SOFTWARE TESTING, VERIFICATION AND RELIABILITY
Softw. Test. Verif. Reliab. 0000; **00**:1–53
Published online in Wiley InterScience (www.interscience.wiley.com). DOI: 10.1002/stvr

Object-Sensitive Cost Analysis for Concurrent Objects

Elvira Albert[1*], Puri Arenas[1], Jesús Correas[1], Samir Genaim[1],
Miguel Gómez-Zamalloa[1], Germán Puebla[2], Guillermo Román-Díez[2]

[1] *DSIC, Complutense University of Madrid*
[2] *DLSIIS, Technical University of Madrid*

SUMMARY

This article presents a novel cost analysis framework for *concurrent objects*. Concurrent objects form a well established model for distributed concurrent systems. In this model, objects are the concurrency *units* which communicate among them via *asynchronous* method calls. Cost analysis aims at automatically approximating the *resource consumption* of executing a program in terms of its input parameters. While cost analysis for sequential programming languages has received considerable attention, concurrency and distribution have been notably less studied. The main challenges of cost analysis in a concurrent setting are: (1) Inferring precise *size abstractions* for data in the program in the presence of shared memory. This information is essential for bounding the number of iterations of loops. (2) Distribution suggests that analysis must infer the cost of the diverse distributed components separately. We handle this by means of a novel form of *object-sensitive recurrence equations* which use *cost centers* in order to keep the resource usage assigned to the different components separate. We have implemented our analysis and evaluated it on several small applications which are classical examples of concurrent and distributed programming.

Received . . .

KEY WORDS: static analysis; resource usage analysis; cost analysis; concurrency; distributed systems; object-sensitive analysis

1. INTRODUCTION

Distribution and concurrency are currently mainstream. The Internet and the broad availability of multi-processors radically influence software. Many standard desktop programs have to deal with distribution aspects like network transmission delay and failure. Furthermore, many chip manufacturers are turning to multicore processor designs as a way to increase performance in desktop, enterprise, and mobile processors. This brings renewed interest in developing both new concurrency models and associated programming languages techniques that help in understanding, analyzing, and verifying the behaviour of concurrent and distributed programs.

One of the most important features of a program is its resource consumption. By resource, we mean not only traditional cost measures (e.g., number of executed instructions, or memory consumption) but also concurrency-related measures (e.g., number of tasks spawned, number of requests to remote servers). Example 1 will illustrate these types of resources on a fragment of

*Correspondence to: **elvira@sip.ucm.es** (Elvira Albert)

Contract/grant sponsor: This work was funded partially by the EU project FP7-ICT-610582 ENVISAGE: Engineering Virtualized Services (http://www.envisage-project.eu), by the Spanish project TIN2012-38137, and by the CM project S2013/ICE-3006.

our running example. Cost analysis (a.k.a. *resource usage* analysis) aims at *statically* inferring approximations of the resource consumption of executing the program. Automatically inferring the resource usage of concurrent programs is challenging because of the inherent complexity of concurrent behaviours.

In addition to traditional applications, like optimization [45], verification and certification of resource consumption [18], cost analysis opens up interesting applications in the context of concurrent programming. In general, having anticipated knowledge on the resource consumption of the different components which constitute a system, is useful for distributing the load of work. Upper bounds can be used to predict that one component may receive a large amount of remote requests, while other siblings are idle most of the time. Also, our framework allows instantiating the upper bounds obtained for different components with the particular features of the infrastructure on which they are deployed. For instance, we might know that the processing capacity of one component doubles the processing capacity of another component. Then, the upper bounds that we obtain for the different components can be compared by taking into account such particular features of the infrastructure. E.g., the number of steps obtained for the former component should be divided by two if we want to compare the processing load of this component w.r.t. the load of the latter component. Then, analysis can be used to detect the components that consume more resources and may introduce bottlenecks. Lower bounds on the resource usage can be used to decide if it is worth executing locally a task or requesting remote execution.

In order to develop our analysis, we consider a concurrency model based on the notion of concurrently running (groups of) objects, similar to the actor-based and active-objects approaches [38, 42, 1, 44, 31, 37, 15]. These models take advantage of the concurrency implicit in the notion of object in order to provide programmers with high-level concurrency constructs that help in producing concurrent applications more modularly and in a less error-prone way. Concurrent objects communicate via *asynchronous* method calls. Intuitively, each concurrent object is a monitor and allows at most one *active* process to execute within the object. Scheduling among the processes of an object is cooperative, i.e., a process has to release the monitor lock explicitly, except when it terminates. Each object has an unbounded set of pending processes. In case the lock of a concurrent object is free any process in the set of pending processes can grab the lock and start to execute (hence process scheduling is non-deterministic). The notation $f = o \,!\, m(\bar{e})$ is used to denote that an asynchronous call $m(\bar{e})$ has been posted on object o and f is a future variable which allows us to know if the execution of the asynchronous call has finished. In such case the result can be retrieved by means of a get operation $f.get$. Also, we can synchronize the execution with the asynchronous call by means of an await instruction, namely *await* $f?$ is used to check if the asynchronous call has finished, otherwise the processor can be released such that a task which was pending to execute can take it.

Example 1 (notion of resource)**.** Let us consider the concurrent method (which is part of our running example that will be showed later) below.

```
Int process(Int pos) {
  Fut⟨Int⟩ f;
  Int i = 0;
  Int res = 0;
  while (i < elems) {
    f = this ! hdRead(pos + i);
    await f?;
    res = this.update(res,f.get);
    i = i + 1;
  }
  return res;
}
```

In the above method, elems is a class field. Our objective is to measure the resource consumption of executing the above method. One crucial aspect will be to find out if field elems can be modified at the *await* f? instruction. Observe that if the processor is released at the *await* and another process that increases the value of elems takes the processor, the loop above might not terminate. Our method relies on *class invariants* which contain information on the shared

DOI: 10.1002/stvr

memory at processor release points. By assuming that elems is not modified, we can now consider different types of resources of interest. For instance, we can measure traditional cost measures like number of executed instructions or memory consumption. In the former case, we will infer that $3 + \mathsf{elems} * (7 + hdRead_{inst} + update_{inst})$ instructions will be executed. In the expression, elems corresponds to the maximum number of iterations of the loop, at each iteration 7 instructions are executed (the loop condition, the two increments, two method invocations, the *get* and the *await*) and the number of instructions of executing methods hdRead and update, denoted $hdRead_{inst}$ and $update_{inst}$ respectively, are added as well. For simplicity, we ignore here the values of the parameters. Besides, 3 instructions are executed outside the loop. The analysis will need to infer also the cost of the methods hdRead and update and plug them in the expression above. As regards memory consumption, new memory is not created in the loop, hence we would output an expression of the form $\mathsf{elems} * (hdRead_{heap} + update_{heap})$ which relies on the memory created by the invoked methods (denoted $hdRead_{heap}$ and $update_{heap}$). Interestingly, we can also consider concurrency-related measures like the number of tasks spawned. In the above method, we spawn one task directly (but the method invocations might also spawn new tasks transitively). Therefore, we will compute an expression of the form $\mathsf{elems} * (1 + hdRead_{tasks} + update_{tasks})$ which relies on the number of tasks spawned by the calls to hdRead and update. Observe that the types of resources we have considered are platform independent (unlike WCET or energy consumption), i.e., they can be inferred by inspecting the program, and the hardware on which the program will be executed can be ignored. Platform dependent resources are beyond the scope of this work. ■

1.1. Contributions

We propose a static cost analysis for concurrent objects, which is parametric w.r.t. the notion of resource that can be instantiated to measure both traditional and concurrency-related resources. The main contributions of this work are:

1. We present a flow-sensitive object-sensitive points-to analysis for concurrent programs which adapts Milanova's analysis framework [35] for Java to the concurrent object setting;
2. We introduce a sound size analysis for concurrent execution. The analysis is *field-sensitive*, i.e., it tracks data stored in the heap whenever it is sound to do so; the accuracy of the field-sensitive size analysis can be increased by means of *class invariants* [34] which contain information on the shared memory;
3. We leverage the definition of cost used in sequential programming to the distributed setting by relying on the notion of *cost centers* [36], which represent the (distributed) components and allow separating their costs;
4. We present a novel form of *object-sensitive* recurrence relations which relies on information gathered by the previous object-sensitive points-to analysis in order to generate the cost equations. Interestingly, the resulting recurrence relations can still be solved to closed-form upper/lower bounds using standard solvers for cost analysis of sequential programs;
5. We report on the SACO system, a prototype implementation of a cost analyzer for programs written in ABS [30] (an Abstract Behavioural Specification language based on concurrent objects).

It is recognized that performing the analysis on a high-level concurrency model, like the concurrent objects model, makes verification more feasible. This is because analysis in concurrent systems often needs to consider too many interleavings and thus ends up being limited to very small programs in practice. We argue that our approach is of both practical and theoretical relevance.

This work is an extended and revised version of APLAS'11 [3]. Points 2 and 3 in the contribution list can be considered as original contributions of the conference, while this journal paper has points 1 and 4 as original contributions, and also with an implementation of them (last point). However, the treatment of fields in the size analysis (included in point 2 in the contribution list) has been improved in this article.

DOI: 10.1002/stvr

1.2. Organization of the Article

The remainder of the article is organized as follows. Section 2 presents the syntax and semantics of the language on which we develop our analysis. Section 3 defines the notion of cost for the concurrent distributed programs that we aim at approximating by means of the resource analysis.

The next three sections present the resource-analysis framework in several steps. Our starting point is a powerful cost analysis framework for sequential OO programs [7]. When lifting such framework to the concurrent and distributed setting, there are two main difficulties and novelties. First, it is widely recognized that, due to the possible interleaving between tasks, tracking values of data stored in the heap is challenging [14, 33]. In Section 5, we present the basic, novel, *field-sensitive* size analysis for the concurrent setting.

The second difficulty is related to the fact that standard recurrence relations (in the sequential setting) assume a single cost center which accumulates the cost of the whole execution. We propose a novel form of *recurrence relations* which use *cost centers* to split the cost of the diverse distributed components. This requires first the inference of object-sensitive points-to information which approximates the set of objects which each reference variable may point. Section 4 adapts the object-sensitive points-to analysis of Milanova [35, 40] to our setting. Then, the points-to information gathered by the analysis allows us to define in Section 6 object-sensitive recurrence relations which, together with the size abstractions, constitute the core of our analysis.

Section 7 presents SACO, a prototype implementation of our analysis, and evaluates it on a series of typical applications of concurrent and distributed programming. Finally, Section 8 reviews the related work and Section 9 recaps the main conclusions.

2. A LANGUAGE WITH CONCURRENT OBJECTS

The concurrency model of Java and C# is based on threads that share memory and are scheduled preemptively, i.e., they can be suspended or activated at any time. To avoid undesired interleavings, low-level synchronization mechanisms such as locks have to be used. Thread-based programs are error-prone, difficult to debug, verify and maintain. In order to overcome these problems, several higher-level concurrency models that take advantage of the inherent concurrency implicit in the notion of object have been developed [38, 42, 29, 19, 34]. They provide simple language extensions that allow programming concurrent applications with relatively little effort. Concurrent objects [29, 19] form today a well established high-level model for distributed concurrent systems.

2.1. The Concurrency Model

For the sake of concreteness, we develop our analysis on a simple imperative language with concurrent objects, which is the imperative subset of the ABS language [30]. However, our techniques work for other languages that use actors (e.g., there are implementations of actor libraries for Scala, Java, Erlang, among others). The central concept of this concurrency model is that of *concurrent object*. Conceptually, each object has a dedicated processor and encapsulates a *local heap* which is not accessible from outside this object, i.e., fields are always accessed using the *this* object, and any other object can only access such fields through method calls. Concurrent objects live in a distributed environment with asynchronous and unordered communication by means of asynchronous method calls. Thus, an object has a set of tasks (i.e., calls) to execute and, among them, at most one task is *active* and the others are *suspended* on a task queue.

Asynchronous method calls may be seen as triggers of concurrent activity, spawning new tasks (so-called processes) in the called object. After asynchronously calling method m of object o with arguments $\overline{e}$, denoted by $f = o\,!\,m(\overline{e})$, the caller may proceed with its execution without blocking on the call. Here f is a *future variable* which refers to a return value which has yet to be computed. There are two operations on future variables, which control external synchronization. First, *await* f? suspends the active task (allowing other tasks in the object to be scheduled) until the future variable f has been assigned a value. Second, the value stored in f can be retrieved using f.*get*, which blocks all execution in the object until f gets a value (in case it has not been assigned a value yet).

DOI: 10.1002/stvr

```
class FileIS(String fp, Int lth, Int blockS) {
  Int readBlock () {
    Int res = 0; Int i = this.lth;
    Int incr = 0; Int pos = 0;
    while (i > 0) {
      if (this.blockS > i) incr = i;
      else incr = this.blockS;
      Fut⟨Int⟩ f;
      f = this ! readContent(pos,incr);
      await f?;
      res = res + f.get;
      i = i - incr;
      pos = pos + incr;
    }
    return res;
  }
  Int readOnce()  {
    Fut⟨Int⟩ f = this ! readContent(0,this.lth);
    await f?;
    return f.get;
  }
  Int readContent(Int pos, Int elems) {
    Reader rd = new Reader (this.fp,elems);
    Fut⟨Int⟩ f = rd ! process(pos);
    await f?;
    return f.get;
  }
}// end class FileIS
```

```
class Reader(String fp, Int elems) {
  Int hdRead(Int i){ ··· }
  Int update(Int a, Int b){ ··· }
  Int process(Int pos) {
    Fut⟨Int⟩ f;
    Int i = 0;
    Int res = 0;
    while (i < this.elems) {
      f = this ! hdRead(pos + i);
      await f?;
      res = this.update(res,f.get);
      i = i + 1;
    }
    return res;
  }
}// end class Reader

main {
  FileIS ob1 = new FileIS("A.txt",20,2);
  FileIS ob2 = new FileIS("A.txt",20,3);
  Fut⟨Int⟩ f1; Fut⟨Int⟩ f2;
(*)f1 = ob1 ! readOnce();
  f2 = ob2 ! readBlock();
  await f1?;
  Int r1 = f1.get;
  await f2?;
  Int r2 = f2.get;
}
```

Figure 1. Running example

Example 2 (syntax of ABS)**.** Figure 1 shows the source code of our running example which implements a simple *file input stream* (defined in class FileIS) that provides two different ways of processing a file. The class contains three fields (defined as class parameters) which represent, respectively, the name of the file fp, the length of the file lth, and the size of the block to be read from the field blockS. Method readBlock reads file fp block by block (of sizes blockS) and sums the values retrieved using *get*. Method readOnce reads the whole file in just one invocation to readContent. The latter method invokes method process of class Reader which reads and processes elems elements of the file starting at position pos. Method hdRead represents the low-level access to the hard-disk and method update performs some arithmetic operation on its arguments and returns an integer value. We do not show the code of these methods as they are not relevant for the purpose of this article. ■

2.2. A Rule-based Intermediate Language

To contextualize the formalization of the analysis in a simpler model, we develop our analysis on an intermediate representation (IR) similar to those for Java bytecode and .NET [43, 7, 41, 20]. In the IR, *recursion* is the only iterative mechanism and *guards* are the only form of conditional. In the following, given any entity z, we use $\bar{z}$ to denote the tuple $\langle z_1, \ldots, z_n \rangle$. The compilation of a program into the IR is done by building the CFG for the original program and representing each block in the CFG by means of a rule. The following definition establishes the formal syntax of the IR.

Definition 1 (syntax of IR programs)**.** *A* program *in the IR consists of a set of classes* $\bar{C}$*. Each* class C *contains a set of fields* $\bar{f}_C$ *and a set of procedures. A* procedure m *of class* C *is defined by a set of guarded rules. A* guarded rule *for* m *has the form "*$r \equiv C.m(this, \bar{x}, \bar{y}) \leftarrow g, \bar{b}$*.", where* $C.m(this, \bar{x}, \bar{y})$ *is the* head *of the rule,* $this$ *is the identifier of the object on which the method is*

DOI: 10.1002/stvr

executing, g specifies the conditions for the rule to be applicable, and, $\bar{b}$ is the rule's body. Guards g *and* instructions $b \in \bar{b}$ *are defined according to the following grammar:*

$$
\begin{array}{lll}
b & ::= x := rhs \mid this.f := y \mid \textit{await}\ x? \mid \textit{call}(ct, m(rec, \bar{x}, \bar{y})) & \\
g & ::= true \mid x\ op_R\ y & \\
rhs & ::= e \mid \textit{new}\ C \mid x.\textit{get} & \\
e & ::= \textit{null} \mid this.f \mid x \mid n \mid x\ op_A\ y & \\
op_R & ::= <\mid>\mid=\mid\neq\mid\geq\mid\leq & op_A ::= + \mid - \mid / \mid * \qquad ct ::= \textit{m} \mid \textit{b}
\end{array}
$$

where x and y denote variable names*, f a* field name*, $\textit{call}(ct, m(rec, \bar{x}, \bar{y}))$ a call to a method or a block and $n \in \mathbb{Z}$.*

The first argument ct of a call $\textit{call}(ct, m(rec, \bar{x}, \bar{y}))$ can be either *m* or *b*. The identifier *m* is used for asynchronous method calls whereas *b* is used for synchronous method calls or calls to intermediate blocks. For instance, intermediate blocks can correspond to if-then-else statements or loops; rec is a variable that refers to the receiver object. For synchronous method calls or calls to blocks, rec is always $this$; the variables $\bar{x}$ (respectively $\bar{y}$) are the formal parameters (respectively return values). For methods, $\bar{y}$ is either empty $\langle\rangle$ or contains a single output variable $\langle y \rangle$.

An instruction $\mathtt{x} = \textit{new}\ \mathtt{C}(\bar{\mathtt{t}})$ in the target language, is represented in the IR by $x := \textit{new}\ C$ followed by a call to the class constructor with the corresponding parameters $\bar{t}$. For example, assuming that the class C contains a field f of type integer, the instruction $\mathtt{x} = \textit{new}\ \mathtt{C}(2)$ will be translated into $x := \textit{new}\ C, y := 2; C\_init(this, \langle y \rangle, \langle\rangle)$, where $C\_init(this, \langle y \rangle, \langle\rangle) \leftarrow this.f := y$.

The translation from the high-level programs to the IR is (almost) identical to the translation of Java (bytecode) to the IR in [7], where classes and fields in the IR are the same as in the original program and each method m of a class C is represented in the IR by a single procedure named $C.m$ (the method entry). The other rules in the IR are intermediate procedures used only within the method, with $ct = b$. The *main* method does not belong to any class. Without lack of generality, we assume that method names are unique, and we omit C when referring to m. Furthermore, we will omit those guards which are $true$. This happens in the rules corresponding to method entries.

Example 3 (CFG of an IR for the running example)**.** Figure 2 depicts the IR (left) and the CFG (right) of method readBlock. Loops are extracted in separate CFGs to enable compositional cost analysis (e.g., the CFG at the bottom is the one for the *while* loop). The method is represented by four procedures, *readBlock*, *while*, *if* and *if_c*, which have a correspondence with blocks in the CFG and the entry to the loop. Each procedure is defined by means of guarded rules. As notation $\overline{inp}$ stands for $\langle res, i, incr, pos \rangle$ and $\overline{out}$ for $\langle res, i, incr, pos \rangle$. Guards in rules state the conditions under which the corresponding blocks in the CFG can be executed. When there is more than one successor in the CFG, we create a *continuation procedure* and the corresponding call in the rule. Blocks in the continuation will in turn be defined by means of (mutually exclusive) guarded rules. As a result of the translation, observe that all forms of iteration in the program are represented by means of *recursive* calls. The unique parameter of the procedure readBlock is the reference to the *this* object. When calling a block, we pass as arguments all local variables that are needed in the block. The heap remains implicit. ■

As we have seen in the example, in the IR there may be multiple rules with the same name, thus, to distinguish each different rule we use a unique identifier t. Similarly, when needed, in order to uniquely identify the different program points, we use a pair $\langle t, j \rangle$, where t is the identifier of the rule and j is the position of the instruction in the body of the rule.

2.3. Operational Semantics

An *object* is of the form $\mathtt{ob}(o, C, h, \langle tv, \bar{b} \rangle, \mathcal{Q})$, where o is the *object identifier*, C is its class name, h is its local *heap*, $\langle tv, \bar{b} \rangle$ is the *execution context* of the current task, being tv the *table of local variables* and $\bar{b}$ the sequence of instructions to be executed by the current task, and $\mathcal{Q}$ is the set of *pending tasks*, being each of them an execution context. In the following we use ϵ to denote either an empty sequence of instructions or an empty execution context. A heap h maps *field names* $\bar{f}_C$

DOI: 10.1002/stvr

$readBlock(this, \langle\rangle, \langle r\rangle) \leftarrow$
$res := 0, i := this.lth,$
$incr := 0, pos := 0,$
$\mathsf{call}(b, while(this, \overline{inp}, \overline{out})),$
$r := res.$
$while(this, \overline{inp}, \overline{out}) \leftarrow i \leq 0.$
$while(this, \overline{inp}, \overline{out}) \leftarrow i > 0,$
$\mathsf{call}(b, if(this, \overline{inp}, \overline{out})).$
$if(this, \overline{inp}, \overline{out}) \leftarrow this.blockS > i,$
$incr := i,$
$\mathsf{call}(b, if\_c(this, \overline{inp}, \overline{out})).$
$if(this, \overline{inp}, \overline{out}) \leftarrow this.blockS \leq i,$
$incr := this.blockS,$
$\mathsf{call}(b, if\_c(this, \overline{inp}, \overline{out})).$
$if\_c(this, \overline{inp}, \overline{out}) \leftarrow$
$\mathsf{call}(m, readContent(this, \langle pos, incr\rangle, \langle f\rangle)),$
$\mathsf{await}\ f?, v := f.\mathsf{get},$
$res := res + v, i := i - incr,$
$pos := pos + incr,$
$\mathsf{call}(b, while(this, \overline{inp}, \overline{out})).$

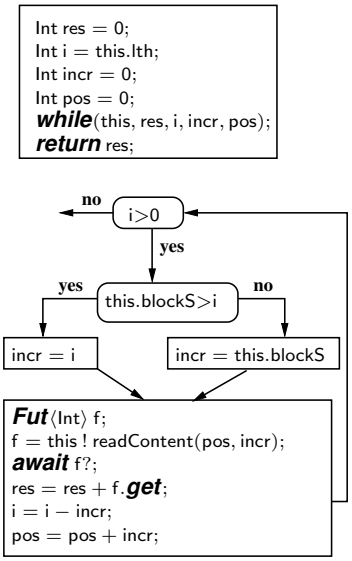

Figure 2. The IR and CFG for method readBlock

declared in C to $\mathbb{V} = \mathbb{Z} \cup \{\mathsf{null}\} \cup Objects$, where $Objects$ denotes the set of object identifiers. A table of variables tv maps local variables to $\mathbb{V}$. It contains the special entry `ret` to associate the return variable of a method to the corresponding future variable. *Future events* have the form $\mathtt{fut}(\mathtt{fn}, v)$ where $v \in \mathbb{V} \cup \{\bot\}$ and `fn` stands for a future variable identifier. The symbol $\bot$ indicates that `fn` does not have a value yet. For simplicity, we assume that all methods return a single value, while intermediate blocks will often have several return values. An *execution state* (or *configuration*) $\mathcal{S}$ has the form $\{a_1, \ldots, a_n\}$, where a_i can be either an object or a future event. Execution states are in fact represented as sets of objects and future events. In the following, we use the notation $\{a|\mathcal{S}\}$ to denote the set $\{a\} \cup \mathcal{S}$.

