$; the model of $\mathfrak{g}() = (\nu t, t', x, y) (x, y)_t \& (y, x)_{t'}$ is $\{\{(x, x)_{\checkmark}\}\}$ (this is the reason why we use the name \varkappa);
- every dependency $(u, v)_t \in \text{proj}_{\bar{x}}^{\checkmark}(\mathbb{R})$ is such that $u, v \in \bar{x}$, except for $(x, x)_{\checkmark}$. For example, if $\mathfrak{f}'(x, y) = (\nu s, z) ((x, y)_s \& (x, z)_s)$ then, if we invoke $\mathfrak{f}'(u, v)$ we obtain $(u, v)_{t'} \& (u, z')_{t'}$, where t' and z' are fresh object names. This lam may be simplified because, being z' fresh and unknown elsewhere, the dependency $(u, z')_{t'}$ will never be involved in a circularity. For example, if we have $\ell = (\nu, u)_t \& \mathfrak{f}'(u, v)$ then we may safely reason on ℓ' -simplified $(u, v)_t \& (u, v)_{t'}$. For this reason we drop the dependencies containing fresh names *after their contribution to the transitive closure has been computed*;
- the same argument does not apply to names used as threads. For example, in the above ℓ' -simplified lam we cannot drop $(u, v)_{t'}$ because t' is fresh. In fact, the context $(\nu, u)_t \& (u, v)_{t'}$ gives a circularity. Therefore, dependencies whose thread names are fresh must be handled in a different way. We take a simple solution: these dependencies all have \bullet as thread name. That is, we assume that they are all generated by the contribution of different threads. For example, $\mathfrak{g}'(x, y) = (\nu t) (x, y)_t$. Then, $I_{\mathcal{L}}^{(1)}(\mathfrak{g}') = \{\{(x, y)_{\bullet}\}\}$.

Example 2.4. The flattening functions of `setTable` are

$$\begin{aligned} I_{\mathcal{L}}^{(0)}(\text{setTable}) &= \{\emptyset\} \\ I_{\mathcal{L}}^{(1)}(\text{setTable}) &= \{\{(y, x)_t\}\} \end{aligned}$$

As another example, consider the function $\mathfrak{g}(x, y, z) = (\nu t, u) (x, y)_t \& \mathfrak{g}(y, z, u)$. Then:

$$\begin{aligned} I_{\mathcal{L}}^{(0)}(\mathfrak{g}) &= \{\emptyset\} \\ I_{\mathcal{L}}^{(1)}(\mathfrak{g}) &= \{\{(x, y)_{\bullet}\}\} \\ I_{\mathcal{L}}^{(2)}(\mathfrak{g}) &= \{\{(x, y)_{\bullet}, (y, z)_{\bullet}, (x, z)_{\checkmark}\}\} \end{aligned}$$

Proposition 2.5. Let $\mathfrak{f}(\bar{x}) = (\nu \bar{z}) \ell_{\mathfrak{f}} \in \mathcal{L}$.

1. For every k , $I_{\mathcal{L}}^{(k)}(\mathfrak{f}) \in \mathcal{P}(\mathcal{P}((\bar{x} \cup \{x\}) \times (\bar{x} \cup \{x\}) \times (\bar{x} \cup \{\checkmark, \bullet\})))$.
2. For every k , $I_{\mathcal{L}}^{(k)}(\mathfrak{f}) \in I_{\mathcal{L}}^{(k+1)}(\mathfrak{f})$.

Proof. (1) follows by definition. As regards (2), we observe that $I(\ell)$ is monotonic on I . That is, if I and I' are two flattening functions such that, for every \mathfrak{f} , $I(\mathfrak{f}) \in I'(\mathfrak{f})$ then it is possible to demonstrate by a standard structural induction that $I(\ell) \in I'(\ell)$. An induction on k gives $I_{\mathcal{L}}^{(k)}(\mathfrak{f}) \in I_{\mathcal{L}}^{(k+1)}(\mathfrak{f})$. \square

Since, for every k , $I_{\mathcal{L}}^{(k)}(\mathfrak{f}_i)$ ranges over a finite lattice, by the fixpoint theory [5], there exists m such that $I_{\mathcal{L}}^{(m)}$ is a fixpoint, namely $I_{\mathcal{L}}^{(m)} \approx I_{\mathcal{L}}^{(m+1)}$ where \approx is the equivalence relation induced by \in . In the following, we let $I_{\mathcal{L}}$, called the *interpretation function* (of a lam), be the least fixpoint $I_{\mathcal{L}}^{(m)}$. In Example 2.4, $I_{\mathcal{L}}^{(1)}$ is the fixpoint of `setTable` and $I_{\mathcal{L}}^{(2)}$ is the fixpoint of \mathfrak{g} .

Proposition 2.6. Let \mathbb{L} be a lam context, ℓ be a lam with no ν binder, and $I(\cdot)$ be a flattening. Then we have:

1. $I(\mathbb{L}[\ell])$ has a circularity if and only if $I(\mathbb{L}[\mathbb{R}])$ has a circularity for some $\mathbb{R} \in I(\ell)$.
2. Let $\mathfrak{f}(\bar{x}) = (\nu \bar{z}) \ell_{\mathfrak{f}} \in \mathcal{L}$ and $\mathbb{R} \in I(\ell_{\mathfrak{f}}\{\bar{z}/\bar{z}\})$ with \bar{z}' fresh. If $I(\mathbb{L}[\mathbb{R}\{\bar{u}/\bar{x}\}])$ has a circularity then $I(\mathbb{L}[(\text{proj}_{\bar{x}}^{\checkmark}(\mathbb{R}^+))\{\bar{u}/\bar{x}\}])$ has a circularity.

Proof. Property 1 follows from the definitions. To see 2, we use a straightforward induction on \mathbb{L} . We analyze the basic case $\mathbb{L} = []$: the general case follows by induction. Let $\mathbb{R} \in I(\ell_{\mathfrak{f}}\{\bar{z}/\bar{z}\})$ such that $\mathbb{R}\{\bar{u}/\bar{x}\}$ has a circularity. There are two cases:

- $(\nu, \nu)_{\checkmark} \in \mathbb{R}^+$. By definition of $\text{proj}_{\bar{x}}^{\checkmark}(\mathbb{R}^+)$ either $(\nu, \nu)_{\checkmark} \in \text{proj}_{\bar{x}}^{\checkmark}(\mathbb{R}^+)$, when $\nu \notin \bar{z}'$, or $(x, x)_{\checkmark} \in \text{proj}_{\bar{x}}^{\checkmark}(\mathbb{R}^+)$, otherwise. In this case the statement 2 follows immediately.
- $(\nu, \nu)_{\checkmark} \in \mathbb{R}\{\bar{u}/\bar{x}\}^+$. By definition of transitive closure $\mathbb{R}\{\bar{u}/\bar{x}\}^+ = (\mathbb{R}^+)\{\bar{u}/\bar{x}\}$. Then there is a dependency $(x_1, x_2)_{\checkmark} \in \mathbb{R}^+$ such that $(x_1, x_2)_{\checkmark}\{\bar{u}/\bar{x}\} = (\nu, \nu)_{\checkmark}$. By definition of proj , $(x_1, x_2)_{\checkmark} \in \text{proj}_{\bar{x}}^{\checkmark}(\mathbb{R}^+)$. Therefore $\text{proj}_{\bar{x}}^{\checkmark}(\mathbb{R}^+)\{\bar{u}/\bar{x}\}$ has also a circularity. \square

Lemma 2.7. Let $\mathcal{L} = \{\mathfrak{f}_1(\bar{x}_1) = (\nu \bar{z}_1) \ell_1, \dots, \mathfrak{f}_n(\bar{x}_n) = (\nu \bar{z}_n) \ell_n\}$ and let

$$\mathbb{L}[\mathfrak{f}_{i_1}(\bar{u}_1)] \cdots [\mathfrak{f}_{i_m}(\bar{u}_m)] \longrightarrow^m \mathbb{L}[\ell_{i_1}\{\bar{z}'_1/\bar{z}_{i_1}\}\{\bar{u}_1/\bar{x}_{i_1}\}] \cdots [\ell_{i_m}\{\bar{z}'_m/\bar{z}_{i_m}\}\{\bar{u}_m/\bar{x}_{i_m}\}]$$

where $\mathbb{L}[\cdot] \cdots [\cdot]$ is a multiple context without function invocations.

If $I_{\mathcal{L}}^{(k)}(\mathbb{L}[\ell_{i_1}\{\bar{z}'_1/\bar{z}_{i_1}\}\{\bar{u}_1/\bar{x}_{i_1}\}] \cdots [\ell_{i_m}\{\bar{z}'_m/\bar{z}_{i_m}\}\{\bar{u}_m/\bar{x}_{i_m}\}])$ has a circularity then $I_{\mathcal{L}}^{(k+1)}(\mathbb{L}[\mathfrak{f}_{i_1}(\bar{u}_1)] \cdots [\mathfrak{f}_{i_m}(\bar{u}_m)])$ has also a circularity.