The *operational semantics* is given in a rewriting-based style, where, at each step, a subset of the state is rewritten according to the rules in Figure 3. Let us intuitively explain the semantics. Function $eval_e$ in rule 1 evaluates an expression e with respect to a heap h and a table of variables tv in the standard way. Note that the heap is required to evaluate expressions of the form $this.f$, that returns $h(f)$ as result. For the case of *null*, $eval_e$ returns *null*. Function $eval_{gd}(g, tv)$ in rule 4 returns $true$ iff $g \equiv true$ or $g \equiv x_1\ op_R\ x_2$ and $tv(x_1)\ op_R\ tv(x_2)$ holds. Finally, the evaluation of the conditions in *await* instructions is done by function $eval_{aw}$. In particular, in rules 9 and 10, this function behaves as follows: $eval_{aw}(x?, tv, \mathcal{S}) = true$ iff $tv(x) = \mathtt{fn}$, $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}$ and $v \neq \bot$. The notation $tv[x \mapsto v]$ (respectively $h[f \mapsto v]$) is used for storing v in the local variable x (respectively field f). In general, given any entity z, we use $vars(z)$ to denote the set of variables occurring in z.

Rules 1 and 2 operate in the expected way. In rule 3, it can be observed that the table of variables tv maps x to o_1. Function $newRef()$ is in charge of generating fresh object identifiers and procedure $newHeap(D, h_1)$ creates a new mapping for fields in D, where each field is initialized to either 0 or *null*. In rule 4, the notation $r \equiv p(this', \bar{x}', \bar{y}) \leftarrow g', b'_1, \ldots, b'_n \ll_p^{\langle \bar{y} \rangle} P$ stands for a renaming of a rule for p in P such that the output variables are mapped to $\bar{y}$, and the remaining variables in r are fresh. Formally the process first selects non-deterministically one rule $r'' \equiv p(this'', \bar{x}'', \bar{y}'') \leftarrow g'', b''_1, \ldots, b''_n$ in P. Afterwards, it defines a renaming ρ such that $\rho(\bar{y}'') = \bar{y}$ and $\rho(z)$ is a fresh variable, for all $z \in vars(r'') - \{\bar{y}''\}$, and finally it applies ρ to r'' resulting in r. Thus, a call to a block, in rule 4, is resolved by finding a matching rule, completely renamed except for the output variables, and adding its body to the sequence of instructions to be executed. The function $newEnv(vars(r) - \{\bar{y}\})$ creates a new mapping for variables in r except for $\bar{y}$ which remain the same, where each variable is initialized to either 0 or *null*. Furthermore,

DOI: 10.1002/stvr

$$(1)\ \frac{v = eval_e(e, h, tv)}{\{\mathtt{ob}(o, C, h, \langle tv, x := e \cdot \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\} \rightsquigarrow \{\mathtt{ob}(o, C, h, \langle tv[x \mapsto v], \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\}}$$

$$(2)\ \frac{v = tv(y)}{\{\mathtt{ob}(o, C, h, \langle tv, this.f := y \cdot \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\} \rightsquigarrow \{\mathtt{ob}(o, C, h[f \mapsto v], \langle tv, \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\}}$$

$$(3)\ \frac{o_1 = newRef(), newHeap(D, h_1)}{\begin{array}{c}\{\mathtt{ob}(o, C, h, \langle tv, x := \textsf{new}\ D \cdot \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\} \rightsquigarrow \\ \{\mathtt{ob}(o, C, h, \langle tv[x \mapsto o_1], \bar{b}\rangle, \mathcal{Q}), \mathtt{ob}(o_1, D, h_1, \epsilon, \emptyset)|\mathcal{S}\}\end{array}}$$

$$(4)\ \frac{\begin{array}{c}r \equiv p(this', \bar{x}', \bar{y}) \leftarrow g', b'_1, \ldots, b'_n \ll_p^{\langle \bar{y}\rangle} P, o \equiv tv(this), \\ tv_1 = newEnv(vars(r) - \{\bar{y}\}), tv_2 = tv_1[this' \mapsto o, \bar{x}' \mapsto tv(\bar{x})], eval_{gd}(g', tv_2) = true\end{array}}{\{\mathtt{ob}(o, C, h, \langle tv, \textsf{call}(\mathsf{b}, p(this, \bar{x}, \bar{y})) \cdot \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\} \rightsquigarrow \{\mathtt{ob}(o, C, h, \langle tv \cup tv_2, b'_1 \cdots b'_n \cdot \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\}}$$

$$(5)\ \frac{\begin{array}{c}r \equiv p(this', \bar{x}', y') \leftarrow b'_1, \ldots, b'_n \ll_p P, o_1 \equiv tv(rec), \mathtt{fn} = newFut(), \\ tv_2 = newEnv(vars(r)), tv_3 = tv_2[this' \mapsto o_1, \bar{x}' \mapsto tv(\bar{x}), \mathtt{ret} \mapsto (y', \mathtt{fn})]\end{array}}{\begin{array}{c}\{\mathtt{ob}(o, C, h, \langle tv, \textsf{call}(\mathsf{m}, p(rec, \bar{x}, y)) \cdot \bar{b}\rangle, \mathcal{Q}), \mathtt{ob}(o_1, D, h_1, \langle tv_1, \bar{b}_1\rangle, \mathcal{Q}')|\mathcal{S}\} \rightsquigarrow \\ \{\mathtt{ob}(o, C, h, \langle tv[y \mapsto \mathtt{fn}], \bar{b}\rangle, \mathcal{Q}), \mathtt{ob}(o_1, D, h_1, \langle tv_1, \bar{b}_1\rangle, \{\langle tv_3, b'_1 \cdots b'_n\rangle\} \cup \mathcal{Q}'), \mathtt{fut}(\mathtt{fn}, \bot)|\mathcal{S}\}\end{array}}$$

$$(6)\ \frac{\mathtt{ret} \notin dom(tv)}{\{\mathtt{ob}(o, C, h, \langle tv, \epsilon\rangle, \mathcal{Q})|\mathcal{S}\} \rightsquigarrow \{\mathtt{ob}(o, C, h, \epsilon, \mathcal{Q})|\mathcal{S}\}}$$

$$(7)\ \frac{\mathtt{ret} \in dom(tv), (y, \mathtt{fn}) = tv(\mathtt{ret}), v = tv(y)}{\{\mathtt{ob}(o, C, h, \langle tv, \epsilon\rangle, \mathcal{Q}), \mathtt{fut}(\mathtt{fn}, \bot)|\mathcal{S}\} \rightsquigarrow \{\mathtt{ob}(o, C, h, \epsilon, \mathcal{Q}), \mathtt{fut}(\mathtt{fn}, v)|\mathcal{S}\}}$$

$$(8)\ \frac{\mathtt{fn} = tv(y),\ v \neq \bot}{\{\mathtt{ob}(o, C, h, \langle tv, x := y.\textsf{get} \cdot \bar{b}\rangle, \mathcal{Q}), \mathtt{fut}(\mathtt{fn}, v)|\mathcal{S}\} \rightsquigarrow \{\mathtt{ob}(o, C, h, \langle tv[x \mapsto v], \bar{b}\rangle, \mathcal{Q}), \mathtt{fut}(\mathtt{fn}, v)|\mathcal{S}\}}$$

$$(9)\ \frac{eval_{aw}(x?, tv, \mathcal{S}) = true}{\{\mathtt{ob}(o, C, h, \langle tv, \textsf{await}\ x? \cdot \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\} \rightsquigarrow \{\mathtt{ob}(o, C, h, \langle tv, \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\}}$$

$$(10)\ \frac{eval_{aw}(x?, tv, \mathcal{S}) = false}{\{\mathtt{ob}(o, C, h, \langle tv, \textsf{await}\ x? \cdot \bar{b}\rangle, \mathcal{Q})|\mathcal{S}\} \rightsquigarrow \{\mathtt{ob}(o, C, h, \epsilon, \{\langle tv, \textsf{await}\ x? \cdot \bar{b}\rangle\} \cup \mathcal{Q})|\mathcal{S}\}}$$

$$(11)\ \frac{b \in \mathcal{Q}}{\{\mathtt{ob}(o, C, h, \epsilon, \mathcal{Q})|\mathcal{S}\} \rightsquigarrow \{\mathtt{ob}(o, C, h, b, \mathcal{Q} - \{b\})|\mathcal{S}\}}$$

Figure 3. Operational semantics

$tv_2 = tv_1[this' \mapsto o, \bar{x}' \mapsto tv(\bar{x})]$ defines a new mapping tv_2 as $tv_2(this') = o$, $tv_2(\bar{x}') = tv(\bar{x})$ and $tv_2(z) = tv_1(z)$ otherwise. Similarly, $tv \cup tv_2$ defines the following mapping: if $x \in dom(tv)$ then $(tv \cup tv_2)(x) = tv(x)$, and $(tv \cup tv_2)(x) = tv_2(x)$ otherwise. It is used to extend a local variable table with the new variables introduced by the block. As notation $dom(tv)$ stands for the set of variables on which tv is defined. When the execution of a block finishes (rule 6) the state is prepared to later apply rule 11 to select a new task from the queue. Note that since asynchronous calls introduce `ret` in the table of variables (rule 5) then the condition $\mathtt{ret} \notin dom(tv)$ in rule 6 ensures that the execution does not correspond to an asynchronous call but rather to a block or synchronous call.

Rule 5 deals with asynchronous method invocations. In this case $r \equiv p(this', \bar{x}', y') \leftarrow b'_1, \ldots, b'_n \ll_p P$ stands for a fresh renaming of a rule for p in P, i.e., "$\ll_p$" stands for "$\ll_p^{\langle\rangle}$". When an object o_1 calls a method $p(\bar{x})$, the information required to execute the call is stored in the queue of the object identified by o_1. Note that parameter passing is done in the construction of tv_3, where the entries for $this'$ and $\bar{x}'$ are assigned a local copy of the value of the actual parameters rec and $\bar{x}$, respectively. Objects are not directly passed as parameters. Instead we pass the corresponding object identifier, which is unique. Function $newFut()$ generates a fresh future variable identifier. Observe

Prepared using stvrauth.cls
DOI: 10.1002/stvr

that tv_3 has the special entry `ret` to store the relation between the future variable `fn` where the result is stored and the output parameter y'. This future variable is initially undefined, thus $\texttt{fut}(\texttt{fn}, \bot)$ is added to the state. When the method returns a value (rule 7), the entry `ret` is used to look for the corresponding future variable and $\bot$ is updated with the returned value.

Rule 9 checks if a future variable is ready. In such case the computation proceeds. Otherwise, in rule 10, the *await* task is introduced in the corresponding queue, and the processor is released. The instruction *get* blocks the execution until the future variable has a value in rule 8. In rule 11 another task is dequeued (because the current one has terminated or released the processor). Note that this rule is applicable after applying rules 6 and 7 which correspond, respectively, to the complete execution of a block and a method, and rule 10 in which the processor is released.

We assume that executions start from a *main* method. Thus, the *initial configuration* is of the form $\{\texttt{ob}(\mathsf{main}, \bot, \bot, \langle tv, \mathit{call}(b, \mathsf{main}(\mathit{this}, \langle\rangle, \langle\rangle))\rangle, \emptyset)\}$ where tv is an empty mapping. Abusing notation, we use $\bot$ to denote an empty heap and an undefined class. The execution then proceeds by applying non-deterministically the execution steps in Figure 3. It is non-deterministic both in task and object selection. The execution finishes in a *final configuration* in which all events are either future events or objects of the form $\texttt{ob}(o, C, h, \epsilon, \emptyset)$. Executions can be regarded as *traces* $\mathcal{T}$ of the form $\mathcal{S}_0 \rightsquigarrow \mathcal{S}_1 \rightsquigarrow \cdots \rightsquigarrow \mathcal{S}_n$.

Example 4 (a trace in the running example)**.** Consider the *main* method of the running example (Figure 1). After executing the constructors we reach a configuration with three objects:
$\{\texttt{ob}(\mathsf{main}, \bot, \bot, \langle tv_{\mathsf{main}}, \bar{b}\rangle, \emptyset), \texttt{ob}(ob1, \mathit{FileIS}, h_{ob1}, \epsilon, \emptyset), \texttt{ob}(ob2, \mathit{FileIS}, h_{ob2}, \epsilon, \emptyset)\}$
where $\bar{b}$ corresponds to the sequence of instructions from the mark ⊛ on in Figure 1. After processing both asynchronous calls (rule 5) consecutively, the new state takes the form:

$$\begin{array}{ll}\{ & \texttt{ob}(\mathsf{main}, \bot, \bot, \langle tv_{\mathsf{main}}[f1 \mapsto \texttt{fn}_1, f2 \mapsto \texttt{fn}_2], \bar{b}'\rangle, \emptyset),\\ & \texttt{ob}(ob1, \mathsf{FileIS}, h_{ob1}, \epsilon, \{\langle tv_{ob1}, \mathit{body}_{ob1}\rangle\}), \texttt{fut}(\texttt{fn}_1, \bot),\\ & \texttt{ob}(ob2, \mathsf{FileIS}, h_{ob2}, \epsilon, \{\langle tv_{ob2}, \mathit{body}_{ob2}\rangle\}), \texttt{fut}(\texttt{fn}_2, \bot) \quad \}\end{array}$$

where body_{ob1} (respectively body_{ob2}) is the renamed body of method readOnce (respectively readBlock). Furthermore, tv_{ob1} (respectively tv_{ob2}) stores the assignment $tv_{ob1}(\texttt{ret}) = (f1, \texttt{fn}_1)$ (respectively $tv_{ob2}(\texttt{ret}) = (f2, \texttt{fn}_2)$). When the event $\langle tv_{ob1}, \mathit{body}_{ob1}\rangle$ is extracted from the queue of $ob1$ (rule 11), its complete processing will replace $\texttt{fut}(\texttt{fn}_1, \bot)$ by $\texttt{fut}(\texttt{fn}_1, v)$ (rule 7), where v is the value returned by the method readOnce. Then, rule 9 can be used to process the instruction *await* $f1?$ of the object main. At this point the new state will take this form:

$$\begin{array}{ll}\{ & \texttt{ob}(\mathsf{main}, \bot, \bot, \langle tv_{\mathsf{main}}[f1 \mapsto \texttt{fn}_1, f2 \mapsto \texttt{fn}_2], \bar{b}''\rangle, \emptyset),\\ & \texttt{ob}(ob1, \mathsf{FileIS}, h_{ob1}, \epsilon, \emptyset), \texttt{fut}(\texttt{fn}_1, v),\\ & \texttt{ob}(ob2, \mathsf{FileIS}, h_{ob2}, \epsilon, \{\langle tv_{ob2}, \mathit{body}_{ob2}\rangle\}), \texttt{fut}(\texttt{fn}_2, \bot) \quad \}\end{array}$$

■

3. COST AND COST MODELS FOR CONCURRENT PROGRAMS

In this section we define the notion of cost we want to approximate by using static analysis. An execution step is annotated as $\mathcal{S} \rightsquigarrow_o^b \mathcal{S}'$, which denotes that we move from a state $\mathcal{S}$ to a state $\mathcal{S}'$ by executing instruction b in object o. Note that from a given state there may be several possible execution steps that can be taken since we make no assumptions on task scheduling and object selection. In order to quantify the cost of an execution step, we use a generic cost model. The following definition formalizes the notion of cost model, where Ins stands for the set of instructions b built using the grammar in Definition 1.

Definition 2 (cost model and cost of execution steps)**.** *A* cost model $\mathcal{M}$ *is a function defined as* $\mathcal{M} : \mathit{Ins} \mapsto \mathbb{R}$*. The cost of an execution step is defined as* $\mathcal{M}(\mathcal{S} \rightsquigarrow_o^b \mathcal{S}') = \mathcal{M}(b)$.

In the execution of sequential programs, the *cumulative cost of a trace* $\mathcal{T}$ is obtained by applying the cost model to each step of the trace. In our setting, this has to be extended because, rather than considering a single machine in which all steps are performed, we have a potentially distributed setting, with multiple objects possibly running concurrently on different CPUs. Thus, rather than

DOI: 10.1002/stvr

aggregating the cost of all executing steps, it is more useful to treat execution steps which occur on different computing infrastructures separately. With this aim, we adopt the notion of *cost centers* [36], proposed for profiling functional programs. Since the concurrency unit of our language is the object, cost centers are used to charge the cost of each step to the cost center associated to the object where the step is performed. For a given set of object identifiers O and a trace $\mathcal{T}$, we define a restricted trace, $\mathcal{T}|_O = \{\mathcal{S}_i \rightsquigarrow_o^b \mathcal{S}_{i+1} \mid \mathcal{S}_i \rightsquigarrow_o^b \mathcal{S}_{i+1} \in \mathcal{T},\ o \in O\}$ to denote the set of execution steps that are attributed to the objects in O. Then, we can define the cost executed by a particular object for a given trace:

Definition 3 (cost attributed to an object)**.** *Given a trace $\mathcal{T}$, a cost model $\mathcal{M}$ and an object identifier o, we define the* cost of $\mathcal{T}$ w.r.t. $\mathcal{M}$ attributed to o *as:*

$$\mathcal{C}(\mathcal{T}, o, \mathcal{M}) = \sum_{t \in \mathcal{T}|_{\{o\}}} \mathcal{M}(t)$$

3.1. Examples of Cost Models

We consider platform independent cost models (e.g., worst-case execution time or energy consumption are excluded). A cost model for approximating the *number of executed instructions* can be defined as $\mathcal{M}_i(b) = 1$, for all $b \in Ins$. Note that also calls of the form $\mathsf{call}(b, \_)$ count 1. This is because the call either corresponds to a synchronous call or to a call block requiring the execution of a guard. By '_', we mean any (valid) expression.

Other interesting cost models can be defined. For instance, a cost model that counts the *total number of objects created* along the execution can be defined as $\mathcal{M}_o(b) = 1$ if $b \equiv x := \mathsf{new}\ C$ and $\mathcal{M}_o(b) = 0$ otherwise. Since objects are the concurrency units, this cost model provides an indication on the amount of parallelism that might be achieved. A cost model that counts $\mathsf{call}(m, \_)$, can be used to infer the number of tasks that are spawned along an execution. We can also count the number of calls to specific methods or objects, e.g., by counting $\mathsf{call}(m, \_(o, \_, \_))$ we obtain bounds on the number of requests to a component o. This is useful for approximating the components' load and finding optimal deployment configurations (e.g., group objects according to the amount of tasks they receive to execute, by also taking into account the infrastructure on which they are deployed). The above cost models can also be used to prove termination of the program by setting the underlying solver [4] to only bound the number of iterations in loops. It is customary to have a cost model for memory consumption. The ABS language has a functional sub-language used to create data types. Hence in our actual implementation, the memory consumption cost model counts the sizes of constructed terms, and the sizes of the objects which are not intended to be concurrency units.

4. POINTS-TO ANALYSIS FOR CONCURRENT PROGRAMS

The aim of points-to analysis is to approximate the set of objects which each reference variable may point to during program execution. An analysis is *object-sensitive* [35, 40] if methods may be analyzed separately for different (sets of) objects on which they are invoked. More precisely, the analysis uses a finite set of *object names* to partition the (possibly infinite) set of objects allocated at runtime into *contexts* which are analyzed separately.

This section presents a flow-sensitive object-sensitive points-to analysis for concurrent programs. It is based on Milanova's analysis framework [35] for Java. As Milanova's analysis is flow-insensitive, it is sound for concurrent programs because it implicitly considers all possible interactions and interleavings between tasks that may happen in a concurrent program. However, our proposed analysis is flow-sensitive since for the inference of the object-sensitive recurrence relations in Section 6, it is fundamental to track flow-sensitive relations among objects.

It is known that flow-sensitive analysis of concurrent programs is challenging due to the complexity of their flow. All possible task interleavings must be considered in order to develop

DOI: 10.1002/stvr

a sound analysis. As our contribution in this regard, we extend the analysis of [35] to make it flow-sensitive in the presence of concurrent behaviours. The concurrency model guarantees that both fields and local variables can only be modified locally in the active task (i.e., in a flow-sensitive way) until the processor is released. At such release points, the values of local variables cannot be changed, whereas the state of the fields might be modified by other tasks. The main idea of our approach is to keep the flow-sensitive and the flow-insensitive information separate. We first describe a transfer function which handles those instructions that are guaranteed to be executed sequentially (flow-sensitive). Then, as in [13], we define our points-to analysis as the solution of a constraint equation system. Such equation system not only handles the flow-sensitive information by means of the transfer function, but it also treats in a flow-insensitive way method invocations and those instructions where the processor can be released (e.g., *await*). Finally, the solution of the equation system obtains the desired points-to information. As we will see in the analysis, not all information about fields has to be lost at the release points. By keeping track of the values of the reference *this*, we can notably reduce information loss.

Section 4.1 presents the basic abstract domain that we use to represent the points-to information; Section 4.2 describes how the different (sequential) instructions change from one abstract state to another; and Section 4.3 introduces the analysis as a solution of a constraint equation system that is built from the program.

4.1. The Abstract Domain

The term *allocation site* refers to program points in which objects are created by executing a *new* instruction. Let $S{=}\{1,\ldots,n\}$ be the set of all allocation sites in a program, we use i (possibly subscripted) to refer to the elements in S. Given $S{=}\{1,\ldots,n\}$ and $q \in \mathbb{N}$, where $1 \leq q \leq n$, we define the set S^q as:

$$S^q = \{s \equiv \langle i_1,\ldots,i_q\rangle \mid \forall j \in \{1,\ldots,q\} \cdot i_j \in \{1,...,n\}\}$$

Finally, for a set $S{=}\{1,\ldots,n\}$ of allocation sites, and a constant $k \in \mathbb{N}$, the analysis considers a finite set of *object names*, denoted $\mathcal{N}^k$, which is defined as:

$$\mathcal{N}^k = \{o_s \mid (s \equiv \epsilon) \vee (\exists q \in \{1,\ldots,k\} \cdot s \in S^q)\}$$

For example, if $S{=}\{1,2\}$, then $S^1{=}\{\langle 1\rangle,\langle 2\rangle\}$, $S^2{=}\{\langle 1,1\rangle,\langle 2,2\rangle,\langle 1,2\rangle,\langle 2,1\rangle\}$ and $\mathcal{N}^2 = \{o_\epsilon, o_{\langle 1\rangle}, o_{\langle 2\rangle}, o_{\langle 1,1\rangle}, o_{\langle 2,2\rangle}, o_{\langle 1,2\rangle}, o_{\langle 2,1\rangle}\}$. Note that k defines the maximum size of sequences of allocation sites, and it allows controlling the precision of the analysis. Allocation sequences have in principle unbounded length and, thus, to ensure termination it is sometimes necessary to lose precision during analysis. This is done by just keeping the k rightmost positions in sequences whose length is greater than k. We define the operation $\oplus_k$ as:

$$\langle i_1,\ldots,i_p\rangle \oplus_k i \quad = \begin{cases} \langle i\rangle & k = 1 \\ \langle i_1,\ldots,i_p,i\rangle & k > 1 \wedge p < k \\ \langle i_{(p+2)-k},\ldots,i_p,i\rangle & k > 1 \wedge p \geq k \end{cases}$$

Note that a variable in a program can be assigned objects with different object names. In order to represent all possible objects pointed to by a variable, sets of object names are used. Given a program, the set S of all allocation sites for it and $k \geq 0$, the *abstraction* of an object created in the program is an element of $\mathcal{N}^k$. Furthermore, $o_{\langle i_1,\ldots,i_p\rangle}$ represents all run-time objects that were created at allocation site i_p when the enclosing instance method was invoked on an object which was in turn created at allocation site i_{p-1}, and the same applies until i_1.