Proof. To show the implication suppose that

$$I_{\mathcal{L}}^{(k)}(\mathbb{L}[\ell_{i_1}\{\bar{z}'_1/\bar{z}_{i_1}\}\{\bar{u}_1/\bar{x}_{i_1}\}] \cdots [\ell_{i_m}\{\bar{z}'_m/\bar{z}_{i_m}\}\{\bar{u}_m/\bar{x}_{i_m}\}])$$

has a circularity. By repeated applications of Proposition 2.6(1), there exists $R_j \in I_{\mathcal{L}}^{(k)}(\ell_{i_1}\{\bar{z}'_j/\bar{z}_{i_1}\}\{\bar{u}_j/\bar{x}_{i_1}\})$ with $1 \leq j \leq m$ such that $I_{\mathcal{L}}^{(k)}(\mathbb{L}[R_1] \cdots [R_m])$ has a circularity. It is easy to verify that every R_j may be written as $R'_j\{\bar{u}_j/\bar{x}_{i_j}\}$, for some R'_j . By repeated application of Proposition 2.6(2), we have that

$$I_{\mathcal{L}}^{(k)}(\mathbb{L}[\text{proj}_{\bar{x}_{i_1}}^{\bar{z}'_1}(R'_1)\{\bar{u}_1/\bar{x}_{i_1}\}] \cdots [\text{proj}_{\bar{x}_{i_m}}^{\bar{z}'_m}(R'_m)\{\bar{u}_m/\bar{x}_{i_m}\}])$$

has a circularity. Since, for every $1 \leq j \leq m$,

$$\text{proj}_{\bar{x}_{i_j}}^{\bar{z}'_j}(R'_j)\{\bar{u}_j/\bar{x}_{i_j}\} \in I_{\mathcal{L}}^{(k+1)}(\mathfrak{f}_{i_j}(\bar{u}_j))$$

and since \mathbb{L} has no function invocation, we derive that $I_{\mathcal{L}}^{(k+1)}(\mathbb{L}[\mathfrak{f}_{i_1}(\bar{u}_1)] \cdots [\mathfrak{f}_{i_m}(\bar{u}_m)])$ has also a circularity. \square

Unlike [11,17], Lemma 2.7 is strict in our case, for every k . For example, consider

$$h(x, y, z) = (\nu t) (x, y)_t \& (y, z)_t .$$

Then, when $k \geq 1$, $I_{\mathcal{L}}^{(k)}(h) = \{\{(x, y)_\bullet, (y, z)_\bullet\}\}$. Notice that $I_{\mathcal{L}}^{(k)}(h(x, y, x)) = \{\{(x, y)_\bullet, (y, x)_\bullet\}\}$, which has a circularity – see Definition 2.1. However

$$I_{\mathcal{L}}^{(k)}((x, y)_t \& (y, x)_t) = \{\{(x, y)_t, (y, x)_t, (x, x)_t\}\}$$

has no circularity (this is a case of reentrant lock).

Theorem 2.8. Let $(\mathcal{L}, (\nu \bar{x}) \ell)$ be a lam program and $\ell \longrightarrow^* \ell'$. If $I_{\mathcal{L}}(\ell')$ has a circularity then $I_{\mathcal{L}}(\ell)$ has also a circularity. Therefore (\mathcal{L}, ℓ) has no circularity if $I_{\mathcal{L}}(\ell)$ has no circularity.

Proof. Let ℓ' have a circularity. Hence, by definition, $I^\perp(\ell')$ has a circularity and, since $I^\perp(\ell') = I_{\mathcal{L}}^{(0)}(\ell')$, by Proposition 2.5(2) $I_{\mathcal{L}}^{(0)}(\ell') \in I_{\mathcal{L}}(\ell')$. Therefore $I_{\mathcal{L}}(\ell')$ has also a circularity. Then, by Lemma 2.7, since $I_{\mathcal{L}}$ is the fixpoint interpretation function, $I_{\mathcal{L}}(\ell)$ has also a circularity. \square

Our algorithm for verifying that a lam will never manifest a circularity consists of computing $I_{\mathcal{L}}(\mathfrak{f})$, for every \mathfrak{f} , and $I_{\mathcal{L}}(\ell)$, where ℓ is the main lam. As discussed in this section, $I_{\mathcal{L}}(\mathfrak{f})$ uses a saturation technique on names based on a powerset construction. Hence it has a computational complexity that is *exponential* on the number of names. We remind that the names we consider are the *arguments* of lam functions (that corresponds to methods' arguments), which are usually not so many. In fact, this algorithm is quite efficient in practice [10].

3. The language and its semantics

In this section we define a simple programming model of concurrent object-oriented languages featuring threads and object synchronizations (the reader may easily recognize the basic operations of thread creation and synchronization used in Java and C#). We give a description of how deadlock may be identified, and discuss few examples. The following section contains the type system that associates lams to the programs.

Our model has two disjoint countable sets of names: there are *integer and object names*, ranged over by x, y, z, t, s, \dots , and *method names*, ranged over by m, n, \dots . A *program* is a tuple

$$(C_1 \mathcal{D}_1, \dots, C_n \mathcal{D}_n, P)$$

where \mathcal{D}_i are *finite sets of method name definitions* $m(\bar{\tau} \bar{x}) = P_m$, with $\bar{\tau} \bar{x}$ being the *formal parameters* \bar{x} and their types $\bar{\tau}$ and P_m being the *body* of m ; we assume that method name definitions have an implicit formal parameter C_i *this* which refers to the carrier. The process P is the *main process*. Types \mathbb{T} are either integers `int` or C_i .

The syntax of processes P and expressions e is defined below

$$\begin{aligned} P &::= 0 \mid (\text{new } C \ x) \ P \mid \text{fork}\{P\} \ P \mid \text{if } e \text{ then } P \text{ else } P \mid x.m(\bar{e}) \\ &\mid \text{sync}(x)\{P\} \\ e &::= x \mid v \mid e \text{ op } e \end{aligned}$$

A process can be the inert process 0 , or a restriction $(\text{new } C \ x) \ P$ that behaves like P except that the external environment cannot access to the object x of class C , or the spawn $\text{fork}\{Q\} \ P$ of a new thread Q by a process P , or a conditional $\text{if } e \text{ then } P \text{ else } Q$ that evaluates e and behaves either like P or like Q depending on whether the value is $\neq 0$ (*true*) or $= 0$ (*false*), or an invocation $x.m(\bar{e})$ of the method m with carrier x and actual parameters \bar{e} . The last process is $\text{sync}(x)\{P\}$ that executes P with exclusive access to x .

An expression e can be a name x , an integer value v , or a generic binary operation on integers op , where op ranges over a set including the usual operators like $+$, \leq , etc. Integer expressions without names (*constant expressions*) may be evaluated to an integer value (the definition of the evaluation of constant expressions is omitted). Let $\llbracket e \rrbracket$ be the evaluation of an expression e : $\llbracket e \rrbracket$ is undefined when e is an integer expression that contains integer names; $\llbracket x \rrbracket = x$ when x is a non-integer name.

To define the semantics of the above language, we use *states*, ranged over by \mathbb{P} , that are multisets of *threads* $P \bullet \sigma$, where σ is a (possibly empty) sequence of object names representing the locks that the process P has already grabbed (in order, from the left to right). When P terminates, the locks of σ must be all released. The thread $P \bullet \sigma$ is *reentrant* on x when σ contains at least two occurrences of x . We usually write states as $P_1 \bullet \sigma_1 \mid \dots \mid P_n \bullet \sigma_n$ and sometime shorten them into $\prod_{i \in 1..n} P_i \bullet \sigma_i$. We write $x \in \mathbb{P}$ if \mathbb{P} contains σ and x occurs in σ . Since states are multiset, we identify those that are equal up-to associativity and commutativity of \mid . Empty sequences σ are written ε . In order to ease the definitions, we will use few notational conventions

- Object names are partitioned into infinitely countable disjoint sets that are addressed by class names. The notation $z = \text{fresh}(C)$ states that z is a fresh name in the set C (e.g. z does not occur in the state).
- We make explicit that z is a name in the set C by writing $z \in C$.
- If a program Ψ has a pair $C \ \mathcal{D}$ such that \mathcal{D} contains a method $m(\bar{T} \ \bar{x}) = P$ then we write $C.m(\bar{T} \ \bar{x}) = P \in \Psi$.

Definition 3.1. The *operational semantics* of a program $\Psi = (C_1 \ \mathcal{D}_1, \dots, C_n \ \mathcal{D}_n, P)$ is a transition system where the initial state is $P \bullet \varepsilon$, and the *transition relation* \longrightarrow_Ψ is the least one closed under the rules:

$$\begin{array}{c} \text{(ZERO)} \\ 0 \bullet \sigma \longrightarrow_\Psi \\ \\ \text{(NEW0)} \quad \frac{z = \text{fresh}(C)}{(\text{new } C \ x) \ P \bullet \sigma \mid \mathbb{P} \longrightarrow_\Psi P\{z/x\} \bullet \sigma \mid \mathbb{P}} \quad \text{(NEWT)} \quad \frac{}{\text{fork}\{P\} \ Q \bullet \sigma \longrightarrow_\Psi P \bullet \varepsilon \mid Q \bullet \sigma} \\ \\ \text{(IFT)} \quad \frac{\llbracket e \rrbracket \neq 0}{\text{if } e \text{ then } P \text{ else } Q \bullet \sigma \longrightarrow_\Psi P \bullet \sigma} \quad \text{(IFF)} \quad \frac{\llbracket e \rrbracket = 0}{\text{if } e \text{ then } P \text{ else } Q \bullet \sigma \longrightarrow_\Psi Q \bullet \sigma} \\ \\ \text{(CALL)} \quad \frac{\llbracket \bar{e} \rrbracket = \bar{v} \quad x \in C \quad C.m(\bar{T} \ \bar{z}) = P \in \Psi}{x.m(\bar{e}) \bullet \sigma \longrightarrow_\Psi P\{x.\bar{v}/\text{this}, \bar{z}\} \bullet \sigma} \quad \text{(SYNC)} \quad \frac{x \notin \mathbb{P}}{\text{sync}(x)\{P\} \bullet \sigma \mid \mathbb{P} \longrightarrow_\Psi P \bullet \sigma \cdot x \mid \mathbb{P}} \end{array}$$

We often omit the subscript of \longrightarrow_Ψ when it is clear from the context. We write \longrightarrow^* for the reflexive and transitive closure of \longrightarrow .

We discuss the relevant rules. Rule (ZERO) models termination: every lock taken by the thread is released and the thread disappears from the state (which is a multiset). This is because, according to the syntax, there is no continuation in this case. Rule (NEW0) creates a new object of class C : we simply take a name that does not occur in the state, replace the bound name with the fresh one, and remove the binder. We observe that the renaming only applies to P : being x new, it cannot have been already locked (σ is not in the scope of the renaming). Rule (NEWT) spawns a new thread that has an empty sequence of locks (because it has not performed any synchronization so far). Rule (CALL) dispatches the invocation to the right method definition and executes the method body once the formal parameters have been replaced by the actual ones. Rule (SYNC) models synchronizations. A thread may lock x provided it is not already locked by another thread in the state – premise $x \notin \mathbb{P}$. We observe that x may occur in σ , that is a thread may lock x several times (thread reentrancy).

Our basic programming model does not retain the explicit instruction for sequential composition, e.g. $P ; Q$. However, this operation may be easily encoded using synchronizations:

$$(\text{new } C \ x) \ \text{sync}(x)\{ \text{fork}\{\text{sync}(x)\{ Q \}\} P \}$$

That is, we sequentialize P and Q by letting P to acquire the lock of a new object x *before* spawning Q . In turn Q will begin *after* the termination of P because the lock of x is grabbed by it. According to our rules, we have

$$\begin{aligned} (\text{new } C \ x) \ \text{sync}(x)\{ \text{fork}\{\text{sync}(x)\{ Q \}\} P \} \bullet \sigma & \\ \longrightarrow \text{sync}(x')\{ \text{fork}\{\text{sync}(x')\{ Q \}\} P \} \bullet \sigma & \\ \longrightarrow \text{fork}\{\text{sync}(x')\{ Q \}\} P \bullet \sigma \cdot x' & \\ \longrightarrow P \bullet \sigma \cdot x' \mid \text{sync}(x')\{ Q \} \bullet \varepsilon & \end{aligned}$$

and, by rule (SYNC), Q cannot start as long as x' is locked by P .

Definition 3.2 (*Deadlock-freedom*). A program $(C_1 \ \mathcal{D}_1, \dots, C_n \ \mathcal{D}_n, P)$ is *deadlock-free* if the following condition holds:

whenever $P \bullet \varepsilon \longrightarrow^* \mathbb{P}$ and $\mathbb{P} = \text{sync}(x)\{ Q \} \bullet \sigma \mid \mathbb{P}'$ then there exists \mathbb{P}'' such that $\mathbb{P} \longrightarrow \mathbb{P}''$.

Example 3.3. We select three processes and discuss their behaviors, highlighting whether they deadlock or not:

- $\text{fork}\{\text{sync}(x)\{ \text{sync}(y)\{ 0 \}\}\} \text{sync}(x)\{ \text{sync}(y)\{ 0 \}\}$. This process spawns a thread that acquire the locks of x and y in the *same order* of the main thread: no deadlock will ever occur.
- On the contrary, the process $\text{fork}\{\text{sync}(y)\{ \text{sync}(x)\{ 0 \}\}\} \text{sync}(x)\{ \text{sync}(y)\{ 0 \}\}$ spawns a thread acquiring the locks in reverse order. This is a computation giving a deadlock:

$$\begin{aligned} \text{fork}\{\text{sync}(y)\{ \text{sync}(x)\{ 0 \}\}\} \text{sync}(x)\{ \text{sync}(y)\{ 0 \}\} \bullet \varepsilon & \\ \longrightarrow \text{sync}(y)\{ \text{sync}(x)\{ 0 \}\} \bullet \varepsilon \mid \text{sync}(x)\{ \text{sync}(y)\{ 0 \}\} \bullet \varepsilon & \\ \longrightarrow \text{sync}(x)\{ 0 \} \bullet y \mid \text{sync}(x)\{ \text{sync}(y)\{ 0 \}\} \bullet \varepsilon & \\ \longrightarrow \text{sync}(x)\{ 0 \} \bullet y \mid \text{sync}(y)\{ 0 \} \bullet x & \end{aligned}$$

- The following method

$$\begin{aligned} \text{m}(C \ x, C \ y, \text{int } n) = & \\ \text{if } (n = 0) \text{ then } \text{fork}\{\text{sync}(y)\{ \text{sync}(x)\{ 0 \}\}\} \text{sync}(y)\{ 0 \} & \\ \text{else } \text{sync}(x)\{ \text{this.m}(x, y, n - 1) \} & \end{aligned}$$

performs n -nested synchronizations on *this* (reentrancy) and then spawns a thread acquiring the locks y and x in this order, while the main thread acquire the lock y . This method deadlocks for every $n \geq 1$, however it never deadlocks when $n \leq 0$.

4. Static semantics

Environments, ranged over by Γ , contain the types of objects, e.g. $x : C$, the type of integer variables e.g. $x : \text{int}$, and the types of method names, e.g. $C.m : [\bar{T}]$. Without loss of generality, in method definitions of Fig. 1 we order formal parameters of methods: those of type class are in front of those of type integer. Therefore $\Gamma(C.m)$ will have shape $[\bar{C}, \bar{\text{int}}]$.

Let $\text{dom}(\Gamma)$ be the domain of Γ and let

- $\Gamma, x:T$, when $x \notin \text{dom}(\Gamma)$

$$(\Gamma, x:T)(y) \stackrel{\text{def}}{=} \begin{cases} T & \text{if } y = x \\ \Gamma(x) & \text{otherwise} \end{cases}$$

- $\Gamma + \Gamma'$, when $x \in \text{dom}(\Gamma) \cap \text{dom}(\Gamma')$ implies $\Gamma(x) = \Gamma'(x)$:

$$(\Gamma + \Gamma')(x) \stackrel{\text{def}}{=} \begin{cases} \Gamma(x) & \text{if } x \in \text{dom}(\Gamma) \\ \Gamma'(x) & \text{if } x \in \text{dom}(\Gamma') \\ \text{undefined} & \text{otherwise} \end{cases}$$

$$\begin{array}{c}
\text{Processes:} \quad \begin{array}{l}
\text{(T-ZERO)} \quad \frac{}{\Gamma; \sigma \vdash_t 0 : (\sigma)^t} \\
\text{(T-NEW)} \quad \frac{\Gamma, x:C; \sigma \vdash_t P : \ell \quad x \notin \sigma}{\Gamma; \sigma \vdash_t (\text{new } C \ x) \ P : (v \ x) \ \ell} \\
\text{(T-IF)} \quad \frac{\Gamma \vdash_t e : \text{int} \quad \Gamma; \sigma \vdash_t P : \ell \quad \Gamma; \sigma \vdash_t P' : \ell'}{\Gamma; \sigma \vdash_t \text{if } e \text{ then } P \text{ else } P' : \ell + \ell'} \\
\text{(T-IF)} \quad \frac{\Gamma; \sigma \cdot x \vdash_t P : \ell}{\Gamma; \sigma \vdash_t \text{sync}(x)\{ P \} : \ell} \\
\text{(T-IF)} \quad \frac{\Gamma; \sigma \vdash_t P : \ell \quad \Gamma, t' : \text{Thread}, u' : \text{Obj}; u' \vdash_{t'} P' : \ell'}{\Gamma; \sigma \vdash_t \text{fork}\{P'\} \ P : \ell \& (v \ t', u') \ \ell'} \\
\text{(T-IF)} \quad \frac{\Gamma(x) = C \quad \Gamma(C.m) = [\overline{C'}, \overline{\text{int}}] \quad \Gamma \vdash \overline{u} : \overline{C'} \quad \Gamma \vdash \overline{e} : \overline{\text{int}}}{\Gamma; \sigma \cdot z \vdash_t x.m(\overline{u}, \overline{e}) : f_{C.m}(t, z, x, \overline{u}) \& (\sigma \cdot z)^t}
\end{array} \\
\text{Expressions:} \quad \begin{array}{l}
\text{(T-INT)} \quad \Gamma \vdash v : \text{int} \\
\text{(T-VAR)} \quad \Gamma, x : T \vdash x : T \\
\text{(T-OP)} \quad \frac{\Gamma \vdash e : \text{int} \quad \Gamma \vdash e' : \text{int}}{\Gamma \vdash e \text{ op } e' : \text{int}} \\
\text{(T-SEQ)} \quad \frac{(\Gamma \vdash e_i : T_i)^{i \in 1..n}}{\Gamma \vdash e_1, \dots, e_n : T_1, \dots, T_n}
\end{array} \\
\text{Programs:} \quad \begin{array}{l}
\text{(T-PROG)} \quad \frac{\Gamma = \{C.m : [\overline{C'}, \overline{\text{int}}] \mid C.m(\overline{C'} \ \overline{x}, \overline{\text{int}} \ \overline{y}) = P_{C.m} \in \Psi\} \\
\left(\Gamma, \text{this} : C, \overline{x} : \overline{C'}, \overline{y} : \overline{\text{int}}, t : \text{Thread}, u : \text{Obj}; u \vdash_t P_{C.m} : \ell_{C.m} \right)^{C.m(\overline{C'} \ \overline{x}, \overline{\text{int}} \ \overline{y}) = P_{C.m} \in \Psi} \\
\mathcal{L} = \bigcup_{C.m(\overline{C'} \ \overline{x}, \overline{\text{int}} \ \overline{y}) = P_{C.m} \in \Psi} \{ f_{C.m}(t, z, \text{this}, \overline{x}) = \ell_{C.m} \} \\
\Gamma, t : \text{Thread}, u : \text{Obj}; u \vdash_t \Psi.\text{main} : \ell}{\Gamma \vdash \Psi : (\mathcal{L}, \ell)}
\end{array}
\end{array}$$

Fig. 1. The type system (we assume a function name $f_{C.m}$ for every method name m of class C).