As notation, we use $\mathcal{V}$ to represent the set of all possible *reference* local variables that may occur in a program and $\mathcal{F}^*$ to represent all possible pairs (o_s, f) which denote all possible accesses to the *reference* field f through the objects $o_s \in \mathcal{N}^k$. In what follows, such pairs are represented as $o_s.f$.

Following Milanova's approach, context sensitivity is achieved by maintaining multiple replicas of each reference variable x for each possible context in which x may be used for calling a method. Let x be a local variable and o an object name to which *this* may point to, we use the fresh variable name x^o to store the analysis information for x and context o. We drop the superscript o when it is not relevant. The set of *replicas* is defined by $\mathcal{R} : \mathcal{V} \times \mathcal{N}^k \mapsto \mathcal{V}^*$, where $\mathcal{V}^* = \{x^o \mid o \in \mathcal{N}^k, x \in \mathcal{V}\}$.

DOI: 10.1002/stvr

	b	$\tau(b,\pi)$	
(1)	(i) $x := \textit{new } C$	$\pi[x^l \mapsto \{l \oplus_k i\}]$	$\forall l \in \pi(this)$
(2)	$x := y$	$\pi[x^l \mapsto \pi(y^l)]$	$\forall l \in \pi(this)$
(3)	$x := \textit{null}$	$\pi[x^l \mapsto \emptyset]$	$\forall l \in \pi(this)$
(4)	$x := y.\textit{get}$	$\pi[x^l \mapsto \pi(y^l)]$	$\forall l \in \pi(this)$
(5)	$x := this.f$	$\pi[x^l \mapsto \pi(l.f)]$	$\forall l \in \pi(this)$
(6)	$this.f := y$	$\pi[l.f \mapsto \pi(y^l)]$	$\forall l \in \pi(this)$
(7)	$otherwise$	π	

Figure 4. Transfer function (where $l \equiv o_{\langle i_1,\ldots,i_p\rangle}$, $l \oplus_k i \equiv o_{\langle i_1,\ldots,i_p\rangle \oplus_k i}$)

Definition 4 (points-to abstract state)**.** *A points-to abstract state is a mapping π from $\mathcal{V}^* \cup \mathcal{F}^*$ to $\wp(\mathcal{N}^k)$, i.e., $\pi : \mathcal{V}^* \cup \mathcal{F}^* \mapsto \wp(\mathcal{N}^k)$.*

The application of π to an element of the form x^o, that is, $\pi(x^o)$, corresponds to the the flow-sensitive information regarding all possible objects that may be assigned to the local variable x when $this$ points to the object name o, and similarly for $\pi(o.f)$. The abstract domain is the lattice $\langle AS, \pi_\top, \pi_\bot, \sqcup, \sqsubseteq\rangle$, where AS is the set of abstract states and $\pi_\top$ is the top of the lattice defined as $\forall x^o \in \mathcal{V}^*, \pi_\top(x^o) = \mathcal{N}^k$, and $\forall o.f \in \mathcal{F}^*, \pi_\top(o.f) = \mathcal{N}^k$. The bottom element of the lattice is $\pi_\bot$, i.e., $\forall x^o \in \mathcal{V}^*, \pi_\bot(x^o) = \emptyset$, $\forall o.f \in \mathcal{F}^*, \pi_\bot(o.f) = \emptyset$. Given two abstract states π_1 and π_2, we use $\pi = \pi_1 \sqcup \pi_2$ to denote that π is the least upper bound, defined as $\forall x^o \in \mathcal{V}^*, \pi(x^o) = \pi_1(x^o) \cup \pi_2(x^o)$ and $\forall o.f \in \mathcal{F}^*, \pi(o.f) = \pi_1(o.f) \cup \pi_2(o.f)$. Similarly, $\pi_1 \sqsubseteq \pi_2$ holds iff $\forall x^o \in \mathcal{V}^*, \pi_1(x^o) \subseteq \pi_2(x^o)$ and $\forall o.f \in \mathcal{F}^*, \pi_1(o.f) \subseteq \pi_2(o.f)$.

4.2. The Transfer Function

One of the ingredients of our analysis is a transfer function that describes the effect of each sequential instruction on the local state of the corresponding program point. The treatment of those instructions that can modify the flow of the program, that is, synchronization instructions and method calls, will be detailed in the next section.

Definition 5 (points-to transfer function)**.** *Given the set of abstract states AS and the set of instructions in the program Ins, the points-to transfer function τ is defined as a mapping $\tau : Ins \times AS \mapsto AS$ computed according to the table in Figure 4.*

In Figure 4, we use (i) in row 1 to denote that the considered allocation site is identified with i, we use x, y to denote reference variables and f to denote a reference field. It is important to note that the operations performed by τ affect π in a flow-sensitive way, i.e., the updates performed on variables or fields overwrite their previous abstract values. As the analysis is object-sensitive, all rows of the transfer function modify the state of local variables (or fields) for all objects in $\pi(this)$. Rows 1-3 correspond to local reference variables. Row 4, which corresponds to *get* instructions, is also treated as an assignment between local variables. According to the concurrency model, the execution of a sequence of instructions is guaranteed to be sequential until release points, i.e, *await* instructions and method entries. For this reason, the transfer function treats field accesses (rows 5 and 6 of Figure 4) in a local (flow-sensitive) way, i.e., as local variables. However, the points-to information for reference fields at release points must consider all possible task interleavings (flow-insensitive). This is the reason why the release points are not treated by the transfer function. Therefore, in release points, propagation of field updates to the global state and method calls, are handled by the constraint equation system defined in next section. Other instructions (row 7) do not modify the abstract state.

4.3. The Constraint Equation System

This section presents a method to build a system of constraints whose solution describes the desired points-to information (similar to the one described at [13]). Such equation system is generated

DOI: 10.1002/stvr

from the IR of the program. From now on, given a rule $r \equiv p(this, \bar{x}, \bar{y}) \leftarrow g, b_1, \ldots, b_n$ in the IR, we assume that the rule r is annotated by a unique identifier $t \in \mathbb{N}$, and we write r_t instead of r. Now given a rule $r_t \equiv p(this, \bar{x}, \bar{y}) \leftarrow g, b_1, \ldots, b_n$, we use $\langle t, j\rangle$ to denote the *program point* of the instruction b_j, $1 \leq i \leq n$. The tuple $\langle t, 0\rangle$ (resp. $\langle t, n\rangle$) corresponds to the entry (resp. exit) of the rule. Given a program P, $\mathsf{pp}(P) = \{\langle t, j\rangle \mid r_t \equiv p(this, \bar{x}, \bar{y}) \leftarrow g, b_1, \ldots, b_n \in P, 0 \leq j \leq n\}$. We write $b_{\langle t,j\rangle}$ to refer to the instruction b at program point $\langle t, j\rangle$ and $b_{\langle t,0\rangle}$ to refer to the head of the rule r_t.

Definition 6 (points-to constraint equation system)**.** *The* points-to constraint equation system *is a set of the form* $\mathcal{L}_P = \bigcup\{C_{\langle t,j\rangle} \mid \langle t, j\rangle \in \mathsf{pp}(P)\}$*, where* $C_{\langle t,j\rangle}$ *is defined as:*

	$b_{\langle t,j\rangle}$	$C_{\langle t,j\rangle}$
(1)	$\overbrace{p(this, \bar{x}, \bar{y})}^{b_{\langle t,0\rangle}} \leftarrow g, b_1, \ldots, b_n \in P$	$\mathcal{X}_{\langle t,0\rangle} \sqsupseteq \mathcal{X}_p^{call}$ *if p is a block* $\mathcal{X}_{\langle t,0\rangle} \sqsupseteq \mathsf{include\_global}(\mathcal{X}_\mathcal{G}, \mathcal{X}_p^{call})$ *if p is a method* $\mathcal{X}_p^{exit} \sqsupseteq \mathcal{X}_{\langle t,n\rangle}$
(2)	$this.f := y$	$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \tau(b_{\langle t,j\rangle}, \mathcal{X}_{\langle t,j-1\rangle})$ $\mathcal{X}_\mathcal{G} \sqsupseteq \mathsf{update\_global}(\mathcal{X}_\mathcal{G}, f, \mathcal{X}_{\langle t,j-1\rangle})$
(3)	*await* x*?*	$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathsf{include\_global}(\mathcal{X}_\mathcal{G}, \mathcal{X}_{\langle t,j-1\rangle})$
(4)	*call*$(b, q(this, \bar{z}, \bar{w}))$ *where* $r' \equiv q(this, \bar{u}, \bar{v}) \leftarrow \ldots \in P$	$\mathcal{X}_q^{call} \sqsupseteq \mathsf{restrict\_block}(\mathcal{X}_{\langle t,j-1\rangle}, \bar{z}, \bar{u})$ $\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathsf{extend\_block}(\mathcal{X}_{\langle t,j-1\rangle}, \mathcal{X}_q^{exit}, \bar{v}, \bar{w})$
(5)	*call*$(m, m(rec, \bar{z}, w))$ *where* $r' \equiv m(rec, \bar{u}, v) \leftarrow \ldots \in P$	$\mathcal{X}_m^{call} \sqsupseteq \mathsf{restrict\_method}(\mathcal{X}_{\langle t,j-1\rangle}, rec, \bar{z}, \bar{u})$ $\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathsf{extend\_method}(\mathcal{X}_{\langle t,j-1\rangle}, \mathcal{X}_m^{exit}, rec, v, w)$
(6)	*otherwise*	$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \tau(b_{\langle t,j\rangle}, \mathcal{X}_{\langle t,j-1\rangle})$

where:

$\mathsf{include\_global}(\mathcal{X}_\mathcal{G}, \mathcal{X}) \equiv \mathcal{X}[o.f \mapsto \mathcal{X}_\mathcal{G}(o.f)], \forall o.f \in dom(\mathcal{X}) \mid o \in \mathcal{X}(this)$
$\mathsf{update\_global}(\mathcal{X}_\mathcal{G}, f, \mathcal{X}) \equiv \mathcal{X}_\mathcal{G}[o.f \mapsto \mathcal{X}_\mathcal{G}(o.f) \cup \mathcal{X}(o.f)], \forall o.f \in dom(\mathcal{X}) \mid o \in \mathcal{X}(this)$
$\mathsf{restrict\_block}(\mathcal{X}, \bar{z}, \bar{u}) \equiv \pi_\perp[this \mapsto \mathcal{X}(this), \bar{u}^l \mapsto \mathcal{X}(\bar{z}^l), l.f \mapsto \mathcal{X}(l.f)], \forall l \in \mathcal{X}(this)$
$\mathsf{extend\_block}(\mathcal{X}_1, \mathcal{X}_2, \bar{v}, \bar{w}) \equiv \mathcal{X}_1[\bar{w}^l \mapsto \mathcal{X}_2(\bar{v}^l), l.f \mapsto \mathcal{X}_2(l.f)], \forall l \in \mathcal{X}(this)$
$\mathsf{restrict\_method}(\mathcal{X}, rec, \bar{z}, \bar{u}) \equiv \pi_\perp[this \mapsto This', \bar{u}^{l'} \mapsto \mathcal{X}(\bar{z}^l)], \forall l \in \mathcal{X}(this), l' \in \mathcal{X}(rec^l)$
$\mathsf{extend\_method}(\mathcal{X}_1, \mathcal{X}_2, rec, v, w) \equiv \mathcal{X}_1[w^l \mapsto \bigsqcup_{l' \in \mathcal{X}_1(rec^l)} \mathcal{X}_2(v^{l'})], \forall l \in \mathcal{X}_1(this)$
$This' \equiv \bigsqcup_{l \in \mathcal{X}(this)} \mathcal{X}(x^l)$

Intuitively, each program point in the IR contributes some constraints to $\mathcal{L}_P$. The elements in $\mathcal{L}_P$ contain variables named $\mathcal{X}$, possibly sub- or super-scripted, which are points-to abstract states (see Definition 4). In $\mathcal{L}_P$ we have variables of the form $\mathcal{X}_{\langle t,j\rangle}$, which stand for the points-to information after program point $\langle t, j\rangle$. In addition, $\mathcal{L}_P$ includes a global variable $\mathcal{X}_\mathcal{G}$, which gathers the global information related to all reference fields in the program. Moreover, $\mathcal{L}_P$ contains constraint variables of the form $\mathcal{X}_p^{call}$ (resp. $\mathcal{X}_p^{exit}$), which store the points-to information used for calling (resp. returning from) rule p. As notation, $\mathcal{X}_{\langle t,j\rangle}(x)$ refers to the mapping of the reference variable x in the abstract state stored in the constraint variable $\mathcal{X}_{\langle t,j\rangle}$. The analysis includes a constraint $\mathcal{X}_{\mathsf{main}}^{call} \sqsupseteq \pi_\perp[this \mapsto o_\epsilon]$ which corresponds to the initial execution of main.

Let us explain in detail the equations generated by Definition 6. For the head of a rule p (row 1), we have one constraint variable $\mathcal{X}_p^{call}$ that represents the points-to information on how p can be called. As blocks start their execution immediately and methods might be postponed, they are treated separately. The equation corresponding to blocks (row 1) directly includes $\mathcal{X}_p^{call}$, which stands for the flow-sensitive information from the calling program point. However, as methods start the execution asynchronously and other tasks might interleave, the equation for methods uses include_global to incorporate the global information for fields. This global information is stored at $\mathcal{X}_\mathcal{G}$, as the analysis cannot rely on the flow-sensitive information. In addition, row 1 includes the constraint $\mathcal{X}_p^{exit} \sqsupseteq \mathcal{X}_{\langle t,n\rangle}$ to capture the points-to information in the last instruction of the rule.

Assignments to fields in row 2 apply the transfer function τ to the points-to information of the previous instruction in order to update the local state. In addition, the global state $\mathcal{X}_\mathcal{G}$ is updated by means of update_global, which includes the new information for the field f for all abstract objects

DOI: 10.1002/stvr

that are pointed by $this$. Note that update_global updates $\mathcal{X}_\mathcal{G}$ with $\mathcal{X}_\mathcal{G}(o.f) \cup \mathcal{X}_{\langle t,j\rangle}(o.f)$, joining the local information regarding fields into $\mathcal{X}_\mathcal{G}$, so that $\mathcal{X}_\mathcal{G}$ gathers the points-to information for fields in a flow-insensitive way. The equation in (3) corresponds to *await* instructions, and it considers that at this program point the processor can be released so that other tasks might modify the contents of the fields. We use include_global to propagate the global information about fields to the flow-sensitive state.

A call to a block, row 4, adds two equations to the system. In the first one, restrict_block projects the points-to information of fields $l.f$ and actual parameters $\bar{z}$ from $\mathcal{X}^p_{\langle t,j-1\rangle}$ into the fields and formal parameters $\bar{u}$ of the called rule $\mathcal{X}^{call}_q$. The second equation for calling a block uses extend_block to project the returned values $\bar{v}$ and the fields $l.f$, taken from $\mathcal{X}^{exit}_q$, into the corresponding local variables $\bar{w}$ and fields of the current state $\mathcal{X}_{\langle t,j\rangle}$. Calls to methods in row 5 are treated similarly to block calls, but, since the analysis is object-sensitive, we also handle the calling context creation. When a method is called we first use restrict_method to project the arguments of a method call to the parameters of the called method. Observe that, whereas in called blocks there are copies of the variables for each element l in $\mathcal{X}(this)$, called methods require copies of the variables for each element l' in $This'$. Moreover, we use the function extend_method to update the abstract state of the calling program point with the results of analyzing the called method. Finally, in contrast to extend_block, as methods are called asynchronously, fields are not updated upon method exit. Other instructions directly apply the transfer function of Definition 5 to the previous state to get the flow-sensitive points-to information.

Example 5 (points-to constraint equation system)**.** Figure 5 illustrates, in the third column, the constraint equation system which encapsulates the desired points-to information for the IR program in the second column. We mark allocation sites with ⓐ, ⓑ and ⓒ instead of using numbers. This allow to distinguish clearly program points in the first column from allocation sites. Column $\mathsf{pp}(P)$ shows the program point of each instruction. Note that the identifiers for rules $main$, $m1$, $m2$ and $m3$ are 0, 1, 2 and 3, respectively. The rule for $main$ first creates an object of class A (program point $\langle 0,1\rangle$) and afterwards, methods $m1$ and $m2$ are called on the newly created object. The field f of class A is modified at program points $\langle 1,2\rangle$ and $\langle 2,2\rangle$.
The initial constraint of the analysis is shown at row 1, $\mathcal{X}^{call}_{main} \sqsupseteq \pi_\perp[this \mapsto o_\epsilon]$. Flow-insensitive information is constrained at program points where f is written, i.e., the new information about f is added to $\mathcal{X}_\mathcal{G}$ at rows 13 (program point $\langle 1,2\rangle$) and 22 (program point $\langle 2,2\rangle$). The global information stored in $\mathcal{X}_\mathcal{G}$ is used in row 17, where there is an *await* instruction, to include in $\mathcal{X}_{\langle 1,5\rangle}$ the global information about fields. As methods are invoked asynchronously, global information in $\mathcal{X}_\mathcal{G}$ is also used at the beginning of method rules at rows 2, 9, 18 and 23. Besides, rows 3, 10, 19, 24 contain the equations that capture the method exit information, gathering the points-to information in the last program-point of the corresponding rule. Additionally, for each method call there are two constraints, one which restricts the current information to the calling context of the corresponding method (rows 5, 7 and 14), and another constraint to project the information from the exit of the methods to the state of the calling instruction (rows 6, 8, 15). The remaining rows directly apply τ with the instruction and the state of the previous program point. ■

When the constraint equation system is solved, constraint variables over-approximate the points-to information for the program. In particular, variables of the form $\mathcal{X}_{\langle t,j\rangle}$ store the points-to information after the execution of $b_{\langle t,j\rangle}$, $\mathcal{X}_\mathcal{G}$ contains the points-to information that might have been stored in fields along the program execution, and, $\mathcal{X}^{call}_m$ ($\mathcal{X}^{exit}_m$) gathers the points-to information used to call (returning from) m. Solving such system can be done iteratively. A naïve algorithm consists in first initializing each constraint variable to $\pi_\perp$, and then iteratively refining the values of these variables as follows:

1. substitute the current values of the constraint variables in the right hand side of each constraint, and then evaluate the right-hand side if needed;
2. if each constraint $\mathcal{X} \sqsupseteq E$ holds, where E is the value of the evaluation of the right-hand side of the previous step, then the process finishes; otherwise

DOI: 10.1002/stvr

	pp(P)	IR	Constraint Equation System
(1)			$\mathcal{X}^{call}_{main} \sqsupseteq \pi_\perp[this \mapsto o_\epsilon]$
(2)	$\langle 0,0\rangle$	$main(this, \langle\rangle, \langle\rangle) \leftarrow$	$\mathcal{X}_{\langle 0,0\rangle} \sqsupseteq$ include_global$(\mathcal{X}_\mathcal{G}, \mathcal{X}^{call}_{main})$
(3)			$\mathcal{X}^{exit}_{main} \sqsupseteq \mathcal{X}_{\langle 0,3\rangle}$
(4)	$\langle 0,1\rangle$	ⓐ$w :=$ *new* A,	$\mathcal{X}_{\langle 0,1\rangle} \sqsupseteq \tau(w :=$ *new* $A, \mathcal{X}_{\langle 0,0\rangle})$
(5)	$\langle 0,2\rangle$	$call(m, m1(w,\langle\rangle,\langle\rangle))$,	$\mathcal{X}^{call}_{m1} \sqsupseteq$ restrict_method$(\mathcal{X}_{\langle 0,1\rangle}, w, \langle\rangle, \langle\rangle)$
(6)			$\mathcal{X}_{\langle 0,2\rangle} \sqsupseteq$ extend_method$(\mathcal{X}_{\langle 0,1\rangle}, \mathcal{X}^{exit}_{m1}, w, \langle\rangle, \langle\rangle)$
(7)	$\langle 0,3\rangle$	$call(m, m2(w,\langle\rangle,\langle y\rangle))$.	$\mathcal{X}^{call}_{m2} \sqsupseteq$ restrict_method$(\mathcal{X}_{\langle 0,2\rangle}, w, \langle\rangle, \langle\rangle)$
(8)			$\mathcal{X}_{\langle 0,3\rangle} \sqsupseteq$ extend_method$(\mathcal{X}_{\langle 0,2\rangle}, \mathcal{X}^{exit}_{m2}, w, \langle u\rangle, \langle y\rangle)$
(9)	$\langle 1,0\rangle$	$m1(this, \langle\rangle, \langle\rangle) \leftarrow$	$\mathcal{X}_{\langle 1,0\rangle} \sqsupseteq$ include_global$(\mathcal{X}_\mathcal{G}, \mathcal{X}^{call}_{m1})$
(10)			$\mathcal{X}^{exit}_{m1} \sqsupseteq \mathcal{X}_{\langle 1,5\rangle}$
(11)	$\langle 1,1\rangle$	ⓑ$y :=$ *new* B,	$\mathcal{X}_{\langle 1,1\rangle} \sqsupseteq \tau(w :=$ *new* $B, \mathcal{X}_{\langle 1,0\rangle})$
(12)	$\langle 1,2\rangle$	$this.f := y$,	$\mathcal{X}_{\langle 1,2\rangle} \sqsupseteq \tau(this.f := y, \mathcal{X}_{\langle 1,1\rangle})$
(13)			$\mathcal{X}_\mathcal{G} \sqsupseteq$ update_global$(\mathcal{X}_\mathcal{G}, f, \mathcal{X}_{\langle 1,2\rangle})$
(14)	$\langle 1,3\rangle$	$call(m, m3(this, \langle\rangle,\langle r\rangle))$,	$\mathcal{X}^{call}_{m3} \sqsupseteq$ restrict_method$(\mathcal{X}_{\langle 1,2\rangle}, t, \langle\rangle, \langle\rangle)$
(15)			$\mathcal{X}_{\langle 1,3\rangle} \sqsupseteq$ extend_method$(\mathcal{X}_{\langle 1,2\rangle}, \mathcal{X}^{exit}_{m3}, t, \langle\rangle, \langle\rangle)$
(16)	$\langle 1,4\rangle$	$x := this.f$,	$\mathcal{X}_{\langle 1,4\rangle} \sqsupseteq \tau(x := this.f, \mathcal{X}_{\langle 1,3\rangle})$
(17)	$\langle 1,5\rangle$	*await* $r?$.	$\mathcal{X}_{\langle 1,5\rangle} \sqsupseteq$ include_global$(\mathcal{X}_\mathcal{G}, \mathcal{X}_{\langle 1,4\rangle})$
(18)	$\langle 2,0\rangle$	$m2(this, \langle\rangle, \langle u\rangle) \leftarrow$	$\mathcal{X}_{\langle 2,0\rangle} \sqsupseteq$ include_global$(\mathcal{X}_\mathcal{G}, \mathcal{X}^{call}_{m2})$
(19)			$\mathcal{X}^{exit}_{m2} \sqsupseteq \mathcal{X}_{\langle 2,2\rangle}$
(20)	$\langle 2,1\rangle$	ⓒ$u :=$ *new* B,	$\mathcal{X}_{\langle 2,1\rangle} \sqsupseteq \tau(w :=$ *new* $B, \mathcal{X}_{\langle 2,0\rangle})$
(21)	$\langle 2,2\rangle$	$this.f := u$.	$\mathcal{X}_{\langle 2,2\rangle} \sqsupseteq \tau(this.f := u, \mathcal{X}_{\langle 2,1\rangle})$
(22)			$\mathcal{X}_\mathcal{G} \sqsupseteq$ update_global$(\mathcal{X}_\mathcal{G}, f, \mathcal{X}_{\langle 2,2\rangle})$
(23)	$\langle 3,0\rangle$	$m3(this, \langle\rangle, \langle v\rangle) \leftarrow$	$\mathcal{X}_{\langle 3,0\rangle} \sqsupseteq$ include_global$(\mathcal{X}_\mathcal{G}, \mathcal{X}_{call})$
(24)			$\mathcal{X}^{exit}_{m3} \sqsupseteq \mathcal{X}_{\langle 3,1\rangle}$
(25)	$\langle 3,1\rangle$	$z := this.f$.	$\mathcal{X}_{\langle 3,1\rangle} \sqsupseteq \tau(z := this.f, \mathcal{X}_{\langle 3,0\rangle})$

Figure 5. Points-to system of constraints for a program with fields

3. for each $\mathcal{X} \sqsupseteq E$ which does not hold, let E' be the current value of $\mathcal{X}$. Then update the current value of $\mathcal{X}$ to $E \sqcup E'$. Once all these updates are (iteratively) applied we repeat the process at step 1.