Let also

$$(x_1 \cdots x_n \cdot x_{n+1})^t \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } n = 0 \\ (x_1 \cdots x_n)^t & \text{if } x_{n+1} \in \{x_1, \dots, x_n\} \\ (x_1 \cdots x_n)^t \& (x_n, x_{n+1})_t & \text{otherwise} \end{cases}$$

The static semantics has three judgments:

- $\Gamma \vdash e : T$ – the expression e has type T in Γ ;
- $\Gamma; \sigma \vdash_t P : \ell$ – the thread P with name t has lam ℓ in $\Gamma; \sigma$. We assume that the class of t is a predefined class `Thread` and Γ contains the bindings $t : \text{Thread}$;
- $\Gamma \vdash \Psi : (\mathcal{L}, \ell)$ the program Ψ has lam program (\mathcal{L}, ℓ) in Γ .

In addition to `Thread`, the static semantics also uses another predefined class, called `Obj`. The objects of this class are used for typing new threads (cf. rule (T-PAR) in Fig. 1; this expedient allows us to have a simpler rule for typing function invocations).

Few notational conventions are in order. When $\Psi = (C_1 \ \mathcal{D}_1, \dots, C_n \ \mathcal{D}_n, P)$, we let $\Psi.\text{main} = P$ and we write $C_{i.m}(\overline{T} \ \overline{x}) = P_{C_{i.m}} \in \Psi$ whenever \mathcal{D}_i contains $m(\overline{T} \ \overline{x}) = P_{C_{i.m}}$.

The type system is defined in Fig. 1. A few key rules are discussed. Rule (T-ZERO) types the process 0 in a thread t that has locked the objects in σ in (inverse) order. The lam is the conjunction of dependencies in σ with thread t – cf. notation $(\sigma)^t$. Rule (T-SYNC) types the critical section P with a sequence of locks extended with x . Rule (T-PAR) types a parallel composition of processes by collecting the lams of the components. We use a new name t' for the spawned thread and a new name u' for the object that is already grabbed by t' . This object has class `Obj`. (According to our convention, when we write $\Gamma, t' : \text{Thread}, u' : \text{Obj}$ we assume that $t', u' \notin \text{dom}(\Gamma)$.) Rule (T-CALL) types a method name invocation in terms of a (lam) function invocation and constrains the types of actual arguments to match with the types of formal parameters in the method declaration. The arguments of the lam function invocation are extended with the thread name of the caller and the name of the last object locked by it (and not yet released). In addition we also conjunct the dependencies $(\sigma \cdot x)^t$ created by the caller because the method invocation run in a thread that holds those locks. This rule must be read in conjunction

with (T-PROG). (T-PROG) associates a lam function to every method – c.f. its second and third premises. The body of the lam function is defined by typing the body of the method in the hypothetical context where the thread only owns one lock, which is given as argument of the lam function. Of course, this is not the general case: a method may be invoked by a thread that grabs many locks. To address the general case we record the last object synchronized by the thread and build the dependencies of the invocation starting from it. It turns out that taking these dependencies in conjunction with those of the caller – see rule (T-CALL) – does suffice for keeping track of the overall dependencies (see Proposition A.1). Because of this technicality, we require that every thread, even those that are just created, locks at least one object.

Example 4.1. Let us show the typing of the method `setTable` in the Introduction. Let

$$\begin{aligned}\Gamma &= C.\text{setTable}:[C, C, \text{int}], \text{this}:C, x:C, y:C, n:\text{int}, t:\text{Thread}, u:\text{Obj} \\ P &= \text{sync}(y)\{\text{sync}(x)\{0\}\} \\ Q &= \text{fork}\{\text{sync}(x)\{\text{sync}(z)\{0\}\}\} z.\text{setTable}(z, y, n-1)\end{aligned}$$

Then

$$\frac{\frac{\Gamma, z:C; u \cdot y \cdot x \vdash_t 0:\ell_1 \quad \ell_1 = (u, y)_t \& (y, x)_t}{\Gamma, z:C; u \cdot y \vdash_t \text{sync}(x)\{0\}:\ell_1} \quad (*)}{\frac{\Gamma, z:C; u \vdash_t P:\ell_1 \quad \Gamma, z:C; u \vdash_t Q:\ell_2}{\Gamma, z:C; u \vdash_t \text{if } n=0 \text{ then } P \text{ else } Q:\ell_1 + \ell_2}}{\Gamma; u \vdash_t (\text{new } C \ z) (\text{if } n=0 \text{ then } P \text{ else } Q):(\nu z) (\ell_1 + \ell_2)}$$

where (*) are the two proof trees

...

$$\Gamma, z:C; u \vdash_t z.\text{setTable}(z, y, n-1):\ell'_2$$

and

...

$$\Gamma, z:C, t':\text{Thread}, u':\text{Obj}; w' \vdash_{t'} \text{sync}(x)\{\text{sync}(z)\{0\}\}:\ell''_2$$

(as an exercise, the reader may try to complete them). By rule (T-PAR), we derive $\ell_2 = \ell'_2 \& (\nu t', u') \ell''_2$. After completing the proof tree, one obtains the lam function

$$\begin{aligned}\text{setTable}(t, u, \text{this}, x, y) = \\ (\nu z, t', u') \left((u, y)_t \& (y, x)_t + (u', x)_{t'} \& (x, z)_{t'} \& \text{setTable}(t, u, x, z, y) \right)\end{aligned}$$

which has two additional arguments with respect to the one in the Introduction: u and this . In fact, these arguments, as well as u' , do not play any role in the analysis of circularities of $\text{setTable}(t, u, \text{this}, x, y)$.

The following theorem states the soundness of our type system. The technical details of the proof are reported in the Appendix.

Theorem 4.2. Let $\Gamma \vdash (C_1 \mathcal{D}_1, \dots, C_n \mathcal{D}_n, P):(\mathcal{L}, \ell)$. If (\mathcal{L}, ℓ) has no circularity then $(C_1 \mathcal{D}_1, \dots, C_n \mathcal{D}_n, P)$ is deadlock-free.

Example 4.3. Let us verify whether the process $x.\text{setTable}(x, x, n)$ is deadlock-free. The lam function associated by the type system is detailed in Example 4.1. The interpretation function $I_{\mathcal{L}}(\text{setTable})$ is computed as follows:

$$\begin{aligned}I_{\mathcal{L}}^{(0)}(\text{setTable}) &= \{\emptyset\} \\ I_{\mathcal{L}}^{(1)}(\text{setTable}) &= \{\{(u, y)_t, (y, x)_t, (u, x)_t\}\} \\ I_{\mathcal{L}}^{(2)}(\text{setTable}) &= \{\{(u, y)_t, (y, x)_t, (u, x)_t\}, \{(u, y)_t\}\}.\end{aligned}$$

Since $I_{\mathcal{L}}^{(2)} = I_{\mathcal{L}}$, we are reduced to compute $I_{\mathcal{L}}^{(2)}(\text{setTable}(t, u, x, x, x))$. That is

$$\{\{(u, y)_t, (y, x)_t, (u, x)_t\}, \{(u, y)_t\}\} \{x/y\} = \{\{(u, x)_t, (x, x)_t\}, \{(u, x)_t\}\}$$

which has no circularity, therefore the process $x.\text{setTable}(x, x, n)$ is deadlock-free.

It is interesting to verify whether the process $x.setTableD(x, x, n)$ is deadlock-free, where $setTableD$ is the method having philosophers with symmetric strategies:

```

setTableD(C x, C y, int n) = (new C z) (
    if (n=0) then sync(x){ sync(y){ 0 } }
    else fork { sync(x){ sync(z){ 0 } } }
    z.setTableD(z, y, n-1)
)

```

In this case $I_{\varnothing}(setTableD) = \{(u, x)_t, (x, y)_t, (u, y)_t\}, \{(x, y)_{\checkmark}, (u, y)_t\}$. It is easy to verify that $I_{\varnothing}(setTableD(t, u, x, x, x))$ has a circularity, therefore the process $x.setTableD(x, x, n)$ may have (and actually has) a deadlock.

To conclude, we observe that the deadlock freedom property guaranteed by Theorem 4.2 is stronger than the one defined in Definition 3.2 that admits a subset of processes to deadlock provided that some other process progresses. In fact, as usual [16,17], our type system guarantees that no subset of processes deadlocks.

5. The extension of inheritance

The basic language in Section 3 does not admit derivations of classes from other classes. As a consequence, when a method is invoked, it is possible to uniquely locate the method definition. This feature is expressed in the type system in Fig. 1 by the strict correspondence between types of arguments and types of formal parameters (e.g. the constraint $\Gamma \vdash \bar{u} : \bar{C}$ in the type rule (T-CALL)). On the contrary, object-oriented programming languages admit subclasses, thereby allowing redefinitions of methods in the subclasses. In this context, locating a method definition is performed at run-time through a mechanism called dynamic dispatch that uses the actual value of the called object.

For example, let C be a class with a method foo that is overridden in the subclasses C_1 and C_2 . Then, consider a method with arguments x, y, z of types $C, C,$ and int , respectively. If this snippet

```
if z then x.foo() else y.foo()
```

is part of its body and the method is invoked with objects of types C_1 and C_2 , the method foo of C_1 or of C_2 will be performed according to the value of z . Clearly, dispatching an invocation to a definition rather than another may change the overall result of the deadlock analysis.

In order to study the impact of inheritance on the deadlock analysis, we extend the language in Section 3 as follows. Let a *program* be a tuple

$$(D_1 \mathcal{D}_1, \dots, D_n \mathcal{D}_n, P)$$

where every \mathcal{D}_i is as in Section 3 and D is a term of the following grammar

$$D ::= C \mid C \text{ extends } C$$

This extension carries a subtyping relation $<$: that is reflexive, transitive and contains $C <: C'$ if $C \text{ extends } C'$ is in the program. Processes are extended with a more flexible object creation operation:

$$P ::= \dots \mid (\text{new } C_1, \dots, C_m \ x) P$$

meaning that the type of x may be one of C_1, \dots, C_m (it is non-deterministic) and it is intended that there is C such that $C_i <: C$ for every C_i . This apparently exotic operation corresponds to the pattern

```

C x ;
if e1 x = new C1() ;
else if e2 x = new C2() ;
...
else x = new Cm() ;

```

that we cannot use because our basic language does not feature assignments (see Section 6).