Termination is guaranteed since the abstract domain does not have infinitely ascending chains. In the implementation, we apply several optimizations to improve the performance of the above process. We omit the details as they are not important for explaining the analysis.

Example 6 (points-to results for Example 5)**.** Figure 6 presents the solutions computed for the constraint system in Figure 5. We have used $k = 2$, but, for this particular example, any value of k achieves the maximum precision. For simplicity, in the examples we omit tuples when it is clear from the context, i.e., for example we use o_{ab} to denote $o_{\langle a,b\rangle}$. Also we keep the calling context as a superscript to the variable such that x^{o_a} denotes the abstract value for x when $this$ points to o_a. When solving the constraint system, we assume that there exists a function allocSite which, given a program P and a program point in pp(P) corresponding to a *new* instruction, it returns the corresponding allocation site. Thus allocSite$(P, \langle 0,1\rangle) =$ ⓐ, allocSite$(P, \langle 1,1\rangle) =$ ⓑ, and allocSite$(P, \langle 2,1\rangle) =$ ⓒ. When we write o_a we refer to an object created at the allocation site

DOI: 10.1002/stvr

	pp(P)	IR	Points-to Results
(0)			$\mathcal{X}_{\mathcal{G}} = \{o_a.f \mapsto \{o_{ab}, o_{ac}\}\}$
(1)			$\mathcal{X}^{call}_{main} = \{this \mapsto \{o_\epsilon\}\}$
(2)	$\langle 0,0\rangle$	$main(this, \langle\rangle, \langle\rangle) \leftarrow$	$\mathcal{X}_{\langle 0,0\rangle} = \{this \mapsto \{o_\epsilon\}\}$
(3)			$\mathcal{X}^{exit}_{main} = \{this \mapsto \{o_\epsilon\}, w^{o_\epsilon} \mapsto \{o_a\}, y^{o_\epsilon} \mapsto \{o_{ac}\}\}$
(4)	$\langle 0,1\rangle$	ⓐ$w :=$ *new* A,	$\mathcal{X}_{\langle 0,1\rangle} = \{this \mapsto \{o_\epsilon\}, w^{o_\epsilon} \mapsto \{o_a\}\}$
(5)	$\langle 0,2\rangle$	$call(m, m1(w, \langle\rangle, \langle\rangle))$,	$\mathcal{X}^{call}_{m1} = \{this \mapsto \{o_a\}\}$
(6)			$\mathcal{X}_{\langle 0,2\rangle} = \{this \mapsto \{o_\epsilon\}, w^{o_\epsilon} \mapsto \{o_a\}\}$
(7)	$\langle 0,3\rangle$	$call(m, m2(w, \langle\rangle, \langle y\rangle))$.	$\mathcal{X}^{call}_{m2} = \{this \mapsto \{o_a\}\}$
(8)			$\mathcal{X}_{\langle 0,3\rangle} = \{this \mapsto \{o_\epsilon\}, w^{o_\epsilon} \mapsto \{o_a\}, y^{o_\epsilon} \mapsto \{o_{ac}\}\}$
(9)	$\langle 1,0\rangle$	$m1(this, \langle\rangle, \langle\rangle) \leftarrow$	$\mathcal{X}_{\langle 1,0\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}, o_{ac}\}\}$
(10)			$\mathcal{X}^{exit}_{m1} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}, o_{ac}\}, y^{o_a} \mapsto \{o_{ab}\}, x^{o_a} \mapsto \{o_{ab}\}\}$
(11)	$\langle 1,1\rangle$	ⓑ$y :=$ *new* B,	$\mathcal{X}_{\langle 1,1\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}, o_{ac}\}, y^{o_a} \mapsto \{o_{ab}\}\}$
(12)	$\langle 1,2\rangle$	$this.f := y$,	$\mathcal{X}_{\langle 1,2\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}\}, y^{o_a} \mapsto \{o_{ab}\}\}$
(13)	$\langle 1,3\rangle$	$call(m, m3(this, \langle\rangle, \langle r\rangle))$,	$\mathcal{X}^{call}_{m3} = \{this \mapsto \{o_a\}\}$
(14)			$\mathcal{X}_{\langle 1,3\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}\}, y^{o_a} \mapsto \{o_{ab}\}\}$
(15)	$\langle 1,4\rangle$	$x := this.f$,	$\mathcal{X}_{\langle 1,4\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}\}, y^{o_a} \mapsto \{o_{ab}\}, x^{o_a} \mapsto \{o_{ab}\}\}$
(16)	$\langle 1,5\rangle$	*await* r?.	$\mathcal{X}_{\langle 1,5\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}, o_{ac}\}, y^{o_a} \mapsto \{o_{ab}\}, x^{o_a} \mapsto \{o_{ab}\}\}$
(17)	$\langle 2,0\rangle$	$m2(this, \langle\rangle, \langle u\rangle) \leftarrow$	$\mathcal{X}_{\langle 2,0\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}, o_{ac}\}\}$
(18)			$\mathcal{X}^{exit}_{m2} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ac}\}, u^{o_a} \mapsto \{o_{ac}\}\}$
(19)	$\langle 2,1\rangle$	ⓒ$u :=$ *new* B,	$\mathcal{X}_{\langle 2,1\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}, o_{ac}\}, u^{o_a} \mapsto \{o_{ac}\}\}$
(20)	$\langle 2,2\rangle$	$this.f := u$.	$\mathcal{X}_{\langle 2,2\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ac}\}, u^{o_a} \mapsto \{o_{ac}\}\}$
(21)	$\langle 3,0\rangle$	$m3(this, \langle\rangle, \langle v\rangle) \leftarrow$	$\mathcal{X}_{\langle 3,0\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}, o_{ac}\}\}$
(22)			$\mathcal{X}^{exit}_{m3} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}, o_{ac}\}, z^{o_a} \mapsto \{o_{ab}, o_{ac}\}\}$
(23)	$\langle 3,1\rangle$	$z := this.f$.	$\mathcal{X}_{\langle 3,1\rangle} = \{this \mapsto \{o_a\}, o_a.f \mapsto \{o_{ab}, o_{ac}\}, z^{o_a} \mapsto \{o_{ab}, o_{ac}\}\}$

Figure 6. Solution of the system of constraints for Example 5

ⓐ. Similarly o_{ab} stands for an object that was created at allocation site ⓑ, when the method was called from an object created at allocation site ⓐ.

When the constraint system is solved, $\mathcal{X}_{\mathcal{G}}$, shown at row 0, captures that $o_a.f$ might point to either o_{ab} or o_{ac}. Program points $\langle 1,1\rangle$ and $\langle 1,2\rangle$ (rows 11, 12) show how the analysis handles the local information for fields. At $\mathcal{X}_{\langle 1,1\rangle}$ we have that f might point to either o_{ab} or o_{ac}. As f is set to o_{ab} at $\langle 1,2\rangle$, and the execution is guaranteed to be sequential until $\langle 1,4\rangle$, field f only points to o_{ab} from $\mathcal{X}_{\langle 1,2\rangle}$ to $\mathcal{X}_{\langle 1,4\rangle}$. At row 16, as the processor might be released in the *await*, the information for f is updated with the information from $\mathcal{X}_{\mathcal{G}}$. Then, at $\mathcal{X}_{\langle 1,5\rangle}$, we have $o_a.f \mapsto \{o_{ab}, o_{ac}\}$. At the beginning of methods $m2$ and $m3$, the points-to information for f is initialized with the information from $\mathcal{X}_{\mathcal{G}}$. In $m2$ the information for f is constrained locally when f is assigned (row 20). As method calls do not have input arguments, rows 5, 7 and 13 show that only object $this$ is restricted, in all cases, to o_a. Outgoing states $\mathcal{X}^{exit}_{main}$, $\mathcal{X}^{exit}_{m1}$, $\mathcal{X}^{exit}_{m2}$, $\mathcal{X}^{exit}_{m3}$ (rows 3, 10, 18 and 22) are equal to the state of the last instruction of the corresponding rule. At row 8, as $m2$ is called with $this \mapsto \{o_a\}$, the variable y takes its value from $\mathcal{X}^{exit}_{m2}(u^{o_a}) = \{o_{ac}\}$, adding $y^{o_\epsilon} \mapsto \{o_{ac}\}$ to $\mathcal{X}_{\langle 0,3\rangle}$. ■

Once we have seen the details of the points-to analysis, let us see the results of the points-to analysis for our running example.

DOI: 10.1002/stvr

pp(P)	IR of the Running Example	Points-to Results
		$\mathcal{X}_{\mathcal{G}} = \{\}$
$\langle 0,0\rangle$	$main(this, \langle\rangle, \langle\rangle) \leftarrow$	$\mathcal{X}_{\langle 0,0\rangle} = \mathcal{X}^{call}_{main} = \{this \mapsto \{o_\epsilon\}\}$
$\langle 0,1\rangle$	①FileIS_$init(ob1, \langle$"A.txt", 20, 2$\rangle, \langle\rangle)$,	$\mathcal{X}_{\langle 0,1\rangle} = \{this \mapsto \{o_\epsilon\}, ob1 \mapsto \{o_1\}\}$
$\langle 0,2\rangle$	②FileIS_$init(ob2, \langle$"A.txt", 20, 3$\rangle, \langle\rangle)$,	$\mathcal{X}_{\langle 0,2\rangle} = \{this \mapsto \{o_\epsilon\}, ob1 \mapsto \{o_1\}, ob2 \mapsto \{o_2\}\}$
$\langle 0,3\rangle$	*call*$(m, readOnce(ob1, \langle\rangle, \langle f_1\rangle))$,	$\mathcal{X}_{\langle 0,3\rangle} = \{this \mapsto \{o_\epsilon\}, ob1 \mapsto \{o_1\}, ob2 \mapsto \{o_2\}\}$
$\langle 0,4\rangle$	*call*$(m, readBlock(ob2, \langle\rangle, \langle f_2\rangle))$.	$\mathcal{X}_{\langle 0,4\rangle} = \{this \mapsto \{o_\epsilon\}, ob1 \mapsto \{o_1\}, ob2 \mapsto \{o_2\}\}$
$\langle 1,0\rangle$	$readBlock(this, \langle\rangle, \langle r\rangle) \leftarrow$	$\mathcal{X}_{\langle 1,0\rangle} = \mathcal{X}^{call}_{readBlock} = \{this \mapsto \{o_2\}\}$
$\langle 1,1\rangle$	$\ldots$,	$\mathcal{X}_{\langle 1,1\rangle} = \{this \mapsto \{o_2\}\}$
$\langle 1,2\rangle$	*call*$(b, while(this, \overline{inp}, \overline{out}))$,	$\mathcal{X}_{\langle 1,2\rangle} = \{this \mapsto \{o_2\}\}$
$\langle 1,3\rangle$	$r := res$.	$\mathcal{X}_{\langle 1,3\rangle} = \{this \mapsto \{o_2\}\}$
$\langle 2,0\rangle$	$while(this, \overline{inp}, \overline{out}) \leftarrow i \leq 0$.	$\mathcal{X}_{\langle 2,0\rangle} = \mathcal{X}^{call}_{while} = \{this \mapsto \{o_2\}\}$
$\langle 3,0\rangle$	$while(this, \overline{inp}, \overline{out}) \leftarrow i > 0$,	$\mathcal{X}_{\langle 3,0\rangle} = \mathcal{X}^{call}_{while} = \{this \mapsto \{o_2\}\}$
$\langle 3,1\rangle$	*call*$(b, if(\overline{inp}, \overline{out}))$.	$\mathcal{X}_{\langle 3,1\rangle} = \{this \mapsto \{o_2\}\}$
$\langle 4,0\rangle$	$if(this, \overline{inp}, \overline{out}) \leftarrow \ldots$	$\mathcal{X}_{\langle 4,0\rangle} = \mathcal{X}^{call}_{if} = \{this \mapsto \{o_2\}\}$
$\langle 5,0\rangle$	$if(this, \overline{inp}, \overline{out}) \leftarrow \ldots$,	$\mathcal{X}_{\langle 5,0\rangle} = \mathcal{X}^{call}_{if} = \{this \mapsto \{o_2\}\}$
$\langle 5,1\rangle$	*call*$(b, if\_c(this, \overline{inp}, \overline{out}))$.	$\mathcal{X}_{\langle 5,1\rangle} = \{this \mapsto \{o_2\}\}$
$\langle 6,0\rangle$	$if\_c(this, \overline{inp}, \overline{out}) \leftarrow$	$\mathcal{X}_{\langle 6,0\rangle} = \mathcal{X}^{call}_{if\_c} = \{this \mapsto \{o_2\}\}$
$\langle 6,1\rangle$	*call*$(m, readContent(this, \langle pos, incr\rangle, \langle f\rangle))$,	$\mathcal{X}_{\langle 6,1\rangle} = \{this \mapsto \{o_2\}\}$
$\langle 6,2\rangle$	$\ldots$.	$\mathcal{X}_{\langle 6,2\rangle} = \{this \mapsto \{o_2\}\}$
$\langle 7,0\rangle$	$readOnce(this, \langle\rangle, \langle r\rangle) \leftarrow$	$\mathcal{X}_{\langle 7,0\rangle} = \mathcal{X}^{call}_{readOnce} = \{this \mapsto \{o_1\}\}$
$\langle 7,1\rangle$	*call*$(m, readContent(this, \langle 0, lth\rangle, \langle f\rangle))$,	$\mathcal{X}_{\langle 7,1\rangle} = \{this \mapsto \{o_1\}\}$
$\langle 7,2\rangle$	*await* $f?$,	$\mathcal{X}_{\langle 7,2\rangle} = \{this \mapsto \{o_1\}\}$
$\langle 7,3\rangle$	$r := f$.*get*.	$\mathcal{X}_{\langle 7,3\rangle} = \{this \mapsto \{o_1\}\}$
$\langle 8,0\rangle$	$readContent(this, \langle pos, elems\rangle, \langle r\rangle) \leftarrow$	$\mathcal{X}_{\langle 8,0\rangle} = \mathcal{X}^{call}_{readContent} = \{this \mapsto \{o_1, o_2\}\}$
$\langle 8,1\rangle$	③Reader_$init(\langle fp, elems\rangle, \langle rd\rangle)$,	$\mathcal{X}_{\langle 8,1\rangle} = \{this \mapsto \{o_1, o_2\}, rd^{o_1} \mapsto \{o_{13}\}, rd^{o_2} \mapsto \{o_{23}\}\}$
$\langle 8,2\rangle$	*call*$(m, process(rd, \langle pos\rangle, \langle f\rangle))$,	$\mathcal{X}_{\langle 8,2\rangle} = \{this \mapsto \{o_1, o_2\}, rd^{o_1} \mapsto \{o_{13}\}, rd^{o_2} \mapsto \{o_{23}\}\}$
$\langle 8,3\rangle$	*await* $f?$,	$\mathcal{X}_{\langle 8,3\rangle} = \{this \mapsto \{o_1, o_2\}, rd^{o_1} \mapsto \{o_{13}\}, rd^{o_2} \mapsto \{o_{23}\}\}$
$\langle 8,4\rangle$	$r := f$.*get*.	$\mathcal{X}_{\langle 8,4\rangle} = \{this \mapsto \{o_1, o_2\}, rd^{o_1} \mapsto \{o_{13}\}, rd^{o_2} \mapsto \{o_{23}\}\}$
$\langle 9.0\rangle$	$process(this, \langle pos\rangle, \langle res\rangle) \leftarrow \ldots$	$\mathcal{X}_{\langle 9,0\rangle} = \mathcal{X}^{call}_{process} = \{this \mapsto \{o_{13}, o_{23}\}\}$

Figure 7. Points-to analysis results for the running example

Example 7 (points-to analysis on the running example)**.** Figure 7 shows (part of) the result of applying the points-to analysis to the running example with $k = 2$. This is the smallest k for which no information is lost when handling object names. We label each program point with a tuple that identifies it. Each set in the right column is the value of the constraint variable for the corresponding program point in the left column. As before, the allocation sites of the example are identified by means of function allocSite. In Figure 7 we have allocSite$(P, \langle 0, 1\rangle) =$ ①, allocSite$(P, \langle 0, 2\rangle) =$ ① and allocSite$(P, \langle 8, 1\rangle) =$ ③. As the running example does not have reference fields, $\mathcal{X}_{\mathcal{G}}$ is empty. Keeping track of the value of the *this* reference is crucial for the precision of the points-to analysis of the running example. All object creations use the object name(s) pointed to by *this* to generate new object names by adding the current allocation site. E.g., at allocation site ③, *this* may be either o_1 or o_2; the new object names created are o_{13} and o_{23}, respectively. Observe that we keep the calling context as a superscript to the variable such that rd^{o_1} denotes the abstract value for rd when *this* is

DOI: 10.1002/stvr

o_1. The value of $this$ within a method comes from the object name(s) for the variable used to call the method. ■

The next theorem states the soundness of the analysis. To that end, let us assume that object identifiers in the semantics are of the form $\langle Oid, s\rangle$, where Oid is a unique identifier and s is the (unbounded) allocation sequence for the object. The semantics can be easily adapted to these identifiers by setting the function $newRef$ for generating fresh object identifiers to $newRef(i, o_{this})$, where i is the program point where the object was created and o_{this} is the object identifier for $this$, i.e., $o_{this} = \langle Oid, l\rangle$. This function returns a unique reference as a pair $\langle Oid', l \oplus_\infty i\rangle$ ($\oplus_\infty$ stands for unbounded allocation site concatenation). In order to relate the concrete semantics to the results inferred by the points-to analysis, we use $name(o)$ to refer to the object name in $\mathcal{N}^k$ that represents the concrete object identifier o. Concretely, $name(o)$ is the longest suffix of length at most k of the unbounded allocation sequence of o encoded in the object identifier.

Theorem 1. *Let P be an IR program, $\mathcal{T} \equiv \mathcal{S}_0 \rightsquigarrow \cdots \rightsquigarrow \mathcal{S}_n$ a trace, and b an instruction at program point $\langle t, j\rangle$ in P. For every trace step $\mathcal{S}_k \overset{b}{\rightsquigarrow} \mathcal{S}_{k+1}$, $0 \leq k < n$, and for every object $\mathtt{ob}(o, C, h, \langle tv, \_\rangle, \_) \in \mathcal{S}_{k+1}$, the following holds:*

a) $name(o) \in \mathcal{X}_{\langle t,j\rangle}(this)$;
b) If $x \in dom(tv)$ is a local reference variable, $tv(x) \neq \mathsf{null}$ and $s = name(tv(x))$, then $s \in \mathcal{X}_{\langle t,j\rangle}(x^{name(o)})$;
c) If $f \in dom(h)$ is a reference field of class C, $h(f) \neq \mathsf{null}$ and $q = name(h(f))$, then $q \in \mathcal{X}_{\langle t,j\rangle}(name(o).f)$.

5. FIELD-SENSITIVE SIZE ANALYSIS FOR CONCURRENT OO PROGRAMS

The objective of size analysis is to infer *size abstractions* which allow reasoning on how the sizes of data change along a program's execution, which is fundamental for bounding the number of iterations that loops perform. Intuitively, the cost of executing a loop can be then obtained by multiplying the cost of each iteration by the number of iterations that it performs. This process is formalized by means of the recurrence equations presented in Section 6 which integrate the size relations computed in the section.

5.1. The Basic Size Analysis

We present the size analysis in two steps: we first recall the notion of size measure that maps variables and values to their sizes; and we then present an abstraction which compiles instructions into size constraints, keeping as much information on global data (i.e., fields) as possible, while still being sound in concurrent executions.

Recall that the language on which we develop our analysis is deliberately simplified so that it only considers numerical and reference types, and thus the size analysis that we present in this section will also be restricted to such types. However, our implementation supports other data-types, in particular *String* and *user-defined* algebraic data-types, that we omit for the sake of simplifying the formal presentation. In Section 5.2, we comment on the additional bits required to handle these types in the size analysis.

Size Measures. For numerical data, the size is the actual numerical value. On the contrary, references require a more sophisticated treatment. A commonly used size measure is *path-length* [41], which counts the number of elements of the longest chain of references that can be traversed through the initial object (e.g., length of a list, depth of a tree, etc.). However, in our context, objects are intended to simulate concurrent computing entities and not data structures. Thus, it is not common that they affect the number of iterations that loops perform. Therefore, ignoring their sizes is sound and precise enough in most cases. A slightly more precise abstraction distinguishes between the case in which a reference variable points to an object (size 1) or to *null*

DOI: 10.1002/stvr

(size 0). The size of a future variable is the same as the size of the value it holds. This is sound since such variables can be used only through *get*, and the instruction *get* blocks the execution until the variable has a value.

Next we define formally our notion of size measure. In the definition below, $tv(x) \in \mathbb{Z}$ (resp. $h(f) \in \mathbb{Z}$) means that the static type of x (resp. f) is integer. Similarly, $tv(x) \in Objects$ (resp. $h(f) \in Objects$) means that the static type of x (resp. f) is a reference type different from *null*. And finally $tv(x) = \mathit{null}$ (resp. $h(f) = \mathit{null}$) means that the reference variable x (resp. field) holds *null*.