The static semantics of the foregoing extensions uses (finite) sets of object names $\{C_1, \dots, C_m\}$ such that there is a C with $C_i <: C$ for every C_i . Let $\mathcal{C}, \mathcal{C}', \dots$ range over these sets. Let also environments Γ bind object names to sets \mathcal{C} . The typing rules become ($+_{0 \leq i \leq n} \ell_i$ is an abbreviation for $\ell_1 + \dots + \ell_n$):

(T-SUBT)

$$\frac{C \text{ [extends } C'] \{ \dots, m(\overline{C'} \bar{x}, \overline{\text{int}} \bar{y}) = P, \dots \} \text{ is in the program} \quad E <: C \quad \overline{E'} <: \overline{C'}}{\Gamma \vdash E.m : [\overline{E'}, \overline{\text{int}}]}$$

(T-NEW-INH)

$$\frac{\Gamma, x:\mathcal{C}; \sigma \vdash_t P : \ell}{\Gamma \vdash_t (\text{new } \mathcal{C} \ x) P : (\nu x) \ell}$$

(T-CALL-INH)

$$\frac{\Gamma(z) = \mathcal{C} \quad \Gamma \vdash \bar{u} : \overline{\mathcal{C}'} \quad \left(\Gamma \vdash C.m : [\overline{C}, \overline{\text{int}}] \right)^{C \in \mathcal{C}, \overline{C} \in \overline{\mathcal{C}'}} \quad \Gamma \vdash \bar{e} : \overline{\text{int}}}{\Gamma; \sigma \cdot x \vdash_t z.m(\bar{u}, \bar{e}) : \left(+_{C \in \mathcal{C}, \overline{C} \in \overline{\mathcal{C}'}} f_{C.m[\overline{C}]}(t, x, z, \bar{u}) \right) \& (\sigma \cdot x)^t}$$

In the rule (T-SUBT), the text in $[\dots]$ is optional: the rule applies both to class and to subclass definitions. Our analyzer associates a lam function $f_{C.m[\overline{C}]}$ to every method m such that $\Gamma \vdash C.m : [\overline{C}, \overline{\text{int}}]$. As in the system of Fig. 1, the body ℓ of this function is defined by the judgment (it is assumed that $m(\overline{C'} \bar{x}, \overline{\text{int}} \bar{y}) = P$ belongs to a class C):

$$\Gamma, \bar{x} : \{\overline{C'}\}, \bar{y} : \overline{\text{int}}, t : \{\text{Thread}\}, u : \{\text{Obj}\}; u \vdash_t P : \ell$$

(T-CALL-INH) is the critical rule. Since statically we do not know what method will be actually executed at runtime, we return the disjunctions of the lams of every possible method, in conjunction with the dependencies $(\sigma \cdot x)^t$ created by the caller. We illustrate the point by discussing an example.

Example 5.1. Consider the following program:

```
C1 { m(C x, C y) = sync(x) { sync(y) { 0 } } } ;
C2 extends C1 { m(C x, Cy) = sync(y) { sync(x) { 0 } } } ;
new C x, C y, {C1,C2} u. fork { u.m(x,y) } sync(x) { sync(y) { 0 } }
```

In this case, the type of u may be either $C1$ or $C2$, therefore the method m that will be actually invoked will be either $C1.m$ or $C2.m$. Therefore we may have a deadlock because the main thread locks x and then y , while the spawned thread $C2.m$ locks y and then x . There is no deadlock if the type of u is $C1$.

By rule (T-CALL-INH) and (T-PAR) we derive the following lam for main:

$$\ell = ((u, x)_t \& (x, y)_t) \ \& \ (\nu t') (((u, x)_{t'} \& (x, y)_{t'}) + ((u, y)_{t'} \& (y, x)_{t'}))$$

It is easy to verify that $I^\perp(\ell)$ has a circularity.

The foregoing management of inheritance is generic. In several cases, it is possible to have a more specific solution using explicit typings. In particular, whenever the run-time types of objects are known, it is possible to attach the right type to method invocations and apply concrete type reasoning. `JaDA` uses this technique starting the analysis from the main program.

6. Other extensions

In order to deliver initial assessments of our technique, we have extended the language in Sections 3 and 5 to cover many features of `Java`. The resulting tool, called `JaDA` [10,19], has been prototyped. Overall, the attitude we had in the development of `JaDA` was to stick to solutions that can be easily integrated with the type system of the core language (and possibly simple to implement). In many places, we might have taken more refined solutions (see, for instance, the following paragraph about recursive types), with the danger of further delaying the release of the tool because of issues that are orthogonal to our technique.

In this section we overview two relevant extensions – assignments and recursive data types – and discuss our solutions in `JaDA`; the details of other extensions (e.g. static members, exceptions, and arrays) can be found in Garcia's PhD thesis [9]. The comparison of `JaDA` with other deadlock analysis tools is reported in Section 7.

Assignments. In the language of Sections 3 and 5 it is not possible to store and retrieve values. These operations are usually achieved by admitting objects with *fields* and the operations of read and write of fields. E.g. the syntax of processes and expressions is extended as follows

$$\begin{array}{l} P ::= \dots \mid \mathfrak{f} = e. P \\ e ::= \dots \mid x.m(\bar{e}) \end{array}$$

(method invocations may return values).

Assignments impact the type system in a significative way because field values change during the execution and these changes may be unpredictable because of concurrency. In addition, since fields may store other objects and may be method's arguments, it is possible to create aliases. Overall, tracing the state changes becomes quite hard. The type system of JaDA overcomes these issues by

- *checking data races* of parallel threads (for example, one thread reading an object and the other one writing on the object). In particular, the environment Γ binds fields to annotations in $\{-, r, w\}$, with $-$ meaning no access, r meaning read access, w meaning write (and possible read) access. The set $\{-, r, w\}$ is a total order ($- < r < w$) and we use the operation of least upper-bound \sqcup .
- *analyzing aliases* created by copying mechanisms. This is implemented by decoupling the connection between object names and values in the environment. E.g. we introduce *object pointers*, noted a, b, \dots (that include the pointer \top) and we let the environment to bind object names to object pointers and object pointer to pairs $([\mathfrak{f}_1^{h_1} : a_1, \dots, \mathfrak{f}_m^{h_m} : a_m], \mathcal{C})$, where h_i are in $\{-, r, w\}$ and a_i may also be \top . The pointer \top represents a record with fields that are tagged w and with values \top .

In order to trace data races, we use a companion type system $\Gamma \vdash^e P : \Gamma'$ that stores in Γ' the type of accesses performed by P to fields of objects in Γ (we assume that accesses in Γ are all set to $-$; when this is not the case, we use the function $\Gamma \uparrow^-$ that rewrites accesses into $-$). The *merge operation* on environments, noted $\Gamma \parallel \Gamma'$, returns an environment where fields have the least upper-bounds of accesses in Γ and Γ' while values are set to \top when accesses are conflictual (e.g. either r - w or w - w). For instance, let

$$\Gamma = a \mapsto ([\mathfrak{f}^r : b, \mathfrak{g}^- : b'], \mathcal{C}) \quad \Gamma' = a \mapsto ([\mathfrak{f}^w : c, \mathfrak{g}^w : d], \mathcal{C})$$

then $\Gamma \parallel \Gamma' = a \mapsto ([\mathfrak{f}^w : \top, \mathfrak{g}^w : d], \mathcal{C})$, that is

1. the field \mathfrak{f} has access w because this is the least upper bound of the two accesses; additionally the resulting value of \mathfrak{f} is \top because we do not know whether the read access in Γ will read b or c ;
2. the field \mathfrak{g} has access w because Γ does not accesses to it while Γ' writes it; additionally the resulting value of \mathfrak{g} is d because this is the value written by Γ' .

Updates are recorded by means of judgments that also return environments. Namely, we use $\Gamma; \sigma \vdash_t P : \ell, \Gamma'$, where Γ' is the environment Γ updated by P . We also use the operation $\Gamma[\Gamma']$ that returns the environment Γ updated according to what is prescribed in Γ' . The typing rule for the spawn of new threads is

$$\frac{\text{(T-PAR-EFF)} \quad \begin{array}{c} \Gamma \uparrow^- \vdash^e P : \Gamma_P \quad \Gamma \uparrow^- \vdash^e P' : \Gamma_{P'} \\ \Gamma[\Gamma_P \parallel \Gamma_{P'}]; \sigma \vdash_t P : \ell, \Gamma' \quad \Gamma[\Gamma_P \parallel \Gamma_{P'}], u' : \text{Obj}; u' \vdash_{t'} P' : \ell', \Gamma'' \end{array}}{\Gamma \vdash_t \text{fork}\{P\} P : \ell \& (v t', u') \ell', \Gamma' \parallel \Gamma''}$$

We notice that P and P' are verified in an environment that already record the possible data races (a priori, P might start after P' or conversely). We also notice that the returned environment in the conclusion is the merge of the returned environments in the premises.

The presence of assignments also impacts on lams. First of all, in order to model object values, arguments of lam functions may also be records. This extension is not problematic with respect to Theorem 2.8 because our records are finite trees with fixed structures (this is also the case when object have recursive types, see below). In addition, the special name \top influences the definition of \mathbb{R}^+ in Section 2. In fact, this name, being a place holder for possible object names, may be replaced by every name when transitive closure is computed, thereby overapproximating the set \mathbb{R}^+ with more dependencies (and circularities). For example, $(\top, a)_t \& (a, b)_{t'}$ is a circularity (because \top may be replaced by b), while $(\top, a)_t \& (a, b)_t$ is not a circularity (because, in case, this may be a reentrant thread).

Recursive types. In JaDA , the analysis of objects with recursive types is over-approximated by using finite representations. In particular, a record of a *recursive type* is built by unfolding the type (exactly) up to those nodes containing a name of a class already present in the tree. Nodes inside the tree have values that are object names, nodes in the leaves whose type

Table 1

Comparison with different deadlock detection tools. The inner cells show the number of deadlocks detected by each tool. The output labeled “(*)” are related to modified versions of the original programs: see the text.

Benchmarks	LOC, #Threads	Deadlock	Static		Hybrid	Dynamic	Commercial
			JaDA[tm]	Chord[tm]	GoodLock[tm]	Sherlock[tm]	ThreadSafe[tm]
Sor	1274, 5	yes	1 [135s]	1 [210s]	7 [4s]	1 [39s]	4 [435s]
RayTracer(*)	1292, 5	no	0 [155s]	0 [223s]	8 [2s]	2 [30s]	0 [502s]
MolDyn (*)	1351, 5	no	0 [110s]	0 [191s]	6 [5s]	1 [49s]	0 [423s]
MonteCarlo (*)	3619, 4	no	0 [231s]	0 [342s]	23 [5s]	2 [102s]	0 [821s]
BuildNetworkN	40, N+1	yes	3 [8s]	0 [50s]			0 [50s]
PhilosophersN	60, N+1	yes	3 [12s]	0 [51s]			0 [51s]
ThreadArraysN	23, N+1	yes	1 [6s]	1 [40s]			1 [40s]
ThreadArraysJoinsN	37, N+1	yes	1 [6s]	1 [41s]			0 [41s]

is not recursive are marked `int` (a primitive type value); leaves of a recursive type have values that are names subscribed by a $\$$; the name corresponds to the nodes of the classes that are already present in the tree.

For instance, if C is a class whose type is $[\text{val} : \text{Thread}, \text{next} : C]$ (a list of threads) then, in correspondence of a `new C` instruction, we produce an environment $r_\$ \mapsto [\text{val} : (a[], \text{Thread}), \text{next} : r_\$]$. Lists like the foregoing one are therefore abstracted by a single node. The analyzer replaces names indexed by $\$$ with T and uses the same approximation described in previous paragraph. In particular, if the types of $r_\$$ and of $r'_\$$ are the same, a term like $(r_\$, a)_t \& (b, r'_\$)_{t'}$ is a circularity because it corresponds to $(T, a)_t \& (b, T)_{t'}$.

7. JaDA and its assessment

Since JaDA covers many features of Java, it has been possible to deliver an initial assessment of it with respect to existing deadlock analysis tools. In particular, we have considered tools using different techniques: Chord for static analysis [23], Sherlock for dynamic analysis [6], and GoodLock for hybrid analysis [3]. We have also considered a commercial tool, ThreadSafe³ [2]. Out of these tools, we were able to install and effectively test only two of them – Chord and ThreadSafe – because the other twos have been discontinued; the results corresponding to GoodLock and Sherlock come from [6]. We also had problems in testing Chord with some of the examples in the benchmarks, perhaps due to some misconfigurations, that we were not able to solve because Chord has been also discontinued.

The analyzed programs exhibit a variety of sharing patterns. The source of all benchmarks in Table 1 is available either at [6,23] or in the JaDA-deadlocks repository.⁴ Since the current release of JaDA does not completely cover Java, in order to gain preliminary experience, we modified the Java libraries and the multithreaded server programs of RayTracer, MolDyn and MonteCarlo (labeled with “(*)” in the Table 1) and implemented them in our system. This required little programming overhead; in particular, we removed volatile variables, avoided the use of Runnable interfaces for creating threads, and reduced the invocations of native methods involved in I/O operations. For every program, we give the lines of code (LOC), the number #Threads of threads explicitly created (in the second and third block this number depends on the argument N). We also state whether the program under examination has a deadlock or not and the time in seconds (tm) each tool took to perform the analysis. The times for GoodLock and Sherlock were taken from the literature [6].

Here are our remarks. The first block of programs belongs to a well known group used as benchmarks for several Java analysis tools; the second block corresponds to examples designed to test JaDA against complex deadlock scenarios. First of all, at least for the benchmark codes in Table 1, JaDA is the unique tool that never returns false negatives (in general, JaDA may return false positives). Chord and ThreadSafe are unsound because they return false negatives (see the second block). The execution time of the tools are similar (JaDA appears more efficient), except for GoodLock and Sherlock, which appear however much less precise (they return a lot of false positives). As regards the second block, we observe that JaDA returns few deadlocks, which do not depend from N. This is because our analysis is symbolic and does not consider numeric values (most of the deadlock are considered “to be similar”).

We retain that the results of Table 1 are encouraging and solicit further programming efforts for covering Java features that have not been considered so far, as well as more elaborate comparisons. This is left for future work.

It is also worth to observe that JaDA actually addresses the Java bytecode (we compile Java programs into bytecode and analyze it). This means that JaDA also analyzes every programming language that is compiled into Java bytecode, such as Scala [24]. As regards Scala, to the best of our knowledge, JaDA is the first static deadlock analysis tool.

³ <http://www.contemplateld.com/threadsafe>.

⁴ <https://github.com/abelunibo/Java-Deadlocks>.

8. Remarks about the analysis technique

The deadlock analysis technique presented in this paper is lightweight because it is compact, intelligible and theoretically manageable. The technique is also very powerful because we can successfully verify processes like *buildTable* and its variant where every philosopher has a symmetric strategy. However, there are processes for which our technique of collecting dependencies is too rough (and we get false positives).

One example is

$$\text{fork}\{\text{sync}(x)\{\text{sync}(z)\{\text{sync}(y)\{0\}\}\}\}\text{sync}(x)\{\text{sync}(y)\{\text{sync}(z)\{0\}\}\}.$$

This process has two threads: the first one locks x and then z and y in order; the second one locks x and then grabs y and z in order. Since the two threads initially compete on the object x , they will be executed *in sequence* and no deadlock will ever occur. However, if we compute the dependencies, we obtain

$$(x \cdot z \cdot y)^t \& (x \cdot y \cdot z)^s$$

that is equal to $((x, z)_t \& (z, y)_t) \& ((x, y)_s \& (y, z)_s)$ where the reader may easily recognize the circularity $(z, y)_t \& (y, z)_s$. This inaccuracy follows by the fact that our technique does not record the dependencies between threads and their state (of locks) when the spawns occur: this is the price we pay to simplicity. In [10] we overcome this issue by associating line codes to symbolic names and dependencies. Then, when a circularity is found, we can exhibit an abstract witness computation that, at least in the simple cases as the above one, can be used to manually verify whether the circularity is a false positive or not.

Another problematic process is $m(C\ x, C\ y, \text{int}\ n)$ in Example 3.3. This process never deadlocks when $n \leq 0$. However, since our technique drops integer values, it always return a circularity (this is a correct result when $n > 0$ and it is a false positive otherwise). To cope with these cases, it suffices to complement our analysis with standard techniques of data-flow analysis and abstract evaluation of expressions.

9. Related works and conclusions

We have defined a simple technique for detecting deadlocks in object-oriented programs. This technique uses an extension of the lam model in order to cope with reentrant locks, a standard feature of object-oriented programs. We have specified an algorithm for verifying the absence of circularities in extended lams and we have applied this model to a simple concurrent object-oriented calculus that is intended to be a minimal concurrent subset of Java. We demonstrate the effectivity of our analysis by applying it to Java. In particular, we discuss a prototype verifier, called JaDA, and deliver initial assessments that seem quite encouraging.

The lam model has been introduced and studied in [12] for detecting deadlocks of an object-oriented language with futures (and no lock and lock reentrancy) [13], but the extension discussed in this paper is new as well as the algorithm for the circularity of lams. We have prototyped this algorithm in JaDA, where we use it for the deadlock analysis of Java [10]. As we discussed in the Introduction, the model has been also applied to process calculi [11,17].

Several techniques have been developed for the deadlock detection of concurrent object-oriented languages. The technique [2] uses a data-flow analysis that constructs an execution flow graph and searches for cycles within this graph. Some heuristics are used to remove likely false positives. No alias analysis to resolve object identity across method calls is attempted. This analysis is performed in [6,23], which can detect reentrance on restricted cases, such as when lock expressions concern local variables (the reentrance of formal parameters, as in *buildTable*($x, x, 0$) is not detected). The technique in [3] and its refinement [6] use a theory that is based on monitors. Therefore the technique is a runtime technique that tags each segment of the program reached by the execution flow and specifies the exact order of lock acquisitions. Thereafter, these segments are analyzed for detecting potential deadlocks that might occur because of different scheduler choices (than the current one). This kind of technique is partial because one might overlook sensible patterns of methods' arguments (cf. *buildTable*, for instance). A powerful static techniques that is based on abstract interpretation is SACO [8]. SACO has been developed for ABS [14], an object-oriented language with a concurrent model different from Java. A comparison between SACO and a tool using a technique similar to the one in this paper can be found in [13].

Our future work includes the analysis of concurrent features of object-oriented calculi that have not been studied yet. A relevant one is thread coordination, which is usually expressed by the methods *wait* and *notify* (and *notifyAll*). These methods modify the scheduling of processes: the thread executing *wait*(x) is suspended, and the corresponding lock on x is released; the thread executing *notify*(x) wakes up one thread suspended on x , which will attempt again to grab x . A simple deadlock in programs with *wait* and *notify* is when the *wait* operation is either mismatched or *happens-after* the matching notification. Initial steps toward the deadlock analysis of coordination primitives can be found in [20], where we design a behavioral type system for coordination primitives. This system resorts to simple object protocols – called *usages* – to determine whether objects are used reliably, so as to guarantee deadlock freedom. In turn, this reliability is reduced to a reachability problem in a Petri net that, in the case of coordination primitives, have also inhibitor arcs. While the technique in [20] seems complementary to our extended lam model, the actual integration must still be studied.