Definition 7 (size measure)**.** *Given a configuration $\mathcal{S}$ and an object $\mathtt{ob}(o, C, h, \langle tv, \_\rangle, \mathcal{Q}) \in \mathcal{S}$, the size of $a \in dom(l)$, where $l \in \{tv, h\}$, is defined as follows:*

$$\alpha(a, \mathcal{S}) \quad = \quad \begin{cases} l(a) & \text{if } l(a) \in \mathbb{Z} \\ 1 & \text{if } l(a) \in Objects \\ 0 & \text{if } l(a) = \mathit{null} \\ v & \text{if } l(a) = \mathtt{fn},\ \mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S},\ v \in \mathbb{Z} \\ 1 & \text{if } l(a) = \mathtt{fn},\ \mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S},\ v \in Objects \\ 0 & \text{if } l(a) = \mathtt{fn},\ \mathtt{fut}(\mathtt{fn}, \mathit{null}) \in \mathcal{S} \end{cases}$$

Abstract Compilation. This section describes how to transform a program P into an abstract program P^α, which can be seen as an abstraction of P with respect to the chosen size measure α. The translation is based on replacing each instruction by (linear) *constraints* which describe its behaviour with respect to the size measure. For example, the instruction $x := \mathit{new}\ C$ can be replaced by the constraint $x = 1$ which indicates that the reference variable is different from *null*. The fact that the formula takes this form is due to the previous choice of size measure for references in Definition 7. In order to simplify the presentation, when it is clear from the context, the same name is used for the original variables (possibly primed or subscripted) and their symbolic sizes. That is, given an integer variable x, the name x is also used in the abstract compilation to denote its value.

An important issue in the presented setting is to be able to obtain relations between the size of a variable at different program points. For example, in the size analysis of $x := x + 1$, the interest is in the relation "the value of x after the instruction is equal to the value of x before the instruction plus 1". This important piece of information can be obtained by using a *Static Single Assignment* (SSA) transformation, which, together with the abstract compilation, produces the constraint $x' = x + 1$, where x and x' refer to, respectively, the value of x before and after executing the instruction. To implement the SSA transformation, a mapping ρ of variable names to new variable names (constraint variables) is maintained. Such mapping is referred to as a *renaming*. The expression $\rho[\bar{x} \mapsto \bar{y}]$ denotes the update performed by ρ, such that it maps variables $\bar{x}$ to the new variables $\bar{y}$.

Modeling shared memory is a main challenge in static analysis of OO programs. Our starting point is [5], which models fields as local variables when the field to be tracked satisfies two soundness conditions: (1) its memory location does not change; and (2) it is always accessed through the same reference (i.e., not through aliases). Both conditions can often be proven statically and the transformation of fields into local variables can then be applied for many fragments of the program. If we ignore concurrency, this approach could be directly adopted for our language. However, concurrency introduces new challenges.

Example 8 (treatment of fields at release points)**.** Consider the loop in the readBlock method in Figure 1. Ignoring the *await* instruction, the above soundness conditions (1) and (2) hold for the field blockS, and hence, we can track it as if it was a local variable. In a concurrent setting, however, while readBlock is executing, another task in the same object might modify the field blockS. Therefore, when analyzing readBlock, we cannot assume that the value of blockS is locally trackable. For instance, readBlock might introduce non-termination if we add a method void p() {blockS = blockS − 2; } to class FileIS. When the *await* is executed inside the loop, method p might change the value of blockS to a non-positive value, and thus the loop counter i would not decrement. ■

DOI: 10.1002/stvr

	B	$\alpha_\rho(B)$	ρ'
(1)	$x\ op\ y$	$\rho(x)\ op\ \rho(y)$	$\rho'{=}\rho$
(2)	$x\ op'\ y$	_	$\rho'{=}\rho$
(3)	*null* $\mid x \mid this.f \mid n$	$0 \mid \rho(x) \mid \rho(f) \mid n$	$\rho'{=}\rho$
(4)	*await* x?	$\bot$	$\rho'{=}\rho[\bar{f}_C \mapsto fresh(\bar{f}_C)]$
(5)	$x := y.$*get* $\mid x := e$	$\rho'(x){=}\rho(y) \mid \rho'(x){=}\alpha_\rho(e)$	$\rho'{=}\rho[x \mapsto fresh(x)]$
(6)	$this.f := y$	$\rho'(f){=}\rho(y)$	$\rho'{=}\rho[f \mapsto fresh(f)]$
(7)	$x :=$ *new* C	$\rho'(x){=}1$	$\rho'{=}\rho[x \mapsto fresh(x)]$
(8)	*call*$(b, q(rec, \bar{x}, \bar{y}))$	*call*$(b, q(\rho(rec), \rho(\bar{x} \cdot \bar{f}_C), \rho'(\bar{y} \cdot \bar{f}_C)))$	$\rho'{=}\rho[\bar{y}{\cdot}\bar{f}_C{\mapsto}fresh(\bar{y}{\cdot}\bar{f}_C)]$
(9)	*call*$(m, q(rec, \bar{x}, y))$	*call*$(m, q(\rho(rec), \rho(\bar{x}), \rho'(y)))$	$\rho'{=}\rho[y \mapsto fresh(y)]$
(10)	$otherwise$	$true$	$\rho'{=}\rho$

where in case (1) $op \in \{<, >, =, \geq, \leq, +, -\}$ and in (2) $op' \in \{\neq, *, /\}$

Figure 8. Abstract compilation. $\texttt{ABST}(B_{k:i}, \rho){=}\langle\alpha_\rho(B_{k:i}), \rho'\rangle$

In order to handle fields, we need to identify the program points at which the shared memory might be modified by other tasks. This can happen when: (1) an *await* is explicitly executed, and thus allows other tasks (of the same object) to run; and (2) an asynchronous invocation is made, and until the called method starts its execution, the fields of the called object might be modified by other tasks. We refer to such program points as *release points*. The above observation suggests that in a sequence of instructions not including *await*, the shared memory can be tracked locally. However, the values in the shared memory when a method starts to execute may not be identical to those when it was called. We first present a safe abstraction which loses all information at release points, and, in a second step, in Section 5.3, we discuss accuracy improvements at such release points.

In the size analysis, an abstract state is a set of linear constraints whose solutions define possible concrete states. This representation allows describing relations that are essential for inferring cost and proving termination, e.g., the size of x decreases by 1 in two consecutive states. The building blocks for this representation are constraints that describe the effect of each instruction b on a given state.

Definition 8 (abstract compilation)**.** *Let B be an instruction or a guard. We define its* abstract compilation $\texttt{ABST}(B, \rho)$ *w.r.t. a mapping ρ and a symbolic size measure α_ρ as $\langle\alpha_\rho(B), \rho'\rangle$, where $\langle\alpha_\rho(B), \rho'\rangle$ is computed according to the abstraction in Figure 8.*

Let us describe the abstraction of some selected instructions. First, note that, except for method/block calls and *await* instructions, $\alpha_\rho(B)$ returns a constraint and ρ' is a new mapping that refers to the sizes in the state after executing B. In Figure 8, given the variables $x_1, \ldots, x_n$ (respectively the fields $f_1, \ldots, f_n$), function $fresh(x_1, \ldots, x_n)$ (respectively $fresh(f_1, \ldots, f_n)$) returns n fresh variable names (respectively field names). Recall that $\rho'(x)$ (respectively $\rho(x)$) refers to the size of x *after* (respectively *before*) executing the instruction. For the assignments, row 5 of Figure 8, the instruction $x := e$ is abstracted into the equality $\rho'(x) = \alpha_\rho(e)$, where $\alpha_\rho(e)$ is the size of e w.r.t. ρ. The abstraction of *await* at row 4 "forgets" sizes of those fields $\bar{f}_C$ of class C. This is because fields might be updated by other methods that take the control when the current task is suspended. When abstracting a call to a block in row 8, the class fields $\bar{f}_C$ are added as arguments in order to track their values within the block called. On the contrary, when abstracting calls to methods (row 9) the fields are not added because their values at call time might not be equal to their values when the method actually starts its execution. Since we use linear constraints, non-linear arithmetic expressions (row 2) are abstracted to a fresh constraint variable "_" that represents any value. A program P is transformed into an abstract program P^α, that approximates its behaviour w.r.t. a size measure, by abstracting its rules as follows.

Definition 9 (abstract compilation of a rule)**.** *Let $r \equiv m(this, \bar{x}, \bar{y}) \leftarrow g, b_1, \ldots, b_n$ be a renaming of a rule in P and let ρ_0 be a renaming over $vars(r) \cup \bar{f}_C$. The abstract compilation of r w.r.t. a size measure α is $r^\alpha \equiv m(\rho_0(this), \bar{I}, \bar{O}) \leftarrow g^\alpha, b_1^\alpha, \ldots, b_n^\alpha$ where:*

DOI: 10.1002/stvr

1. $\mathtt{ABST}(g, \rho_0){=}\langle g^\alpha, \rho_1\rangle$, $\mathtt{ABST}(b_i, \rho_i){=}\langle b_i^\alpha, \rho_{i+1}\rangle$, $1 \leq i \leq n$;
2. *If* m *is a block then* $\bar{I}{=}\rho_0(\bar{x}{\cdot}\bar{f}_C)$ *and* $\bar{O}{=}\rho_{n+1}(\bar{y}{\cdot}\bar{f}_C)$*; and*
3. *If* m *is a method then* $\bar{I}{=}\rho_0(\bar{x})$ *and* $O{=}\rho_{n+1}(y)$.

The size abstraction $\mathtt{ABST}(r)$ *for the rule* r *is* $g^\alpha \wedge b_1^\alpha \wedge \ldots \wedge b_n^\alpha$ *and* $\mathcal{C}(\mathtt{ABST}(r)) = g^\alpha \wedge \bigwedge_{i=1}^n \{b_i^\alpha \mid b_i^\alpha$ *is a linear constraint.*$\}$

Note that, according to row 8 in Figure 8, when abstracting a rule corresponding to a block (item 2 in Definition 9), we add the fields $\bar{f}_C$ to the input parameters. In what follows, we sometimes represent conjunctions of linear constraints $\varphi_1 \wedge \ldots \wedge \varphi_n$ as sets of the form $\{\varphi_1, \ldots, \varphi_n\}$.

Once the abstract compilation of the rule has finished, the renaming ρ_{n+1} coming from $\mathtt{ABST}(b_n, \rho_n) = \langle \alpha_{\rho_n}(b_n), \rho_{n+1}\rangle$ is applied to the output variables and fields. However, when abstracting a method rule (item 3 in Definition 9) fields are not considered in the head of the method. This matches row 9 in Figure 8.

Example 9 (abstract compilation for the running example)**.** Let us show the abstract compilation of the rule *if_c* of Figure 2. We use $\overline{inp}$, $\overline{out}$ and $\overline{F}$ to denote, respectively, the input parameters $res, i, incr, pos$, the output parameters $res'', i'', incr', pos''$ and the fields $fp, lth, blockS$. The substitution ρ_0 stands for the identity mapping on $this, \overline{inp}$, $\overline{out}$ and $\overline{F}$. The number to the right of each instruction indicates the row in Figure 8 used to compute the abstract compilation.

$if\_c(this, \langle\overline{inp}, \overline{F}\rangle, \langle\overline{out}, \overline{F}''\rangle) \leftarrow$	ρ_0	
ⓐ *call*$(m, readContent(this, \langle pos, incr\rangle, \langle f'\rangle))$,	$fresh(f) = f', \rho_1 = \rho_0[f \mapsto f']$	(9)
ⓑ $\bot$,	$fresh(\overline{F}) = \overline{F}', \rho_2 = \rho_1[\overline{F} \mapsto \overline{F}']$	(4)
$v' = f'$,	$fresh(v) = v', \rho_3 = \rho_2[v \mapsto v']$	(5)
$res' = res + v'$	$fresh(res) = res', \rho_4 = \rho_3[res \mapsto res']$	(5)
$i' = i - incr$	$fresh(i) = i', \rho_5 = \rho_4[i \mapsto i']$	(5)
$pos' = pos + incr$	$fresh(pos) = pos', \rho_6 = \rho_5[pos \mapsto pos']$	(5)
ⓒ *call*$(b, while(this, \langle res', i', incr, pos', \overline{F}'\rangle,$	$\rho_7 = \rho_6[res{\mapsto}res'', i{\mapsto}i'', incr{\mapsto}incr',$	
$\langle res'', i'', incr', pos'', \overline{F}''\rangle))$.	$pos \mapsto pos'', \bar{F} \mapsto \bar{F}'']$	(8)

In ⓒ, function $fresh(res, i, inc, pos, \bar{F})$ returns $res'', i'', incr', pos'', \bar{F}''$. Furthermore, the output parameters of *if_c* are the result of applying ρ_7 to $res, i, incr, pos, fp, lth, blockS$. According to the abstraction in Figure 8, at ⓑ *await* is abstracted to $\bot$ and the information on fields is lost. At ⓒ the fields are added to the call in order to keep track of their values, however, when calling a method at ⓐ, the abstraction "forgets" this information.

Let us consider now method readBlock, whose abstract compilation is:

$readBlock(this, \langle\rangle, \langle r'\rangle) \leftarrow$	ρ_0	
$res' = 0$,	$fresh(res) = res', \rho_1 = \rho_0[res \mapsto res']$	(5)
$i' = lth$,	$fresh(i) = i', \rho_2 = \rho_1[i \mapsto i']$	(5)
$incr' = 0$,	$fresh(incr) = incr', \rho_3 = \rho_2[incr \mapsto incr']$	(5)
$pos' = 0$	$fresh(pos) = pos', \rho_4 = \rho_3[pos \mapsto pos']$	(5)
call$(b, while(this, \langle res', i', incr', pos', \overline{F}\rangle,$	$\rho_5 = \rho_4[res \mapsto res'', i \mapsto i'',$	
$\langle res'', i'', incr'', pos'', \overline{F}'\rangle))$,	$incr \mapsto incr'', pos \mapsto pos'', \bar{F} \mapsto \bar{F}']$	(8)
$r' = res''$.	$fresh(r) = r', \rho_6 = \rho_5[r \mapsto r']$	(5)

In ρ_5, function $fresh(res, i, incr, pos, \bar{F})$ returns $res'', i'', incr'', pos'', \bar{F}'$. According to Definition 9, the head of the rule does not contain fields while the call to *while* does. ■

An abstract program P^α basically abstracts the behaviour of the original program with respect to a size measure α. An abstract state has the form $\mathcal{A} \circ \phi$, where $\mathcal{A} \equiv \{a_1^\alpha, \ldots, a_n^\alpha\}$, a_i^α is an *abstract object* of the form $\langle \bar{b}^\alpha, \bar{\rho}\rangle$ and ϕ is a linear constraint. In order to formalize the operational

DOI: 10.1002/stvr

$$(1)_\alpha \ \frac{p(this', \bar{x}', \bar{y}') \leftarrow g^\alpha, b'^\alpha_1, \ldots, b'^\alpha_n \circ \rho_0 \cdots \rho_{n+1} \ll^\alpha_p P^\alpha, g^\alpha \wedge \phi \not\models false}{\{\langle \mathsf{call}(b, p(this, \bar{x}, \bar{y})) \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho} \rangle | \mathcal{A}\} \circ \phi \rightsquigarrow_\alpha \{\langle b'^\alpha_1 \cdots b'^\alpha_n \cdot \bar{b}^\alpha, \rho_1 \cdots \rho_{n+1} \cdot \bar{\rho} \rangle | \mathcal{A}\} \circ \phi \wedge g^\alpha \wedge this = this' \wedge \bar{x}' = \bar{x} \wedge \bar{y}' = \bar{y}}$$

$$(2)_\alpha \ \frac{p(rec', \bar{x}', y') \leftarrow b'^\alpha_1, \ldots, b'^\alpha_n \circ \rho_1 \cdots \rho_{n+1} \ll^\alpha_p P^\alpha}{\{\langle \mathsf{call}(m, p(rec, \bar{x}, y)) \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho} \rangle | \mathcal{A}\} \circ \phi \rightsquigarrow_\alpha \{\langle \bar{b}^\alpha, \bar{\rho} \rangle, \langle b'^\alpha_1 \cdots b'^\alpha_n, \rho_1 \cdots \rho_{n+1} \rangle | \mathcal{A}\} \circ \phi \wedge rec' = rec \wedge \bar{x}' = \bar{x} \wedge y' = y}$$

$$(3)_\alpha \ \frac{\varphi \wedge \phi \not\models false}{\{\langle \varphi \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho} \rangle | \mathcal{A}\} \circ \phi \rightsquigarrow_\alpha \{\langle \bar{b}^\alpha, \bar{\rho} \rangle | \mathcal{A}\} \circ \phi \wedge \varphi}$$

$$(4)_\alpha \ \frac{}{\{\langle \bot \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho} \rangle | \mathcal{A}\} \circ \phi \rightsquigarrow_\alpha \{\langle \bar{b}^\alpha, \bar{\rho} \rangle | \mathcal{A}\} \circ \phi}$$

$$(5)_\alpha \ \frac{}{\{\langle \bot \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho} \rangle | \mathcal{A}\} \circ \phi \rightsquigarrow_\alpha \{\langle \bot \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho} \rangle | \mathcal{A}\} \circ \phi}$$

$$(6)_\alpha \ \frac{}{\{\langle \epsilon, \rho \rangle | \mathcal{A}\} \circ \phi \rightsquigarrow_\alpha \{\epsilon | \mathcal{A}\} \circ \phi}$$

$$(7)_\alpha \ \frac{}{\mathcal{A} \circ \phi \rightsquigarrow_\alpha \mathcal{A} \circ \phi}$$

Figure 9. Semantics of abstract programs

semantics for an abstract program, we modify the presentation of the abstract rules, by storing also the renamings computed in Definition 8. Therefore, abstract rules are now of the form $p(rec, \bar{x}, \bar{y}) \leftarrow g^\alpha, b^\alpha_1, \ldots, b^\alpha_n \circ \rho_0 \cdots \rho_{n+1}$, where $\rho_0, \ldots, \rho_{n+1}$ is the tuple of all renamings that were used during the abstract compilation of that possibly renamed specific rule.

The operational semantics for an abstract program is given in a rewriting-based style in Figure 9. Such semantics simply accumulates the constraints (when possible) and proceeds to execute the calls in the body of the rules. Now, let us relate the abstract semantics of Figure 9 with the operational semantics of the IR shown in Figure 3. Rules $(1)_\alpha$ and $(2)_\alpha$ correspond to calls to blocks and methods respectively. The notation $p(this', \bar{x}', \bar{y}') \leftarrow g^\alpha, b'^\alpha_1, \ldots, b'^\alpha_n \circ \rho_0 \cdots \rho_{n+1} \ll^\alpha_p P^\alpha$ in rule $(1)_\alpha$ stands for the abstract compilation of a possibly renamed rule from P starting from a fresh renaming ρ_0, i.e., the initial renaming ρ_0 in Definition 9 must apply variables to fresh variables. Hence all variables in $p(this', \bar{x}', \bar{y}') \leftarrow g^\alpha, b'^\alpha_1, \ldots, b'^\alpha_n$ are completely fresh. The intended meaning of $\ll^\alpha_p$ in rule $(2)_\alpha$ is similar to that in rule $(1)_\alpha$.

Rules $(4)_\alpha$ and $(5)_\alpha$ correspond to the abstraction of an *await* instruction since *await* is abstracted to $\bot$ (see Figure 8). In particular, when an *await* instruction is evaluated to $true$ (rule (9) of Figure 3) the execution proceeds. This is simulated by rule $(4)_\alpha$. Similarly, if the evaluation of the *await* instruction fails (rule (10) of Figure 3), then the context of the suspended task is introduced in the queue of pending tasks for the current object. Hence, in rule $(5)_\alpha$ the abstraction $\bot$ of the corresponding *await* is kept until the associated *await* succeeds and rule $(4)_\alpha$ can be applied.

Rules (6) and (7) of Figure 3 handle, respectively, those cases in which the execution of a method or a block call has finished. This fact is captured by rule $(6)_\alpha$ by simply removing the abstract execution context, as it is done in rules (6) and (7). Note also that, after applying rules (6) and (7), since the execution context is ϵ, rule (11) can then be applied to select a new task from the queue of the corresponding object for its execution. In the abstract semantics, such task exists inside the abstract configuration and thus it can be selected at any time of the computation. Thus, rule (11) corresponds to rule $(7)_\alpha$ in the abstract semantics. Finally, rule $(3)_\alpha$ basically accumulates constraints when it is possible. The notation $\varphi \wedge \phi \not\models false$ means that $\varphi \wedge \phi$ is satisfiable.

DOI: 10.1002/stvr

main	main_1
$main(this, \langle\rangle, \langle\rangle) \leftarrow$ $a := \mathit{new}\ A(),$ $z := 4,$ $\mathit{call}(m, p(a, \langle z\rangle, \langle f\rangle)),$ $\mathit{await}\ f?.$	$main(this_1, \langle\rangle, \langle\rangle) \leftarrow$ $a_1 := \mathit{new}\ A(),$ $z_1 := 4,$ $\mathit{call}(m, p(a_1, \langle z_1\rangle, \langle f_1\rangle))$ $\mathit{await}\ f_1?.$
p	p_3
$p(this, \langle w\rangle, \langle r\rangle) \leftarrow$ $r := w + 1.$	$p(this_3, \langle w_3\rangle, \langle r_3\rangle) \leftarrow$ $r_3 := w_3 + 1.$

Figure 10. An IR program P (left) together with a possibly renaming (right)

main_1	main_1^α	ρ
$main(this_1, \langle\rangle, \langle\rangle) \leftarrow$ $a_1 := \mathit{new}\ A(),$ $z_1 := 4,$ $\mathit{call}(m, p(a_1, \langle z_1\rangle, \langle f_1\rangle)),$ $\mathit{await}\ f_1?.$	$main(this_2, \langle\rangle, \langle\rangle) \leftarrow$ $a_2 = 1,$ $z_2 = 4$ $\mathit{call}(m, p(a_2, \langle z_2\rangle, \langle f_2\rangle)),$ $\bot.$	$\rho_0(this_1) = this_2$ $\rho_1 = \rho_0[a_1 \mapsto a_2]$ $\rho_2 = \rho_1[z_1 \mapsto z_2]$ $\rho_3 = \rho_2[f_1 \mapsto f_2]$ $\rho_4 = \rho_3$
p_3	p_3^α	ρ
$p(this_3, \langle w_3\rangle, \langle r_3\rangle) \leftarrow$ $r_3 := w_3 + 1.$	$p(this_4, \langle w_4\rangle, \langle r_4\rangle) \leftarrow$ $r_4 = w_4 + 1.$	$\rho_5(this_3){=}this_4, \rho_5(w_3){=}w_4$ $\rho_6 = \rho_5[r_3 \mapsto r_4]$

Figure 11. A renamed program (left) together with its abstract compilation (right)

The following example briefly shows the correspondence between concrete and abstract traces. For simplicity, we ignore renamings in the abstract states and focus on the abstract rules that are applied.