A research direction that is close to the above one is the analysis of concurrent libraries, such as `java.util.concurrent`. This issue has several cases, we discuss the most relevant ones. One case is when a library function can be implemented in `Java` (bytecode) using features that are already covered [9] or that can be covered by our analysis (e.g. coordination primitives). In this case, we derive the corresponding lam either in an automatic way or in a semi-automatic way, by providing the missing judgments, and use *explicit typing* in every place where the function is used. A second case is when a function is written in native code. In this case, we compute all the dependencies by parsing the code manually and type the function with the resulting lam (and again use explicit typing). Another case is when the function uses synchronization primitives that rely on conditional expressions, such as conditional `await`. In this case, a reasonable solution is integrating our technique with some axiomatic formalism using pre- and post-conditions. These axiomatic formalisms are also very powerful for tackling assignments and pointers (see Section 6), cf. separation logic [26]. The study of this integration and the possible definition of an hybrid solution for the deadlock analysis of `Java`-like languages is matter of future research.

Further work will be also necessary for the development of `JaDA`: while the results of Table 1 are encouraging, they also solicit further programming efforts for covering `Java` features that have not been considered so far. This should allow us to address the benchmarks labeled (*) without any programming overhead and possibly other benchmarks as in [23].

Acknowledgements

We thank Abel Garcia and Elena Giachino for the fruitful discussions and useful comments. In particular, Abel has been the main developer of `JaDA`. We also thank the anonymous referees for reading the paper carefully and for their many insightful comments and suggestions.

Appendix A. Proof of Theorem 4.2

We first need to extend the typing to states. Therefore we extend the type system in Fig. 1 with the following two rules:

$$\frac{\text{(T-THREAD)} \quad \Gamma, t : \text{Thread}; \sigma \vdash_t P : \ell}{\Gamma \vdash P \bullet \sigma : (v \ t) \ \ell} \quad \frac{\text{(T-STATE)} \quad (\Gamma \vdash P_i \bullet \sigma_i : \ell_i)_{i \in 1..n}}{\Gamma \vdash \prod_{i \in 1..n} P_i \bullet \sigma_i : \&_{i \in 1..n} \ell_i}$$

Theorem 4.2 is a consequence of the following statements. The first two are obtained by a standard induction on the proof-tree of the judgment in the premise.

Proposition A.1. *If $\Gamma; x \vdash_t P : \ell$ and names in σ belong to $\text{dom}(\Gamma)$ then there is ℓ' such that $\Gamma; \sigma \cdot x \vdash_t P : \ell'$ and $I_{\mathcal{L}}(\ell') \subseteq I_{\mathcal{L}}(\ell \& (\sigma \cdot x)^t)$.*

Let $\Gamma\{\bar{u}/\bar{x}\}$ be defined as follows:

$$\Gamma\{u_1, \dots, u_n / x_1, \dots, x_n\}(z) = \begin{cases} \Gamma(x_i) & \text{if } z = u_i, 1 \leq i \leq n \\ \Gamma(z) & \text{if } z \notin u_1, \dots, u_n \text{ and } z \notin x_1, \dots, x_n \end{cases}$$

In the above definition, the topmost alternative applies before the other one. Therefore, for example, $\Gamma\{x/\bar{x}\} = \Gamma$ and $\Gamma\{u, v/v, u\}$ simply swaps the mappings of u and v , e.g. $\Gamma\{u, v/v, u\}(u) = \Gamma(v)$ and $\Gamma\{u, v/v, u\}(v) = \Gamma(u)$.

Lemma A.2 (Substitution lemma). *Let $\Gamma; \sigma \vdash_t P : \ell$ and $t' = t\{\bar{u}/\bar{x}\}$. Then $\Gamma\{\bar{u}/\bar{x}\}; \sigma\{\bar{u}/\bar{x}\} \vdash_{t'} P\{\bar{u}/\bar{x}\} : \ell\{\bar{u}/\bar{x}\}$. Similarly for $\Gamma \vdash e : \mathbb{T}$.*

Lemma A.3 (Type preservation). *Let $\Gamma \vdash (C_1 \ \mathcal{D}_1, \dots, C_n \ \mathcal{D}_n, P) : (\mathcal{L}, \ell)$ and $P \bullet \varepsilon \longrightarrow^* \mathbb{P}'$. Then there exist ℓ' and Γ' such that $\Gamma + \Gamma' \vdash \mathbb{P}' : \ell'$ and $I_{\mathcal{L}}(\ell') \subseteq I_{\mathcal{L}}(\ell)$.*

Proof. This follows by induction on the derivation of $P \bullet \varepsilon \longrightarrow^* \mathbb{P}'$, with case analysis on the last rule used, let it be $\mathbb{P} \longrightarrow \mathbb{P}'$. The only non-trivial cases are (SYNC) and (CALL).

- Case (SYNC): In this case, we have

$$\mathbb{P} = \text{sync}(x)\{P'\} \bullet \sigma \mid \mathbb{P}'' \quad x \notin \mathbb{P}'' \quad \mathbb{P}' = P' \bullet \sigma \cdot x \mid \mathbb{P}''$$

$$\Gamma + \Gamma'' \vdash \mathbb{P} : \ell''$$

where $\mathbb{P}'' = \prod_{i \in 1..n} P_i \bullet \sigma_i$.

By the condition $\Gamma + \Gamma'' \vdash \mathbb{P} : \ell''$ and rules (T-STATE) and (T-THREAD) we have

$$\Gamma + \Gamma'', t_0 : \text{Thread}; \sigma \vdash_{t_0} \text{sync}(x)\{ P' \} : \ell_0 \quad (\text{A.1})$$

$$\left(\Gamma + \Gamma'' \vdash P_i \bullet \sigma_i : \ell_i \right)^{i \in 1..n} \quad (\text{A.2})$$

$$\ell'' = (\nu t_0) \ell_0 \& (\&_{i \in 1..n} \ell_i) \quad (\text{A.3})$$

By using (T-SYNC), from (A.1) we derive

$$\Gamma + \Gamma'', t_0 : \text{Thread}; \sigma \cdot x \vdash_{t_0} P' : \ell_0$$

By (T-THREAD) applied to the foregoing judgment we derive

$$\Gamma + \Gamma'' \vdash P' \bullet \sigma \cdot x : (\nu t_0) \ell_0$$

Then, by applying (T-STATE) to this last judgment and (A.2), we have the required result (with $\ell' = \ell$).

- Case (CALL): We assume $C \mathcal{D}$ is in the program and $m(\overline{C'} \overline{z}, \overline{\text{int}} \overline{y}) = P_{C.m}$ is in \mathcal{D} . We have

$$\mathbb{P} = x.m(\overline{u}, \overline{e}) \bullet \sigma \mid \mathbb{P}'' \quad \llbracket \overline{e} \rrbracket = \overline{v} \quad \mathbb{P}' = P_{C.m}\{x, \overline{u}, \overline{v} / \text{this}, \overline{y}, \overline{z}\} \bullet \sigma \mid \mathbb{P}''$$

$$\Gamma + \Gamma'' \vdash \mathbb{P} : \ell''$$

where $\mathbb{P}'' = \prod_{i \in 1..n} P_i \bullet \sigma_i$. By the condition $\Gamma + \Gamma'' \vdash \mathcal{S} : \ell''$ and rules (T-STATE) and (T-THREAD) we have

$$\Gamma + \Gamma'', t_0 : \text{Thread}; \sigma \vdash_{t_0} x.m(\overline{u}, \overline{e}) : \ell_0 \quad (\text{A.4})$$

$$\left(\Gamma + \Gamma'' \vdash P_i \bullet \sigma_i : \ell_i \right)^{i \in 1..n}$$

$$\ell'' = (\nu t_0) \ell_0 \& (\&_{i \in 1..n} \ell_i)$$

Let $\sigma = \sigma' \cdot x'$. By (T-CALL) applied to (A.4) we have that $\ell_0 = \mathfrak{f}_{C.m}(t_0, x', x, \overline{u}) \& (\sigma)^{t_0}$. By the hypothesis $\Gamma \vdash (C_1 \mathcal{D}_1, \dots, C_n \mathcal{D}_n, P) : (\mathcal{L}, \ell)$ and (T-PROG) we derive

$$\Gamma, \text{this} : C, \overline{y} : \overline{C'}, \overline{z} : \overline{\text{int}}, t' : \text{Thread}, z' : \text{Thread}; z' \vdash_{t'} P_{C.m} : \ell_{C.m} \quad (\text{A.5})$$

$$\mathfrak{f}_{C.m}(t', z', \text{this}, \overline{y}) = \ell_{C.m} \in \mathcal{L}$$

By applying the Substitution Lemma A.2 to (A.5) we have

$$\Gamma''' = \left(\Gamma, \text{this} : C, \overline{y} : \overline{C'}, \overline{z} : \overline{\text{int}}, t' : \text{Thread}, x' : \text{Thread} \right) \{x, \overline{u}, t_0, x' / \text{this}, \overline{y}, t', z'\}$$

$$\Gamma'''; x' \vdash_{t_0} P_{C.m}\{x, \overline{u}, \overline{v} / \text{this}, \overline{y}, \overline{z}\} : \ell_{C.m}\{x, \overline{u}, t_0, x' / \text{this}, \overline{y}, t', z'\}$$

Then we notice that $\Gamma''' \subseteq \Gamma + \Gamma'', t_0 : \text{Thread}$. Therefore, by (T-THREAD) and (T-STATE) we derive

$$\Gamma + \Gamma'', t_0 : \text{Thread}; x' \vdash P_{C.m}\{x, \overline{u}, \overline{v} / \text{this}, \overline{y}, \overline{z}\} : \ell_{C.m}\{x, \overline{u}, t_0, x' / \text{this}, \overline{y}, t', z'\}$$

and, by Proposition A.1,

$$\Gamma + \Gamma'', t_0 : \text{Thread}; \sigma' \cdot x' \vdash P_{C.m}\{x, \overline{u}, \overline{v} / \text{this}, \overline{y}, \overline{z}\} : \ell'_0$$

$$I_{\mathcal{L}}(\ell'_0) \subseteq I_{\mathcal{L}}(\ell_{C.m}\{\overline{u}, t_0, x' / \overline{y}, t', z'\} \& (\sigma' \cdot x')^{t_0}) \quad (\text{A.6})$$

(we recall that $\sigma = \sigma' \cdot x'$). By applying again (T-THREAD) and (T-STATE) we derive

$$\Gamma + \Gamma'' \vdash \mathbb{P}' : \ell'$$

where $\ell' = ((\nu t_0) \ell'_0) \& (\&_{i \in 1..n} \ell_i)$. By definition of $I_{\mathcal{L}}$ and (A.6), we conclude $I_{\mathcal{L}}(\ell') \subseteq I_{\mathcal{L}}(\ell)$. \square

Lemma A.4 (Correctness). Let $\Gamma \vdash (C_1 \mathcal{D}_1, \dots, C_n \mathcal{D}_n, P) : (\mathcal{L}, \ell)$ such that $I_{\mathcal{L}}(\ell)$ has no circularity. If

$$P \bullet \varepsilon \longrightarrow^* \mathbb{P} \quad \text{and} \quad \mathbb{P} = \text{sync}(x_0)\{ P_0 \} \bullet \sigma_0 \mid \mathbb{P}'$$

then there exists \mathbb{P}'' such that $\mathbb{P} \longrightarrow \mathbb{P}''$.

Proof. If \mathbb{P}' contains object creation or forks or conditionals or method invocations at the top level, then the required property immediately follows. Similarly if it contains threads like $0 \bullet \sigma$. Thus, we can assume that

$$\mathbb{P}' = \prod_{i \in 1..m} \text{sync}(x_i) \{ P_i \} \bullet \sigma_i .$$

We notice that $m \geq 1$, otherwise rule (SYNC) is possible. We also assume that, for every j , there is $i_j \neq j$ such that $x_j \in \sigma_{i_j}$, otherwise (SYNC) is also possible. We demonstrate a contradiction: $I_{\mathcal{L}}(\ell)$ must have a circularity.

By Lemma A.3, there are Γ' and ℓ' such that $\Gamma + \Gamma' \vdash \mathbb{P} : \ell'$ and $I_{\mathcal{L}}(\ell')$ has no circularity. Because of the form of \mathbb{P} , by applying (T-STATE) and (T-THREAD), in order, we have

$$\left(\Gamma + \Gamma', t_i : \text{Thread}; \sigma_i \vdash_{t_i} \text{sync}(x_i) \{ P_i \} : \ell_i \right)^{i \in 0..m} \quad \ell' = \&_{i \in 0..m} (\nu t_i) \ell_i$$

Consider the thread $\text{sync}(x_0) \{ P_0 \} \bullet \sigma_0$ and let $x_0 \in \sigma_1$. For simplicity, assume that $x_1 \in \sigma_0$.

By (T-SYNC) and Proposition A.1, we have $I_{\mathcal{L}}((\sigma_0 \cdot x_0)^{t_0}) \in I_{\mathcal{L}}(\ell_0)$ and $I_{\mathcal{L}}((\sigma_1 \cdot x_1)^{t_1}) \in I_{\mathcal{L}}(\ell_1)$. Since, according to the assumptions, $(x_1, x_0)_{t_0} \in I_{\mathcal{L}}((\sigma_0 \cdot x_0)^{t_0})$ and $(x_0, x_1)_{t_1} \in I_{\mathcal{L}}((\sigma_1 \cdot x_1)^{t_1})$, by definition of ℓ' and of $I_{\mathcal{L}}$, we derive $\{(x_1, x_0)_{t_0}, (x_0, x_1)_{t_1}\} \in I_{\mathcal{L}}(\ell')$ (without loss of generality, we assuming t_0 and t_1 are replaced by identical names). Therefore $I_{\mathcal{L}}(\ell')$ contains a circularity and, by Theorem 2.8, $I_{\mathcal{L}}(\ell)$ must also have a circularity. This contradicts the hypothesis. The general case, when the chain $x_0 \in \sigma_1, x_1 \in \sigma_2, \dots, x_h \in \sigma_0$, is similar. \square

References

- [1] Martin Abadi, Cormac Flanagan, Stephen N. Freund, Types for safe locking: static race detection for Java, *ACM Trans. Program. Lang. Syst.* 28 (2006) 207–255.
- [2] Robert Atkey, Donald Sannella, Threadsafe: static analysis for Java concurrency, *Electron. Commun. ECEASST* 72 (2015).
- [3] Saddek Bensalem, Klaus Havelund, Dynamic deadlock analysis of multi-threaded programs, in: *Hardware and Software Verification and Testing*, in: LNCS, vol. 3875, Springer, 2005, pp. 208–223.
- [4] Chandrasekhar Boyapati, Robert Lee, Martin Rinard, Ownership types for safe programming: preventing data races and deadlocks, in: *OOPSLA*, ACM, 2002, pp. 211–230.
- [5] B.A. Davey, H.A. Priestley, *Introduction to Lattices and Order*, Cambridge University Press, 2002.
- [6] Mahdi Eslamimehr, Jens Palsberg, Sherlock: scalable deadlock detection for concurrent programs, in: *Proc. 22nd International Symposium on Foundations of Software Engineering (FSE-22)*, ACM, 2014, pp. 353–365.
- [7] Cormac Flanagan, Shaz Qadeer, A type and effect system for atomicity, in: *PLDI*, ACM, 2003, pp. 338–349.
- [8] Antonio Flores-Montoya, Elvira Albert, Samir Genaim, May-happen-in-parallel based deadlock analysis for concurrent objects, in: *FORTE/FMOODS 2013*, in: LNCS, vol. 7892, Springer, 2013, pp. 273–288.
- [9] Abel Garcia, *Static Analysis of Concurrent Programs Based on Behavioral Type Systems*, PhD thesis, School in Computer Science and Engineering, 2017. Available at JaDA.cs.unibo.it.
- [10] Abel Garcia, Cosimo Laneve, JaDA – the Java deadlock analyser, in: *Behavioural Types: From Theories to Tools*, River Publishers, 2017, pp. 169–192.
- [11] Elena Giachino, Naoki Kobayashi, Cosimo Laneve, Deadlock analysis of unbounded process networks, in: *CONCUR 2014 – Concurrency Theory – 25th International Conference*, in: *Lecture Notes in Computer Science*, vol. 8704, Springer, 2014, pp. 63–77.
- [12] Elena Giachino, Cosimo Laneve, Deadlock detection in linear recursive programs, in: *Proceedings of SFM-14:ESM*, in: LNCS, vol. 8483, Springer-Verlag, 2014, pp. 26–64.
- [13] Elena Giachino, Cosimo Laneve, Michael Lienhardt, A framework for deadlock detection in core ABS, *Softw. Syst. Model.* 15 (4) (2016) 1013–1048.
- [14] Einar Broch Johnsen, Reiner Hähnle, Jan Schäfer, Rudolf Schlatte, Martin Steffen, ABS: a core language for abstract behavioral specification, in: LNCS, vol. 6957, Springer-Verlag, 2011, pp. 142–164.
- [15] Neil D. Jones, Lawrence H. Landweber, Y. Edmund Lien, Complexity of some problems in Petri nets, *Theor. Comput. Sci.* 4 (3) (1977) 277–299.
- [16] Naoki Kobayashi, A new type system for deadlock-free processes, in: *CONCUR*, in: LNCS, vol. 4137, Springer, 2006, pp. 233–247.
- [17] Naoki Kobayashi, Cosimo Laneve, Deadlock analysis of unbounded process networks, *Inf. Comput.* 252 (2017) 48–70.
- [18] Cosimo Laneve, A lightweight deadlock analysis for programs with threads and reentrant locks, in: *Proc. of 22nd Int. Symposium on Formal Methods (FM 2018)*, in: *Lecture Notes in Computer Science*, vol. 10951, Springer, 2018, pp. 608–624.
- [19] Cosimo Laneve, Abel Garcia, Deadlock detection of Java bytecode, in: *Proc. of 27th Int. Symposium on Logic-Based Program Synthesis and Transformation (LOPSTR 2017)*, in: *Lecture Notes in Computer Science*, vol. 10855, Springer, 2017, pp. 37–53.
- [20] Cosimo Laneve, Luca Padovani, Deadlock analysis of wait-notify coordination, cs.unibo.it/~laneve/papers/DAoWNC.pdf, January 2019.
- [21] Robin Milner, *A Calculus of Communicating Systems*, LNCS, vol. 92, Springer, 1980.
- [22] Robin Milner, Joachim Parrow, David Walker, A calculus of mobile processes, II, *Inf. Comput.* 100 (1992) 41–77.
- [23] Mayur Naik, Chang-Seo Park, Koushik Sen, David Gay, Effective static deadlock detection, in: *Proc. 31st International Conference on Software Engineering (ICSE 2009)*, ACM, 2009, pp. 386–396.
- [24] Martin Odersky, et al., *An Overview of the Scala Programming Language*, Technical Report IC/2004/64, EPFL, Lausanne, Switzerland, 2004.
- [25] Wolfgang Reisig, *Petri Nets: An Introduction*, Springer-Verlag, Berlin, Heidelberg, 1985.
- [26] John C. Reynolds, Separation logic: a logic for shared mutable data structures, in: *Proc. of LICS 2002*, IEEE Computer Society, 2002, pp. 55–74.
- [27] Kohei Suenaga, Type-based deadlock-freedom verification for non-block-structured lock primitives and mutable references, in: *APLAS*, in: LNCS, vol. 5356, Springer, 2008, pp. 155–170.
- [28] Vasco Thudichum Vasconcelos, Francisco Martins, Tiago Cogumbreiro, Type inference for deadlock detection in a multithreaded polymorphic typed assembly language, in: *PLACES*, in: *EPTCS*, vol. 17, 2009, pp. 95–109.