Example 10 (correspondence between concrete and abstract traces)**.** Consider the IR program to the left of Figure 10, which is composed of two methods, main and p. Let us start by showing some steps of the concrete execution of the program in Figure 10. Let $\mathcal{S}_0$ be the initial concrete configuration:

$$\mathcal{S}_0 \equiv \{\mathtt{ob}(\mathsf{main}, \bot, \bot, \langle tv_m, \mathit{call}(b, main(this, \langle\rangle, \langle\rangle))\rangle, \emptyset)\}$$

From $\mathcal{S}_0$, by applying rule (4) in Figure 3 with $\mathsf{main}_1 \ll^{\langle\rangle}_{\mathsf{main}} P$, where main_1 is defined in Figure 10 (right-top), we get:

$$\mathcal{S}_1{\equiv}\{\mathtt{ob}(\mathsf{main}, \bot, \bot, \langle tv_m^1, a_1 := \mathit{new}\ A() \cdot z_1 := 4 \cdot \mathit{call}(m, p(a_1, \langle z_1\rangle, \langle f_1\rangle)) \cdot \mathit{await}\ f_1?\rangle, \emptyset)\}$$

where $tv'_m = newEnv(vars(\mathsf{main}_1))$ and $tv_m^1 = tv_m \cup tv'_m[this_1 \mapsto this]$. From $\mathcal{S}_1$, by applying consecutively rules (3) and (1) in Figure 3, we generate the new concrete state:

$$\mathcal{S}_2{\equiv}\{\mathtt{ob}(\mathsf{main}, \bot, \bot, \langle tv_m^2, \mathit{call}(m, p(a_1, \langle z_1\rangle, \langle f_1\rangle)) \cdot \mathit{await}\ f_1?\rangle, \emptyset), \mathtt{ob}(o_a, A, h_a, \epsilon, \emptyset)\}$$

where $tv_m^2 = tv_m^1[a_1 \mapsto o_a, z_1 \mapsto 4]$. From $\mathcal{S}_2$, selecting $\mathsf{p}_3 \ll_p P$, where p_3 is defined in Figure 10 (right-bottom), we can apply rule (5) to obtain:

$$\mathcal{S}_3{\equiv}\{\mathtt{ob}(\mathsf{main}, \bot, \bot, \langle tv_m^3, \mathit{await}\ f_1?\rangle, \emptyset), \mathtt{ob}(o_a, A, h_a, \epsilon, \{\langle tv_a, r_3 := w_3 + 1\rangle\}), \mathtt{fut}(\mathtt{fn}, \bot)\}$$

where $tv_m^3 = tv_m^2[f_1 \mapsto \mathtt{fn}]$, $tv'_a = newEnv(vars(\mathsf{p}_3))$, $\mathtt{fn} = newFut()$ and $tv_a = tv'_a[this_3 \mapsto o_a, w_3 \mapsto 4, \mathtt{ret} \mapsto (\mathtt{fn}, r_3)]$. Now we apply rule (10) to $\mathcal{S}_3$ on object main. Then, since the future variable is not ready, the task execution context is introduced in the queue of object main:

$$\mathcal{S}_4{\equiv}\{\mathtt{ob}(\mathsf{main}, \bot, \bot, \epsilon, \{\langle tv_m^3, \mathit{await}\ f_1?\rangle\}), \mathtt{ob}(o_a, A, h, \epsilon, \{\langle tv_a, r_3 := w_3 + 1\rangle\}), \mathtt{fut}(\mathtt{fn}, \bot)\}$$

DOI: 10.1002/stvr

Let us apply rule (11) on object o_a in $\mathcal{S}_4$ in order to extract the unique task from its queue, followed by rule (1). The resulting concrete state is:

$$\mathcal{S}_5 \equiv \{\mathtt{ob}(\mathsf{main}, \bot, \bot, \epsilon, \{\langle tv_m^3, \textit{await}\ f_1?\rangle\}), \mathtt{ob}(o_a, A, h, \langle \epsilon, tv_a^1\rangle, \emptyset), \mathtt{fut}(\mathtt{fn}, \bot)\}$$

where $tv_a^1 = tv_a[r_3 \mapsto 5]$. On the concrete state $\mathcal{S}_5$ it holds that $\mathtt{ret} \in dom(tv_a^1)$, $tv_a^1(\mathtt{ret}) = (\mathtt{fn}, r_3)$, $tv_a^1(r_3) = 5$, thus rule (7) can be used to fix the value of the future variable as follows:

$$\mathcal{S}_6 \equiv \{\mathtt{ob}(\mathsf{main}, \bot, \bot, \epsilon, \{\langle tv_m^3, \textit{await}\ f_1?\rangle\}), \mathtt{ob}(o_a, A, h, \epsilon, \emptyset), \mathtt{fut}(\mathtt{fn}, 5)\}$$

and afterwards apply rules (11), (9) and (6) and obtain:

$$\mathcal{S}_7 \equiv \{\mathtt{ob}(\mathsf{main}, \bot, \bot, \epsilon, \emptyset), \mathtt{ob}(o_a, A, h, \epsilon, \emptyset), \mathtt{fut}(\mathtt{fn}, 5)\}$$

Now, in order to show the correspondence between the concrete and the abstract semantics, we start from the initial abstract state:

$$\mathcal{A}_0 \equiv \{\langle \textit{call}(b, main(this, \langle\rangle, \langle\rangle)), \_\rangle\} \circ true$$

By applying rule $(1)_\alpha$, with $\mathsf{main}_1^\alpha \ll^\alpha_{\mathsf{main}} P^\alpha$, to $\mathcal{A}_0$, where main_1^α is defined in Figure 11 (middle-top), we reach the abstract state:

$$\mathcal{A}_1 \equiv \{\langle a_2 = 1 \cdot z_2 = 4 \cdot \textit{call}(m, p(a_2, \langle z_2\rangle, \langle f_2\rangle)) \cdot \bot, \rho_0 \cdots \rho_4\rangle\} \circ this_2 = this$$

From $\mathcal{A}_1$, by applying twice rule $(3)_\alpha$ we can compute:

$$\mathcal{A}_2 \equiv \{\langle \textit{call}(m, p(a_2, \langle z_2\rangle, \langle f_2\rangle)) \cdot \bot, \rho_2 \cdots \rho_4\rangle\} \circ \underbrace{this_2 = this \wedge a_2 = 1 \wedge z_2 = 4}_{\varphi}$$

Similarly, by applying rule $(2)_\alpha$ to $\mathcal{A}_2$ with $\mathsf{p_3}^\alpha \ll^\alpha_p P^\alpha$, where p_3^α is defined in Figure 11 (middle-bottom) we have:

$$\mathcal{A}_3 \equiv \{\langle \bot, \rho_3 \cdot \rho_4\rangle, \langle r_4 = w_4 + 1, \rho_5 \cdot \rho_6\rangle\} \circ \varphi \wedge this_4 = a_2 \wedge w_4 = z_2 \wedge r_4 = f_2$$

The point now is that we can apply rule $(5)_\alpha$ on the abstract state $\mathcal{A}_3$ and obtain exactly the same abstract state $\mathcal{A}_4 \equiv \mathcal{A}_3$, i.e., we delay the abstract trace until the *await* succeeds. On the abstract state $\mathcal{A}_4$ we can apply rule $(7)_\alpha$ followed by rule $(3)_\alpha$, to compute:

$$\mathcal{A}_5 \equiv \{\langle \bot, \rho_3 \cdot \rho_4\rangle, \langle \epsilon, \rho_6\rangle\} \circ \varphi \wedge this_4 = a_2 \wedge w_4 = z_2 \wedge r_4 = f_2 \wedge r_4 = w_4 + 1$$

Note that rules $(7)_\alpha$ and $(3)_\alpha$ correspond to the application of rules (11) and (1) respectively. For the case of the abstract state $\mathcal{A}_5$, it is enough to apply, first rule $(6)_\alpha$ (corresponding to the application of rule (7)) to compute

$$\mathcal{A}_6 \equiv \{\langle \bot, \rho_3 \cdot \rho_4\rangle, \epsilon\} \circ \varphi \wedge this_4 = a_2 \wedge w_4 = z_2 \wedge r_4 = f_2 \wedge r_4 = w_4 + 1$$

Afterwards the application of rule $(7)_\alpha$, which corresponds to the application of rule (11) in the concrete semantics, followed by rule $(4)_\alpha$, which corresponds to the application of rule (9), results in $\{\langle \epsilon, \rho_4\rangle, \epsilon\} \circ \varphi \wedge this_4 = a_2 \wedge w_4 = z_2 \wedge r_4 = f_2 \wedge r_4 = w_4 + 1$. Finally with the application of rule $(6)_\alpha$, which corresponds rule (6) in the concrete semantics, we have:

$$\mathcal{A}_7 \equiv \{\epsilon, \epsilon\} \circ \varphi \wedge this_4 = a_2 \wedge w_4 = z_2 \wedge r_4 = f_2 \wedge r_4 = w_4 + 1$$

■

We now establish the *soundness* of the abstract compilation with respect to the chosen size measure α. Intuitively, we prove that the size of the variables in a given concrete trace is computed in its corresponding abstract trace. As notation, given an object $a \equiv \mathtt{ob}(o, C, h, \langle tv, \bar{b}\rangle, \mathcal{Q})$, we say that $\langle tv, \bar{b}\rangle$ is its *active task* (denoted by $active(a)$) and we define $pending(a) = \{tk \mid tk \in \mathcal{Q}\}$ as the set of pending tasks of a. Thus, we define the set of tasks for an object a, denoted as $tasks(a)$, as $\{active(a)\} \cup pending(a)$. Finally, given an state $\mathcal{S} = \{a_1, \ldots, a_n\}$, we define the set of tasks in $\mathcal{S}$, denoted as $tasks(\mathcal{S})$, as $\cup_{i=1}^n tasks(a_i)$.

DOI: 10.1002/stvr

Definition 10 (relation between abstract and concrete states)**.** *Let $\mathcal{S}$ be a concrete state and α a size measure. We say that an abstract state $\mathcal{A} \circ \phi$* approximates *$\mathcal{S}$, denoted as $\mathcal{A} \circ \phi \approx \mathcal{S}$ if an only if:*

1. *ϕ is satisfiable;*
2. *for all $\langle tv, b_1 \cdots b_n \rangle \in tasks(\mathcal{S})$, from an object $\mathtt{ob}(o, C, h, \_, \_) \in \mathcal{S}$, there exists $\langle b_1^\alpha \cdots b_n^\alpha, \rho_1 \cdots \rho_{n+1} \rangle \in \mathcal{A}$ and an assignment $\sigma : vars(tv) \cup \bar{f}_C \mapsto \mathbb{Z}$ such that $\mathtt{ABST}(b_i, \rho_i) = \langle b_i^\alpha, \rho_{i+1} \rangle$, $1 \leq i \leq n$, $\sigma \models \phi$ and:*
   - *For all $x \in dom(tv)$ it holds that $\sigma(\rho_1(x)) = \alpha(x, \mathcal{S})$;*
   - *For all $f \in dom(h)$ it holds that $\sigma(\rho_1(f)) = \alpha(f, \mathcal{S})$;*
   - *For all $y \in dom(tv)$, if $tv(y) = \mathtt{fn}$, $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}$ and $v \neq \bot$, then $\sigma(\rho_1(y)) = \alpha(y, \mathcal{S})$.*

Example 11 (equivalence between concrete and abstract traces)**.** Consider the concrete and abstract states $\mathcal{S}_2$ and $\mathcal{A}_2$ of Example 10:

$$\mathcal{S}_2 \equiv \{\mathtt{ob}(\mathsf{main}, \bot, \bot, \langle tv_m^2, \mathit{call}(m, p(a_1, \langle z_1 \rangle, \langle f_1 \rangle)) \cdot \mathit{await}\ f_1?\rangle, \emptyset), \mathtt{ob}(o_a, A, h_a, \epsilon, \emptyset)\}$$

$$\mathcal{A}_2 \equiv \{\langle \mathit{call}(m, p(a_2, \langle z_2 \rangle, \langle f_2 \rangle)) \cdot \bot, \rho_2 \cdots \rho_4 \rangle\} \circ \underbrace{this_2 = this \wedge a_2 = 1 \wedge z_2 = 4}_{\varphi}$$

Then, we have that $\mathtt{ABST}(\mathit{call}(m, p(a_1, \langle z_1 \rangle, \langle f_1 \rangle), \rho_2) = \langle \mathit{call}(m, p(a_2, \langle z_2 \rangle, \langle f_2 \rangle)), \rho_3 \rangle$ and $\mathtt{ABST}(\mathit{await}\ f_1?, \rho_3) = \langle \bot, \rho_4 \rangle$, where ρ_2 and ρ_3 are defined in Figure 11. Furthermore, let σ be an assignment such that $\sigma \models \varphi$, then $\sigma(z_2) = 4$. Finally note that $z_1 \in dom(tv_m^2)$ and $\sigma(\rho_2(z_1)) = \sigma(z_2) = 4 = tv_m^2(z_1) = \alpha(z_1, \mathcal{S}_2)$. Hence $\mathcal{A}_2 \circ \varphi \approx \mathcal{S}_2$.
Now, let us focus on $\mathcal{S}_5$ and $\mathcal{A}_5$ of Example 10.

$$\mathcal{S}_5 \equiv \{\mathtt{ob}(\mathsf{main}, \bot, \bot, \epsilon, \{\langle tv_m^3, \mathit{await}\ f_1?\rangle\}), \mathtt{ob}(o_a, A, h, \langle \epsilon, tv_a^1 \rangle, \emptyset), \mathtt{fut}(\mathtt{fn}, \bot)\}$$

$$\mathcal{A}_5 \equiv \{\langle \bot, \rho_3 \cdot \rho_4 \rangle, \langle \epsilon, \rho_6 \rangle\} \circ \varphi \wedge this_4 = a_2 \wedge w_4 = z_2 \wedge r_4 = f_2 \wedge r_4 = w_4 + 1$$

Let σ' be a valuation that satisfies $\varphi \wedge this_4 = a_2 \wedge w_4 = z_2 \wedge r_4 = f_2 \wedge r_4 = w_4 + 1$. Then $\sigma'(w_4) = 4$. Now in order to prove that $\mathcal{A}_5$ approximates $\mathcal{S}_5$, we need to ensure also that $\sigma'(\rho_6(w_3)) = \alpha(w_3, \mathcal{S}_5)$ and $\sigma'(\rho_6(r_3)) = \alpha(r_3, \mathcal{S}_5)$. But this follows from $\sigma'(\rho_6(w_3)) = \sigma'(w_4) = 4 = tv_a^1(w_3) = \alpha(w_3, \mathcal{S}_5)$ and $\sigma'(\rho_6(r_3)) = \sigma'(r_4) = 5 = tv_a^1(r_3) = \alpha(r_3, \mathcal{S}_5)$. Finally, observe that $\mathcal{A}_6$ approximates $\mathcal{S}_6$. In this case, $tv_m^3(f_1) = \mathtt{fn}$, $\mathtt{fut}(\mathtt{fn}, 5) \in \mathcal{S}_6$ and $\sigma'(\rho_3(f_1)) = \sigma'(f_2) = 5 = \alpha(f_1, \mathcal{S}_6)$. ■

The following theorem establishes that given a concrete trace, we can generate an abstract trace of the same length and instantiate it (i.e., give the integer values to all constraints variables using a consistent assignment σ) in such a way that the size of a variable in the i-th concrete state coincides with the value of the corresponding constraint variable in the i-th abstract state.

Theorem 2 (soundness of abstract compilation)**.** *Let P be an IR program, P^α the abstract compilation of P and $\mathcal{S}_0 = \{\mathtt{ob}(\mathsf{main}, \bot, \bot, \langle tv_0, \mathit{call}(b, \mathsf{main}(this, \langle\rangle, \langle\rangle))\rangle, \emptyset)\}$ an initial state. If $\mathcal{S}_0 \rightsquigarrow^n \mathcal{S}_n$ then there exists an abstract trace $\mathcal{A}_0 \circ true \rightsquigarrow_\alpha^n \mathcal{A}_n \circ \phi_n$, such that $\mathcal{A}_0 \equiv \{\langle \mathit{call}(b, \mathsf{main}(this, \langle\rangle, \langle\rangle)), \rho \cdot \rho' \rangle\}$ and for all $\mathcal{S}_i$, it holds that $\phi_n \models \phi_i$ and $\mathcal{S}_i \approx \mathcal{A}_i \circ \phi_i$, $0 \leq i \leq n$.*

5.2. Handling Strings and Algebraic Data-Types

As already mentioned, our implementation handles *String* and *user-defined* algebraic data-types (e.g., lists, trees, etc). Below we describe how such types are handled in the size analysis. For each case (1) we describe corresponding size measures that allow abstracting data of such types to numerical values; and (2) we describe corresponding abstract compilation for instructions that manipulate such types.

DOI: 10.1002/stvr

String data-type. Strings are abstracted to their length. This is a classical size measure that allows bounding the number of iterations of loops that traverse strings. In the abstract compilation phase, instructions that manipulate strings are abstracted to linear constraints that describe relations between the lengths of the strings on which they operate. For example, the instruction $\mathsf{c} = \mathsf{strapp}(\mathsf{a}, \mathsf{b})$, which concatenates strings a and b into a new one c, is compiled into the constraint $c = a + b$ to indicate that the length of c is as the sum of the lengths of a and b. In addition, we add the constraint $a \geq 0 \wedge b \geq 0 \wedge c \geq 0$ to indicate that the length is a non-negative measure. Other string manipulating instructions are treated similarly.

Algebraic data-types. A classic size measure used for algebraic data-types, mainly in the context of termination analysis, is the *term-size* norm [16] which abstracts data-structures to the number of occurrences of type constructs in the data-structure. For example, the size of the list $\mathsf{Cons}(\mathsf{F}(\mathsf{a}, \mathsf{b}), \mathsf{Cons}(\mathsf{F}(\mathsf{a}, \mathsf{c}), \mathsf{Nil}))$ is 9, where each occurrence of a type construct from $\{\mathsf{Cons}, \mathsf{Nil}, \mathsf{F}, \mathsf{a}, \mathsf{b}, \mathsf{c}\}$ contributes 1. In the abstract compilation phase, instructions that manipulate data-structures are abstracted to linear constraints that describe relations between the term-size of the data-structures on which they operate. For example, the instruction $\mathsf{ys} = \mathsf{Cons}(\mathsf{x}, \mathsf{xs})$, which constructs a list whose head is x and whose tail is xs, is abstracted to $ys = 1 + x + xs \wedge ys \geq 0 \wedge x \geq 0 \wedge xs \geq 0$. Our implementation allows choosing between the *term-size* and the *term-depth* measure (which abstracts data-structures to their depth in a similar way). Besides, a recent extension includes state-of-the-art size measures that are automatically extracted from the user-defined types [12]. Such *type-based* norms do not count all type constructs of a given data-structures, but rather only those that are potentially traversed by loops. For example, the size of the list above would be 3 because it will only count Cons and Nil, which would lead to a more precise bound for a loop that traverses this list but not its internal elements. Moreover, our implementation allows using several type-based size measures simultaneously, which in turn allows abstracting one data-structure using several size measures. This is particularly useful when the different parts of a data-structure are traversed by different parts of the program.

5.3. Class Invariants in Cost Analysis

The accuracy of the size analysis can be improved by using a generalization of *class invariants* (see, e.g., [34]). As discussed in Section 5.1, release points are problematic since at these points other task(s) may modify the values of shared fields. However, it is often possible to gather useful information about the shared variables, in the form of class invariants, which must hold at those points. In sequential programs, class invariants have to be established by constructors and must hold on termination of all (public) methods of the class. They can be assumed at (public) method entry but may not hold temporarily at intermediate states not visible outside the object. In our context, we need that such invariants hold on method termination and also at all release points of all methods. This way, we can use them to improve the abstraction at the release points. In the following, given a class C, Ψ_C denotes the class invariant for class C, which is a set of linear constraints over the fields of C and possibly some constant symbols.

Definition 11 (abstract compilation with class invariants)**.** *Let B be an instruction or a guard and Ψ_C a class invariant for C. We define the* abstract compilation $\mathtt{ABST}^I(B, \rho, \Psi_C)$ *of B w.r.t. a mapping ρ, a symbolic size measure α_ρ and a class invariant Ψ_C as $\mathtt{ABST}^I(B, \rho, \Psi_C) = \mathtt{ABST}(B, \rho)$, if $B \neq \mathsf{await}\, x?$ and $\mathtt{ABST}^I(B, \rho, \Psi_C) = \langle \bot \wedge \Psi_C[\bar{f}_C \mapsto \rho'(\bar{f}_C)], \rho' \rangle$ otherwise, where $\mathtt{ABST}(B, \rho) = \langle \bot, \rho' \rangle$.*

The definition above allows us to define the *abstract compilation* of a possibly renamed rule $r \equiv m(this, \bar{x}, \bar{y}) \leftarrow g, b_1, \ldots, b_n \in P$ of class C w.r.t. a class invariant Ψ_C defined in terms of $\bar{f}_C$ as $r^\alpha \equiv m(\rho_0(this), \bar{I}, \bar{O}) \leftarrow \Psi_C \wedge g^\alpha, b_1^\alpha, \ldots, b_n^\alpha$, where r^α is computed as in Definition 9, using $\mathtt{ABST}^I$ instead of $\mathtt{ABST}$.

Example 12 (class invariants for the running example)**.** The following invariants will be required in order to obtain the cost of all methods of our running example: (1) In class Reader, we need

DOI: 10.1002/stvr

to know that field $elems$ is bounded, i.e., $0 \leq elems \leq elems_{max}$ where $elems_{max}$ is a constant symbol that bounds the value of $elems$; and (2) in class FileIS, we also need to know that field lth is bounded, i.e., $0 \leq lth \leq lth_{max}$. Furthermore, we need the invariant $blockS = blockS_{init}$ for the loop in method readBlock. Thus, $\Psi_{Reader} = \{0 \leq elems \leq elems_{max}\}$ and $\Psi_{FileIS} = \{0 \leq lth \leq lth_{max} \wedge blockS = blockS_{init}\}$. Now, if we consider the abstract compilation of block if_c in Example 9, and we use as invariant Ψ_{FileIS}, (b) would be replaced $\bot \wedge \rho_2(\Psi_{FileIS})$, where $\rho_2(\Psi_{FileIS}) = 0 \leq lth' \leq lth_{max} \wedge blockS' = blockS_{init}$. ■

The invariants above can be inferred automatically, for instance, by means of a syntactic analysis that simply checks that the corresponding fields are initialized and never updated again. Note that even if several processes modify lth we can still obtain the upper bounds that we have computed before. This is because there is no loop whose termination relies on the value of lth. Observe that the loop in readBlock first copies the value of lth into variable i and then the termination depends on variable i. Thus, if we have several instances of method readBlock interleaving their computations, we can still prove their termination and infer their resource consumption.

Furthermore, even if a field is modified at a release point, we can use the points-to analysis of Section 4 to determine if the field is read by means of references different from those used to write the field. If this is the case, then such a field can be preserved in rule 4 of Figure 8. In the following, given a rule $p(this, \bar{x}, \bar{y}) \leftarrow g, b_1, \ldots, b_n$, we use $body(p)$ to refer to the multiset of instructions $\{b_1, \ldots, b_n\}$. Now, we define the set $Read(C, f) = \{o \mid x := this.f \in body(p), o \in \mathcal{X}_p^{call}(this), p$ is a rule in C, $p \not\equiv C\_init\}$ of references used to read a field f in class C. Similarly, we can define the set $Write(C, f)$ but considering instructions of the form $this.f := x$ in rules of class C. Then we define the set $trackable(C)$ as the set $\{f \in \bar{f}_C \mid Read(C, f) \cap Write(C, f) = \emptyset\}$.

Example 13 (automatic inference of class invariants)**.** Let us consider the field blockS in Figure 1. Then, since $\mathcal{X}^{call}_{readBlock}(this) = \{o_2\}$ (see Example 7), and the field is only read, then $Read(\mathsf{FileIS}, \mathsf{blockS}) = \{o_2\}$ but $Write(\mathsf{FileIS}, \mathsf{blockS}) = \emptyset$. Hence $\mathsf{blockS} \in trackable(\mathsf{FileIS})$ i.e., this field is not lost when processing the instruction *await* f? inside the while loop in $readBlock$. For fields lth and fp, it holds that $Read(\mathsf{FileIS}, \mathsf{lth}) = \{o_2\}$, $Write(\mathsf{FileIS}, \mathsf{lth}) = \emptyset$, $Read(\mathsf{FileIS}, \mathsf{fp}) = \{o_1, o_2\}$ and $Write(\mathsf{FileIS}, \mathsf{fp}) = \emptyset$. Hence both fields belong to $trackable(\mathsf{FileIS})$. Thus $trackable(\mathsf{FileIS}) = \{\mathsf{blockS}, \mathsf{lth}, \mathsf{fp}\}$. ■

Differently to the sequential setting in [5], a field satisfying that $Read(C, f) = Write(C, f) = \{o\}$, i.e., the field is read and written using the same reference, cannot be considered trackable. The following example illustrates this.

Example 14 (comparison with the sequential setting)**.** Consider the following two methods:

```
void m1(A o){                    void m2(A o){
Fut<Int> x;                      Fut<Int> x;
  while (this.f > 0) {             while (this.f < 0) {
    x = o ! p();                     x = o ! p();
    await x?;                        await x?;
    this.f = this.f - 1;             this.f = this.f + 1;
  }                                }
}                                }
```

that belong to the same class C, where f is a field in C. Assume that the points-to analysis computes $\mathcal{X}^{call}_{m1}(this) = \mathcal{X}^{call}_{m2}(this) = \{o_1\}$. Then $Read(\mathsf{C}, \mathsf{f}) = Write(\mathsf{C}, \mathsf{f}) = \{o_1\}$. However, in a setting in which the execution of m1 and m2 are continuously interleaved because the corresponding *await* instructions do not hold, termination is not guaranteed, since m1 decreases f what endangers the termination of m2. Similarly, as m2 increases f, the termination of m1 cannot be guaranteed. ■

The same idea can be also applied to points-to analysis where class invariants can be used to state which fields remain unchanged at release points. In this case, the class invariant Ψ^{pt}_C is a set of field names that are guaranteed to remain unchanged after their initialization. To take this information into account, the equations in Definition 6 for *await* instructions (row 3) are changed to update only for those fields that are not in the class invariant, that is:

DOI: 10.1002/stvr

$$\mathsf{include\_global\_await}(\mathcal{X}_{\mathcal{G}}, \mathcal{X}) \equiv \mathcal{X}[o.f \mapsto \mathcal{X}_{\mathcal{G}}(o.f)], \forall o.f \in dom(\mathcal{X}) \mid o \in \mathcal{X}(this) \wedge f \notin \Psi^{pt}_{C}$$

Example 15 (class invariants for points-to analysis)**.** If $\Psi^{pt}_{\mathsf{FieIS}} = \{f\}$ was inferred for the program of Example 5, local information for the field f would be valid even after the release point at row 17 of Figure 5. Thus, using $\Psi^{pt}_{\mathsf{FieIS}}$ we have $\mathcal{X}_{\langle 1,5\rangle} = \{this{\mapsto}o_a, o_a.f{\mapsto}o_{ab}, y^{o_a}{\mapsto}o_{ab}, x^{o_a}{\mapsto}o_{ab}\}$.

6. OBJECT-SENSITIVE RESOURCE ANALYSIS

In this section, we present the process of obtaining upper bounds on the resource consumption. Our analysis follows the classical two-fold approach to cost analysis [45] in which: (1) a program is first transformed into a set of cost relations [7] which (2) can then be solved into closed-form upper/lower bounds [4], i.e., cost expressions that are not in recursive form. The cost relations we generate can be solved using [4] without requiring any modification; thus we do not describe this second phase and, in what follows, we focus exclusively on the first phase of the resource analysis.

After showing an intuitive example in Section 6.1, the presentation of the analysis is performed in two steps. First, we illustrate in Section 6.2 how an object-insensitive analysis can be defined as in sequential programming, by using the size abstraction computed in Section 5, and point out its limitations. Then, Section 6.3 defines the object-sensitive analysis which, by relying on the object-sensitive points-to information of Section 4, overcomes the limitations of the object-insensitive analysis.

6.1. An Intuitive Example

Let us consider the following simple code (left) and its IR (right):

void m(A a,int n) {	$m(this, \langle a, n\rangle, \langle\rangle) \leftarrow$
a.p(n);	$\mathit{call}(m, p(a, \langle n\rangle, \langle\rangle)),$
n++;	$n := n + 1,$
A b=*new* A();	$b := \mathit{new}\ A,$
b.p(n);	$\mathit{call}(m, p(b, \langle n\rangle, \langle\rangle).$
}	
void p(int n){	$p(this, \langle n\rangle, \langle\rangle) \leftarrow$
n++;	$n := n + 1.$
}	

We want to infer automatically the number of instructions executed by m. The most relevant point is that method p is invoked from two different objects and with two different arguments. Intuitively, the transformation of an IR program into cost relations can be formalized by transforming each rule in the program into a cost equation, which accumulates the cost of the instructions in the rule and contains the applicability conditions for the equations resulting from the abstract program. Our analysis is based on the three components introduced in the previous sections as follows. First, (1) the cost models introduced in Section 3 are used by the analysis to determine the cost of each instruction. For instance, we now use the cost model that counts the number of executed instructions and apply it to each rule. We have that the rule for m, denoted m, accumulates 4 instructions and the one for p accumulates 1. Second, (2) the size relations in Section 5 are necessary to generate the applicability conditions (guards) for cost relations and to determine how the size of data changes when the equations are applied. In particular, we infer the equations:

$$\begin{array}{lcll} m(n) & = & 4 + p(n) + p(n') & \{n' = n + 1\} \\ p(n) & = & 1 & \end{array}$$

In the equation for m, we can observe that the size relation $\{n' = n + 1\}$ tells us that the size of n is increased by one in the second call to p. In this case, the equations have no guards as they apply unconditionally.

DOI: 10.1002/stvr

Up to this point, we have obtained object-insensitive equations, because we do not distinguish the object which is executing the instructions. In the third step, (3) the points-to analysis in Section 4 is necessary to define the *cost centers*, which are the artifacts used by the analysis to separate the cost of the distributed components so that we can then distinguish the resource consumption of each component. In the example, three cost centers are obtained: $c(this)$ for the object executing m, $c(a)$ for the object that is passed as parameter to m, and $c(b)$ for the object created in m. Now, when we account for the cost of executing an instruction, we take into account the object that executes it. In particular, we obtain the following equations:

$$\begin{array}{lcll} m(n) & = & 4{*}c(this) + p_a(n) + p_b(n') & \{n' = n+1\} \\ p_a(n) & = & c(a){*}1 & \\ p_b(n) & = & c(b){*}1 & \end{array}$$

We can observe that attached to the cost we add the cost center for the object executing the instruction. The 4 instructions for method m are attributed to $this$. When a method is executed from different objects, we create object-sensitive equations which distinguish all possible calling contexts. In the example, we generate two equations for method p, one in which the cost is attributed to a and one to b. Solving the equations results in the upper bound:

$$m(n) = 4 * c(this) + 1{*}c(a) + 1{*}c(b)$$

While the cost in this example is constant, in general the cost is a function of the data input sizes, what makes the problem more interesting and challenging.

6.2. *Object-Insensitive Analysis*

The generation of cost relations from our concurrent and distributed programs, for a generic cost model $\mathcal{M}$, can be done exactly as for sequential programs [7], by using the size abstractions in Section 5 which already take the concurrent behaviour into account, and then simply applying the generic cost model in Section 3 to each instruction of each rule.

Definition 12 (object-insensitive resource analysis)**.** *Let $\mathcal{M}$ be a cost model, P an IR program, $r \equiv p(this, \bar{x}, \bar{y}) \leftarrow g, b_1, \ldots, b_n$ a rule in P and $r^\alpha \equiv p(this, \bar{x}, y') \leftarrow g^\alpha, b_1^\alpha, \ldots, b_n^\alpha$ its abstract compilation. Let $calls(r^\alpha) = \{q(\bar{w}) \mid b_i^\alpha \equiv \mathsf{call}(ct, q(rec, \bar{w}, \_)), 1 \leq i \leq n\}$ be the multiset of calls to methods or blocks in r^α. The* cost equation *associated to the rule r is defined as:*

$$p(\bar{x}) = \langle cexpr, \varphi \rangle$$

where $cexpr = \sum_{i=1}^{n} \mathcal{M}(b_i) + \sum_{q(\bar{w}) \in calls(r^\alpha)} q(\bar{w})$ *and* $\varphi \equiv \mathcal{C}(g^\alpha \wedge b_1^\alpha \wedge \ldots \wedge b_n^\alpha)$.

Given a program P, its cost relation system (CRS for short) is obtained by applying the above definition to all rules. The CRS is like a standard CRS for sequential programs with the following features: (i) equations do not have output arguments, as we aim at obtaining the cost as a function of the input argument sizes $\bar{x}$; (ii) given a rule being analyzed, its cost expression $cexpr$ is obtained by applying the cost model $\mathcal{M}$ to each of the basic instructions in the body (first summation in the $cexpr$); (iii) a call in the program is substituted by a call to its corresponding cost equation (second summation in the $cexpr$); (iv) the linear constraints φ, obtained from the size abstraction of the rule, are attached to the rule to define its applicability conditions and the size relations among the variables in the equation. When we have class invariants available, they are added to the constraints φ in the equations. Finally, the CRS is called *object-insensitive* because it does not separate the cost per object, but rather it accumulates the cost carried out by all objects in the program.

Example 16 (object-insensitive cost equations)**.** Let us see the application of Definition 12 to the rule $readBlock$ (shown in Figure 2) w.r.t. the cost model $\mathcal{M}_i$ (see Section 3), which counts the number of executed instructions. First we apply $\mathcal{M}_i$ to all instructions in the body of $readBlock$ (see Example 9), what results in the constant 6. In addition, we have to include the cost of the *while* loop, i.e., we add the call to $while$ with its corresponding arguments:

DOI: 10.1002/stvr

$$readBlock() \quad = \quad \langle \underbrace{6 + while(i', blockS)}_{cexpr}, \quad \underbrace{\{i'{=}lth, 0 \leq lth \leq lth_{max}, blockS = blockS_{init}\}}_{\varphi} \rangle$$

where φ comes from the abstract compilation in Example 9 and the class invariant Ψ_{FileIS} in Example 12. In what follows we only include the constraints of the class invariants that are relevant to compute the upper bounds. Thus, for rule *readBlock* we only include in Ψ_{FileIS} the constraint $0 \leq lth \leq lth_{max}$. Moreover, note that, according to the abstract compilation of *readBlock* in Example 9, the call to *while* should have as input arguments $\langle res', i', incr', pos', fp, lth, blockS \rangle$. For the sake of clarity, we only include those arguments that are relevant for the cost, i.e., those involved in guards (see [6] for more details). In the case of *while*, we only include i' and $blockS$. Now, by applying Definition 12 to all rules in the IR of the running example, we obtain the following CRS:

$$\begin{array}{lcll}
main() & = & \langle 14 + readOnce() + readBlock(), & \{\}\rangle \\
readBlock() & = & \langle 6 + while(i', blockS), & \{i' = lth, 0 \leq lth \leq lth_{max}\}\rangle \\
while(i, blockS) & = & \langle 0, & \{i \leq 0\}\rangle \\
while(i, blockS) & = & \langle 1 + if(i, blockS), & \{i > 0, blockS = blockS_{init}\}\rangle \\
if(i, blockS) & = & \langle 2 + if_c(i, blockS, incr), & \{blockS > i, incr = i\}\rangle \\
if(i, blockS) & = & \langle 2 + if_c(i, blockS, incr), & \{blockS \leq i, incr = blockS\}\rangle \\
if_c(i, blockS, incr) & = & \langle 7 + readContent() + while(i', blockS) & \{i' = i - incr\}\rangle \\
readOnce() & = & \langle 4 + readContent(), & \{\}\rangle \\
readContent() & = & \langle 7 + process(), & \{\}\rangle \\
process() & = & \langle 4 + while_1(i, elems), & \{i = 0, 0 \leq elems \leq elems_{max}\}\rangle \\
while_1(i, elems) & = & \langle 0, & \{i \geq elems\}\rangle \\
while_1(i, elems) & = & \langle 15 + while_1(i', elems'), & \{i < elems, i' = i + 1\}\rangle
\end{array}$$

We assume that the execution of methods hdRead and update have constant costs, which are accounted in the constant 15 of the second equation for $while_1$. Likewise, the constant 7 in the equation for *readContent* includes the cost of executing the constructor of class Reader which is assumed to be 2 (two fields are initialized). Note that the constraints of the equation *process* include the class invariant of the class Reader, i.e., $\Psi_{Reader} = \{0 \leq elems \leq elems_{max}\}$. For brevity, we do not include Ψ_{FileIS} in the equations *readContent* and *readOnce* as it is not relevant for solving the equations. The constraint $blockS = blockS_{init}$ from Ψ_{FileIS} is only relevant in the equation *while*. Observe that the constraints capture (1) the conditions required to apply the rule, as well as (2) the constraints that state how their values are modified along the program execution, e.g., in the first equation for *if* we see that, (1) when $blockS > i$, (2) the value of i is updated, $i' = i - incr$. This CRS is solved using [4] (without requiring any modification to the solving process) into the following closed-form upper bounds:

$$\begin{array}{lcl}
UB_{process}() & = & 4 + 15 * \mathsf{nat}(elems_{max}) \\
UB_{readContent}() & = & 11 + 15 * \mathsf{nat}(elems_{max}) \\
UB_{readOnce}() & = & 15 + 15 * \mathsf{nat}(elems_{max}) \\
UB_{readBlock}() & = & 6 + \mathsf{nat}(lth_{max}) * (21 + 15 * \mathsf{nat}(elems_{max}))
\end{array}$$

Before explaining the above bounds, let us first explain the role of the nat function that we use above. Suppose we are given a loop $while(x > 0)\ \{x{=}x{-}1\}$, and suppose that each iteration costs 1. Then x is an upper bound on the cost that is valid for non-negative values of x, and 0 is an upper bound that is valid for negative values of x. This bound can be expressed using a piecewise function that considers these two cases, or, for compactness, we can use $\max(0, x)$ instead, which we abbreviate as $\mathsf{nat}(x)$. Let us now explain the different parts of the upper bound computed for readBlock. The constant 6 comes from the constant in the equation for *readBlock*. The cost of the loop is the following quadratic expression:

$$\mathsf{nat}(lth_{max}) * (21 + 15 * \mathsf{nat}(elems_{max}))$$

where $\mathsf{nat}(lth_{max})$ is an upper bound on the number of iterations of the loop and $21{+}15{*}\mathsf{nat}(elems_{max})$ is the worst-case cost of each iteration. At each iteration, method

DOI: 10.1002/stvr

$readContent$ is invoked. This method contains a loop whose cost is linear on $elems_{max}$. Thus, the component $\mathsf{nat}(elems_{max})$ in the upper bound above is due to such method invocation. As the cost equation for main includes the cost of $readOnce$ and $readBlock$, its upper bound is:

$$UB_{main}() = 14 + \underbrace{15{+}15 * \mathsf{nat}(elems_{max})}_{readOnce} + \underbrace{6{+}\mathsf{nat}(lth_{max}){*}(21{+}15{*}\mathsf{nat}(elems_{max}))}_{readBlock}$$

■

The above analysis has a main drawback: it is not capable of distinguishing the different distributed components. Instead, the resource usage contributed by all objects is accumulated in a single cost center which corresponds to the whole execution of the distributed system.

6.3. Adding Cost Centers to the Equations

As its main novelty, a CRS in our object-sensitive resource analysis uses cost centers in order to keep the resource usage corresponding to the different components. The main idea is to take advantage of the object-sensitive points-to information to generate cost equations for all possible contexts (and thus objects). In particular, the object-sensitive equations will allow us to count separately the cost that corresponds to different instances of objects that are created at the same allocation site but correspond to different object names and potentially different distributed components. Given a program rule $r_t \equiv p(this, \bar{x}, \bar{y}) \ \leftarrow g, b_1, \ldots, b_n$, we annotate the rule as follows: $r_t \equiv [p(this, \bar{x}, \bar{y})]^{\mathcal{X}_p^{call}(this)} \ \leftarrow g, [b_1]^{O_1}, \ldots, [b_n]^{O_n}$, where $1 \leq j \leq n$ and O_j is defined as:

(a) $O_j = \{\langle o_l, \mathcal{X}_{\langle t,j\rangle}(rec^{o_l})\rangle \mid o_l \in \mathcal{X}_p^{call}(this)\}$ if b_j is of the form $\mathsf{call}(m, q(rec, \bar{w}, z))$;

(b) $O_j = \emptyset$, otherwise.

Observe that we are annotating the head of the rule with the set of object names that the object $this$ might take for this rule, that is, $\mathcal{X}_p^{call}(this)$. In addition we annotate the calls to methods with a set of tuples of the form $\langle o_l, \mathcal{X}_{\langle t,j\rangle}(rec^{o_l})\rangle$, where, for each element o_l in $\mathcal{X}_p^{call}(this)$, we obtain the set of possible object names that might be pointed by rec^{o_l}.

Example 17 (annotated rules)**.** The annotated rules using the information provided by the points-to analysis with $k = 2$ are the following:

$[main(this, \langle\rangle, \langle\rangle)]^{\{o_\epsilon\}} \leftarrow \ldots,$
$\quad [\mathsf{call}(m, readOnce(ob1, \langle\rangle, \langle f_1\rangle))]^{\{\langle o_\epsilon, \{o_1\}\rangle\}},$
$\quad [\mathsf{call}(m, readBlock(ob2, \langle\rangle, \langle f_2\rangle))]^{\{\langle o_\epsilon, \{o_2\}\rangle\}}, \ldots$

$[readBlock(this, \langle\rangle, \langle r\rangle)]^{\{o_2\}} \leftarrow \ldots,$
$\quad [\mathsf{call}(b, while(\overline{inp}, \overline{out}))]^{\{\langle o_2, \{o_2\}\rangle\}}, \ldots$

$[if\_c(\overline{inp}, \overline{out})]^{\{o_2\}} \leftarrow \ldots,$
$\quad [\mathsf{call}(m, readContent(this, \langle pos, incr\rangle, \langle f\rangle))]^{\{\langle o_2, \{o_2\}\rangle\}}, \ldots$

$[readOnce(this, \langle\rangle, \langle r\rangle)]^{\{o_1\}} \leftarrow \ldots,$
$\quad [\mathsf{call}(m, readContent(this, \langle 0, lth\rangle, \langle f\rangle))]^{\{\langle o_1, \{o_1\}\rangle\}}, \ldots$

$[readContent(this, \langle pos, incr\rangle, \langle f\rangle)]^{\{o_1, o_2\}} \leftarrow \ldots,$
$\quad [\mathsf{call}(m, process(rd, \langle pos\rangle, \langle f\rangle))]^{\{\langle o_1, \{o_{13}\}\rangle, \langle o_2, \{o_{23}\}\rangle\}}, \ldots$

$[process(this, \langle pos\rangle, \langle r\rangle)]^{\{o_{13}, o_{23}\}} \leftarrow \ldots$

Some blocks of method readBlock are omitted since all of them are annotated with $\{o_2\}$. Also we omit those calls not affecting further explanations. As it is shown in Figure 7, at program point $\langle 8, 0\rangle$ we have that $\mathcal{X}_{readContent}^{call}(this) = \{o_1, o_2\}$, then we annotate the rule readContent with the set $\{o_1, o_2\}$. Similarly, since at program point $\langle 9, 0\rangle$ it holds $\mathcal{X}_{process}^{call}(this) = \{o_{13}, o_{23}\}$, then the

DOI: 10.1002/stvr

corresponding annotated rule for process is annotated with $\{o_{13}, o_{23}\}$. On the contrary, at program point $\langle 8, 2\rangle$ in Figure 7, depending on the value of *this*, we have different object names which might be pointed by rd, $\mathcal{X}_{\langle 8,2\rangle}(rd^{o_1}) = \{o_{13}\}$ and $\mathcal{X}_{\langle 8,2\rangle}(rd^{o_2}) = \{o_{23}\}$. Then we annotate the call to *process* from *readContent* with $\{\langle o_1, \{o_{13}\}\rangle, \langle o_2, \{o_{23}\}\rangle\}$. ■

Now, given an annotated rule $r \equiv [p(this, \bar{x}, \bar{y})]^{This} \leftarrow g, [b_1]^{O_1}, \ldots, [b_n]^{O_n}$, we also annotate its abstract compilation as $r^\alpha \equiv [p(this, \bar{x}, \_)]^{This} \leftarrow g, [b_1^\alpha]^{O_1}, \ldots, [b_n^\alpha]^{O_n}$, and we use the following functions:

- $methods(r^\alpha)$ to obtain the multiset of annotated elements $[q(rec, \bar{w}, z)]^O$ which correspond to calls to methods of the form $[call(m, q(rec, \bar{w}, z))]^O$ in the body of r^α;
- $blocks(r^\alpha)$ to refer to the multiset of elements $q(this, \bar{w}, \bar{z})$ which are calls to intermediate rules of the form $call(b, q(this, \bar{w}, \bar{z}))$ in the body of r^α.

Definition 13 (object-sensitive resource analysis)**.** *Let $\mathcal{M}$ be a cost model, P an IR program, $r \equiv p(this, \bar{x}, \bar{y}) \leftarrow g, b_1, \ldots, b_n$ a rule in P and $r^\alpha \equiv p(this, \bar{x}, y') \leftarrow g^\alpha, b_1^\alpha, \ldots, b_n^\alpha$ its abstract compilation. Let us consider the annotated abstract rule $[p(this, \bar{x}, \bar{y})]^{This} \leftarrow \_$ for r, where $methods(r^\alpha) = \{[q_1(y_1, \bar{w}_1, z_1)]^{O_1}, \ldots, [q_k(y_k, \bar{w}_k, z_k)]^{O_k}\}$. Then the following set of equations defines the cost of r: for each $o \in This$, and for each $\langle o_1, \ldots, o_k\rangle \in O_1' \times \cdots \times O_k'$ such that $\langle o, O_i'\rangle \in O_i$, $1 \leq i \leq k$, we generate the equation:*

$$m\_o(\bar{x}) = \langle cexpr, \varphi\rangle$$

where $cexpr = \sum_{i=1}^{n} c(o) * \mathcal{M}(b_i) + \sum_{p(this,\bar{w},\bar{z}) \in blocks(r^\alpha)} p\_o(\bar{w}) + q_1\_o_1(\bar{w}_1) + \cdots + q_k\_o_k(\bar{w}_k)$,

$\varphi \equiv \mathcal{C}(g^\alpha \wedge b_1^\alpha \wedge \ldots \wedge b_n^\alpha)$, and $q_i\_o_i$ is the name of the equation that represents a call to method q_i from object o_i, and $c(o)$ denotes the cost center associated to o.

Intuitively, the above definition generates, from one rule, as many equations as needed for defining its cost such that all possible contexts (i.e., object names of callees) are considered. The new names are obtained by concatenating the corresponding object name to the rule name. Each generated equation corresponds to one object name o on which the call is performed, and such object name is used to block invocations, $p\_o(\bar{w})$. The cost expressions we accumulate, i.e. $\mathcal{M}(b_i)$, are multiplied by a symbolic expression $c(o)$ which denotes the cost center of the corresponding object name. Besides, as regards method invocations, all combinations have to be generated. This is done in the definition by means of the Cartesian product $O_1' \times \cdots \times O_k'$ which gives us all possible combinations for the elements in the sets. As an example, if we have a rule, where m_2 and m_3 are calls to methods:

$$[m_1(this, \langle x\rangle, \langle y\rangle)]^{\{o_1,o_2\}} \leftarrow [m_2(\_, \langle x, u\rangle, \langle\rangle)]^{\{\langle o_1,\{o_3\}\rangle,\langle o_2,\{o_4,o_5\}\rangle\}} + [m_3(\_, \langle u, y\rangle, \langle\rangle)]^{\{\langle o_1,\{o_6,o_7\}\rangle,\langle o_2,\{o_8\}\rangle\}}$$

Intuitively, when *this* points to o_1, then x (in m_2) may point to o_3 and u to o_6 or o_7 (in m_3). Similarly, when *this* points to o_2, then x (in m_2) may point to o_4 or o_5 and u (in m_3) may point to o_8. From the above rule, the following four equations are generated to cover all cases:

$$m_1\_o_1(x) = m_2\_o_3(x, u) + m_3\_o_6(u, y)$$
$$m_1\_o_1(x) = m_2\_o_3(x, u) + m_3\_o_7(u, y)$$
$$m_1\_o_2(x) = m_2\_o_4(x, u) + m_3\_o_8(u, y)$$
$$m_1\_o_2(x) = m_2\_o_5(x, u) + m_3\_o_8(u, y)$$

Multiple rules for the same procedure are interpreted as non-deterministic choices and the upper bound solver computes the maximum over them. Therefore, the fact that multiple non-deterministic rules are introduced (e.g., two rules for $m_1\_o_1(x)$) does not degrade the quality of the upper bound obtained.

Example 18 (object sensitive cost equations for the running example)**.** In the running example, method readContent is executed by two different objects. This is captured in the points-to analysis by means of two object names, o_1 and o_2 for the *this* reference of readContent (see Example 7).

DOI: 10.1002/stvr

Nevertheless, the upper bound for main shown in Example 16, accumulates the cost executed by all objects together. By applying Definition 13 to the annotated rules in Example 17, the equations in the CRS (Example 16) are replicated for all possible object names that could execute the equations. For example, as the rule readContent is annotated with $\{o_1, o_2\}$ and the call to process is annotated with $\{\langle o_1, \{o_{13}\}\rangle, \langle o_2, \{o_{23}\}\rangle\}$, the replicated equations obtained by the object sensitive cost analysis are as follows:

$$\begin{array}{lcll} readContent\_o_1() & = & \langle c(o_1) * 7 + process\_o_{13}(), & \{\}\rangle \\ readContent\_o_2() & = & \langle c(o_2) * 7 + process\_o_{23}(), & \{\}\rangle \\ process\_o_{13}() & = & \langle c(o_{13}) * 4 + while_1\_o_{13}(i, elems), & \{i{=}0, 0 \leq elems{\leq}elems_{max}\}\rangle \\ process\_o_{23}() & = & \langle c(o_{23}) * 4 + while_1\_o_{23}(i, elems), & \{i{=}0, 0 \leq elems{\leq}elems_{max}\}\rangle \end{array}$$

Note that the replication of the rule *process* is analogous, but for the set $\{o_{13}, o_{23}\}$. The equations include the cost center corresponding to the equation responsible of accumulating such cost, i.e., the cost accumulated by the equation $readContent\_o_1$ is multiplied by $c(o_1)$. Using such cost centers, the closed-form upper bounds now keep separate the resource consumption associated to each cost center o_i by means of a symbolic constant $c(o_i)$. From the above equations for readContent the solver obtains the following upper bounds:

$$\begin{array}{lcl} UB_{readContent\_o_1}() & = & c(o_1) * 7 + \underbrace{c(o_{13}) * (4 + 15 * \mathsf{nat}(elems_{max}))}_{process\_o_{13}} \\ UB_{readContent\_o_2}() & = & c(o_2) * 7 + \underbrace{c(o_{23}) * (4 + 15 * \mathsf{nat}(elems_{max}))}_{process\_o_{23}} \end{array}$$

Similarly for *readOnce* and *readBlock*, which are annotated with o_1 and o_2 (respectively), we obtain the upper bounds:

$$\begin{array}{lcl} UB_{readOnce\_o_1}() & = & c(o_1) * 4 + \underbrace{c(o_1) * 7 + c(o_{13}) * (4 + 15 * \mathsf{nat}(elems_{max}))}_{readContent\_o_1} \\ UB_{readBlock\_o_2}() & = & c(o_2) * 6 + c(o_2) * \mathsf{nat}(lth_{max}) * 10 + \\ & & \mathsf{nat}(lth_{max}) * \underbrace{(c(o_2) * 7 + c(o_{23}) * (4 + 15 * \mathsf{nat}(elems_{max})))}_{readContent\_o_2} \end{array}$$

With the object sensitive cost analysis, in contrast to the upper bound obtained in Example 16, the cost centers added in the cost expressions the closed-form upper bound for main keeps the number of instructions executed on each object multiplied by its corresponding cost center:

$$\begin{array}{lcl} UB_{main}() & = & c(o_\epsilon) * 14 + \underbrace{c(o_1) * 4 + c(o_1) * 7 + c(o_{13}) * (4 + 15 * \mathsf{nat}(elems_{max}))}_{readOnce\_o_1} \\ & & \underbrace{c(o_2) * 6 + c(o_2) * \mathsf{nat}(lth_{max}) * 10 +}_{readBlock\_o_2} \\ & & \underbrace{\mathsf{nat}(lth_{max}) * (c(o_2) * 7 + c(o_{23}) * (4 + 15 * \mathsf{nat}(elems_{max})))}_{readBlock\_o_2} \end{array}$$

■

The upper bound for a set of objects $\mathcal{O}$, $UB_p|_{\mathcal{O}}$, is obtained by setting $c(o)$ to 1 for all object names $o \in \mathcal{O}$ and to 0 for the remaining ones. Note that, if we replace $c(o)$ by 1 (for all object names o), the accuracy of object-insensitive CRS in Def. 12 coincides with that of object-sensitive CRS.

Example 19 (object sensitive upper bound accuracy)**.** If we are interested in the number of instructions performed by the cost centers $\{o_{13}\}$ and by $\{o_{23}\}$, we replace the symbolic expression $c(o_{13})$, $c(o_{23})$, respectively, by 1 and the rest of cost centers by 0. Then,

$$UB_{main}()|_{\{o_{13}\}} = 4 + 15 * \mathsf{nat}(elems_{max})$$
$$UB_{main}()|_{\{o_{23}\}} = \mathsf{nat}(lth_{max}) * (4 + 15 * \mathsf{nat}(elems_{max}))$$

Such upper bound captures the instructions executed by process when we call it from readOnce. The main observation is that the accuracy of the upper bound for main is significantly better when the

DOI: 10.1002/stvr

analysis is performed with $k = 2$ than with $k = 1$. If the points-to analysis is performed for $k = 1$, the object names o_{13} and o_{23} are merged in a single object name o_3, resulting in the upper bound:

$$\begin{aligned} UB^{k=1}_{main}() \quad = \quad & c(o_\epsilon) * 14 + c(o_1) * 4 + c(o_1) * 7 + \underline{c(o_3)} * (4 + 15 * \mathsf{nat}(elems_{max})) \\ & c(o_2) * 6 + c(o_2) * \mathsf{nat}(lth_{max}) * 10 + \\ & \mathsf{nat}(lth_{max}) * (c(o_2) * 7 + \underline{c(o_3)} * (4 + 15 * \mathsf{nat}(elems_{max}))) \end{aligned}$$

Therefore, it would not be possible to distinguish between the objects created at allocation site ③ and the costs are aggregated together resulting in a less precise upper bound that accumulates the expressions for $c(o_{13})$ and for $c(o_{23})$:

$$UB^{k=1}_{main}()|_{\{o_3\}} \quad = \quad (4 + 15 * \mathsf{nat}(elems_{max})) + \mathsf{nat}(lth_{max}) * (4 + 15 * \mathsf{nat}(elems_{max})) \qquad \blacksquare$$

The following theorem relates the concrete cost $\mathcal{C}(\mathcal{T}, o, \mathcal{M})$ defined in Definition 3 with the upper bound inferred by the analysis.

Theorem 3 (soundness)**.** *Let P be a program and $\mathcal{S}_0$ an initial state. If $\mathcal{T} \equiv \mathcal{S}_0 \rightsquigarrow^n \mathcal{S}_n$, then for all object identifier o such that $\mathsf{ob}(o, \_, \_, \_, \_) \in \mathcal{S}_i$, $0 \leq i \leq n$, it holds that $\mathcal{C}(\mathcal{T}, o, \mathcal{M}) \leq UB_{\mathsf{main}}()|_{\{name(o)\}}$.*

Given the soundness of the size and points-to analyses used to generate the equations, soundness of the object-sensitive cost analysis is proved by simply showing that the above CRS can be obtained by cloning the program as many times as determined by the number of object names computed by the points-to analysis and applying the standard object-insensitive cost analysis to each of the versions.

Finally, the use of cost centers easily allows us to instantiate our analysis with different deployment strategies. Such strategies determine the groups of objects that share the processor (see, e.g., JCobox [38]). The resource consumption of each group can be obtained by our approach by setting $c(o)$ to 1 for all object names o that belong to the group, and to 0 for the remaining ones.

7. EXPERIMENTAL EVALUATION

We have implemented our analysis in SACO [2], an analyzer of ABS programs which can be tried out online at: http://costa.ls.fi.upm.es/saco/web/. This section presents our experimental evaluation using the SACO system with a set of typical concurrent programs as benchmarks. The overall goal of the evaluation is to measure the accuracy and performance of our cost analysis. First, in Section 7.1 we evaluate the accuracy of the object insensitive analysis, by evaluating the obtained upper bounds against the actual cost obtained in real runs using a profiler. Then, Section 7.2 evaluates the accuracy and performance of the object sensitive analysis and the impact of using more precise points-to analysis. Finally, we evaluate and discuss the potential applications of our cost analysis using a larger and real application, namely, the TradingSystem, a case study based on the *Common Component Modelling Example* [27] (CoCoME) and developed within the FP7 HATS project http://www.hats-project.eu.

7.1. Object-Insensitive Experiments

In this section we evaluate the accuracy of the object insensitive analysis. This is done by comparing, using a set of benchmarks, the actual number of executed instructions in real runs with a random series of inputs using the aPET system [9] as profiler, against the estimated cost obtained by evaluating the generated upper bounds with the $\mathcal{M}_i$ cost model for the corresponding abstractions of the inputs. E.g., if a concrete list in algebraic form is used as input for one run, the corresponding evaluation of the upper bound is done with its term-size abstraction (see Section 5.2). The following typical concurrent applications have been used as benchmarks: BBuffer, a classical bounded-buffer for communicating several producers; DistHT, a distributed implementation of a hash table; MailServer, a simple model of a mail server; BookShop, a web shop client-server application; and, PeerToPeer, a pure peer-to-peer file sharing application. The source code of all benchmarks is available at the SACO web page.

DOI: 10.1002/stvr

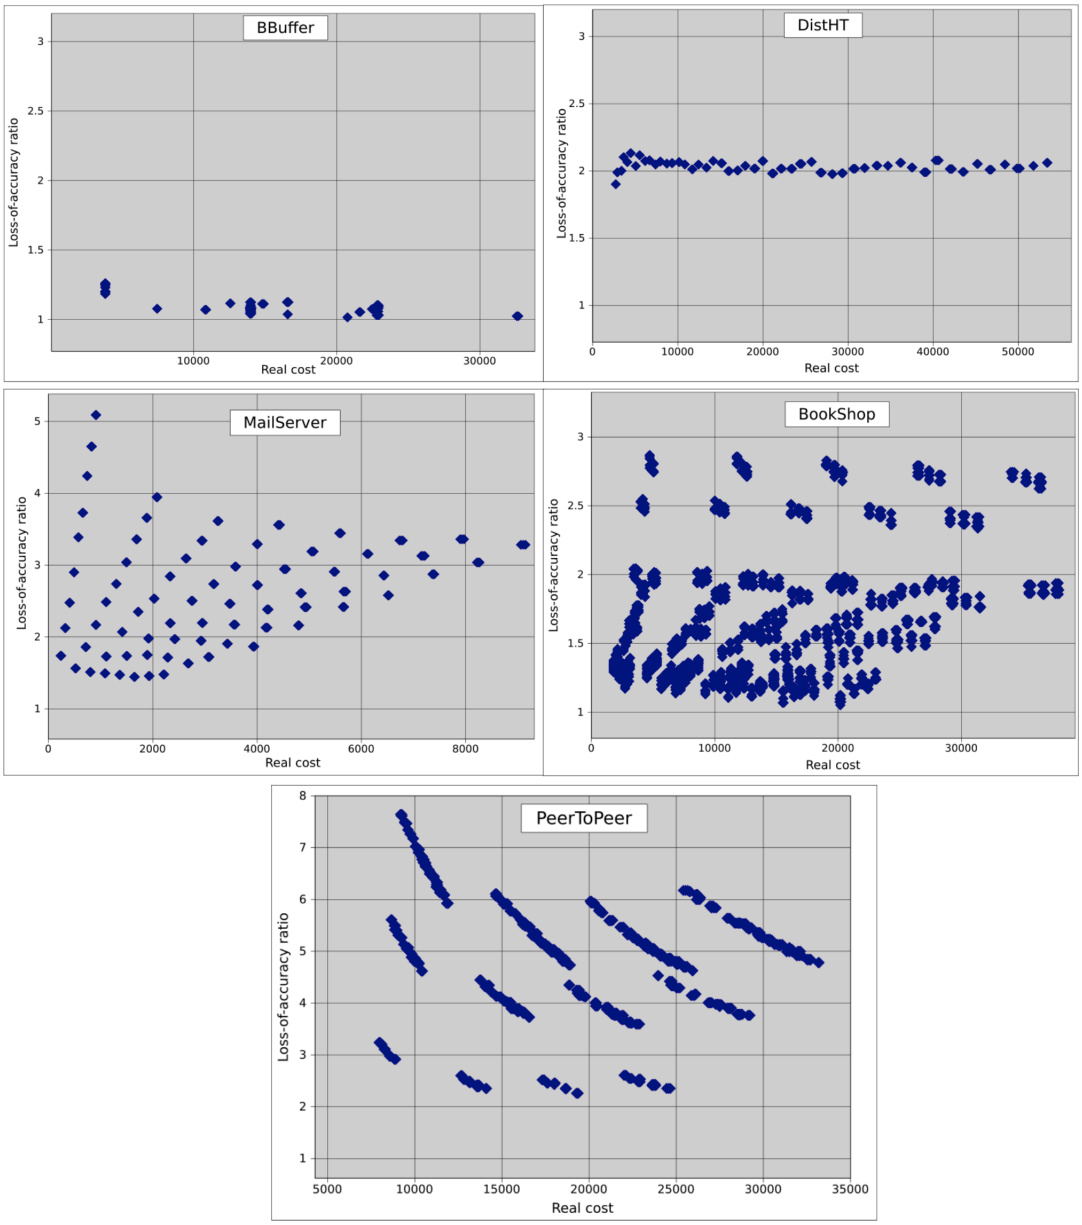

Figure 12. Evaluating precision of object-insensitive analysis

Figure 12, which aims at showing the accuracy of the analysis (y-axis) and how it varies with the cost (x-axis), shows a chart for each studied benchmark, Concretely, each point in the charts corresponds to the *loss-of-accuracy ratio* (y-component) of one run and its real cost (x-component). The loss-of-accuracy ratio for one run is obtained by dividing the estimated cost (with the corresponding abstraction of the inputs) by the real cost. Thus, the higher the ratio the less accurate the upper bound is, and the closer the ratio to one, the more precise it is.

The loss-of-accuracy ratio seems reasonable, namely from 1.1 to 7.7. For most runs it is less than 3, except for some runs of the MailServer and PeerToPeer benchmarks. It is unavoidable that the loss-of-accuracy-ratio fluctuates with different inputs. This fluctuation can be more or less severe depending on the concrete program, even for the most precise upper bounds. That is because we infer upper bounds on the worst-case cost, and thus it is clear that their quality varies when applying them to different inputs (because they must over-approximate the cost of a program for all possible

DOI: 10.1002/stvr

Benchmark	#ev	k=1			k=2						
		#c	#e	T	#c	#e	T	$\%_c$	$\%_m$	$\%_M$	$\%_A$
BBuffer	1280	15	96	602	21	143	1078	40.0	0.2	36.2	12.9
DistHT	1458	10	78	1036	17	131	1808	70.0	20.4	44.0	30.8
MailServer	1250	7	93	644	13	159	1012	71.4	3.3	42.9	19.1
Chat	1250	12	81	231	15	104	275	25.0	13.3	20.6	16.4
BookShop	3072	7	116	1485	10	162	1969	42.9	1.8	17.4	10.5
PeerToPeer	864	11	306	14648	19	581	28019	72.7	39.4	48.5	45.0

Benchmark	k=3							k=4						
	#c	#e	T	$\%_c$	$\%_m$	$\%_M$	$\%_A$	#c	#e	T	$\%_c$	$\%_m$	$\%_M$	$\%_A$
BBuffer	35	253	1620	66.7	3.9	48.8	21.8	47	301	1645	68.6	2.5	47.0	24.4
DistHT	27	229	3324	58.8	13.4	44.0	32.9	39	385	4165	44.4	7.7	39.9	29.1
MailServer	13	159	1023	0.0	0.0	0.0	0.0	13	159	1027	0.0	0.0	0.0	0.0
Chat	17	118	288	76.0	3.8	19.2	10.0	19	118	305	2.4	0.4	0.5	0.4
BookShop	11	195	2047	40.0	5.6	24.3	18.2	11	195	2042	0.0	0.0	0.0	0.0
PeerToPeer	19	581	28312	0.0	0.0	0.0	0.0	19	581	28371	0.0	0.0	0.0	0.0

Table I. Statistics about the object-sensitive resource analysis (times in ms.)

inputs). The important aspect is that those fluctuations remain stable and that they do not increase for executions of higher cost; in fact they tend to decrease in most benchmarks.

Apart from this, there are several reasons why our analysis can lose precision, namely, the size measures for algebraic data structures (see Section 5.2), field maximizations, the precision of the underlying analyses, etc. All these factors directly affect the precision of the inferred loop bounds, which are the basic ingredients used to build upper bounds. The MailServer and PeerToPeer benchmarks are especially affected by those factors mainly because they manipulate more complex, nested data-structures. Also, their costs depend on more input parameters (and fields), what increases the loss of precision of the obtained UBs.

7.2. *Object-Sensitive Experiments*

In this section we evaluate the accuracy and performance of object-sensitive cost analysis. We use the benchmarks of Section 7.1, plus an additional one, named Chat, a chat application, which could not be used in the previous section since it cannot be handled by the aPET profiler. In general, the more precise the points-to analysis is, the results of the cost analysis can be more precise as well. In this experimental evaluation, our objectives are: (1) to measure experimentally how an improvement in the precision of points-to analysis affects the precision of cost analysis; (2) to evaluate the impact of using more precise (and thus more costly) points-to analysis on the efficiency of the overall cost-analysis; and, (3) to find out which value, or values, of k achieve the best balance between precision and performance. In order to tackle such questions, we have applied object-sensitive cost analysis with four different values of k (see Section 4), $k = 1$, $k = 2$, $k = 3$ and $k = 4$.

Table I summarizes the results obtained on an Intel Core 2 Duo at 2.53GHz with 4GB of RAM, running Linux 3.2.0. Columns **#c**, **#e** show, for each value of k, the number of cost centers identified by points-to analysis and the number of equations of the CRS, respectively. Column **T** shows the time taken to apply the overall cost analysis, including the generation of the CRS and the time to solve the CRS into a closed-form upper bound. In order to measure the accuracy gained when incrementing the value of k, we evaluate the upper bound for different combinations of the input arguments and compute the average of such evaluations. Column **#ev** shows the number of different combinations evaluated for the benchmark. The result of each evaluation is a positive integer value for each cost center identified by the points-to analysis.

The evaluation of the accuracy gained by improving the precision of the points-to analysis can be done by comparing the upper bounds obtained for the cost centers identified with a value of $k = i$ with the upper bound of its corresponding object with $k = i-1$. Let us illustrate such issue by using the upper bounds obtained in Example 19. If we evaluate the upper bounds with $elems_{max} = 10$ and

DOI: 10.1002/stvr

$lth_{max} = 10$, we have that $UB^{k=1}_{main}()|_{\{o_3\}} = 1694$, $UB^{k=2}_{main}()|_{\{o_{13}\}} = 154$ and $UB^{k=2}_{main}()|_{\{o_{23}\}} = 1540$. It can be seen that the object o_3 identified with $k = 1$ has been split in objects o_{13} and o_{23} for $k = 2$, and the accuracy of the upper bounds has been improved because of that splitting. To evaluate the improvement we compute the ratio between both upper bounds and obtain the gain. Thus, to get the precision gained with o_{13}, we apply the formula $1 - (UB^{k=1}_{main}()|_{\{o_{13}\}} / UB^{k=1}_{main}()|_{\{o_3\}}) * 100$, whose evaluation returns a 90.90% of improvement. Analogously, the application to obtain the improvement of o_{23} returns 9.10%.

So as to evaluate the benchmarks we systematically compute the gain for all cost centers for different values of k, $1 \leq k \leq 4$. Thus, for each cost center identified with $k = i$, we compare its upper bound, with the upper bound of its corresponding cost center identified with the previous value of k, i.e., $k' = i - 1$. The corresponding cost center can be identified when the last i elements of the allocation sequence obtained by using k coincide with the allocation sequence of one cost center obtained using k'. In general, the gain for a given cost center c for $k = i$, for which c' is its corresponding cost center for k', is obtained with the formula $(1 - UB^k|_c \ / \ UB^{k'}|_{c'}) * 100$. As we have seen, $UB^k|_c$ is the upper bound obtained with $k = i$ for the cost center c, and $UB^{k'}|_{c'}$ is the upper bound obtained with $k' = i - 1$ for c'. Columns $\%_m$ and $\%_M$ show, respectively, the minimum and the maximum gains, and column $\%_A$ shows the average of the gains obtained for all cost centers that improve their results.

Let us start by discussing benchmarks BBuffer and DistHT. In their main methods, a structure of objects is created from which the different methods are invoked. The same method is often called from different objects and thus replication of the equations is required. In BBuffer, it can be seen that the number of equations increases from 96, with $k = 1$ up to 301 with $k = 4$. As a consequence, analysis time increases from 602ms to 1645ms. We have a similar behaviour in DistHT, it goes from 78 to 385 equations, and the time goes from 1036 ms to 4165ms. In both benchmarks, the accuracy is improved when k grows. In particular, the number of cost centers in which we improve the precision ranges between 40% and 70%, and the actual gain ranges from 13.4% to 33%. In summary, as expected, an increment on the number of cost centers found by the points-to analysis multiplies the number of equations, leading to more precise bounds and requiring larger analysis time.

As regards MailServer and PeerToPeer, the best precision is achieved with $k = 2$, i.e., incrementing the value of k does not lead to further improvements in the cost analysis. In particular, for both benchmarks, around 72% of the cost centers improve their precision when k is increased from 1 to 2. However, the gain for the MailServer is on an average 19.1%, while for PeerToPeer it is 45.0%. For both benchmarks, the time taken by the analysis increases with the number of new equations created. In MailServer, we need from 644ms for 93 equations to 1027ms for 159 equations and in the PeerToPeer, from 306 equations in 14.6s to 581 equations in 28s. For BookShop, we obtain the best precision with $k = 3$, achieving an improvement in 40% of the cost centers, and a gain of 18.2% in the upper bounds, which is higher than the improvement obtained from $k = 1$ to $k = 2$, on average 10.5% for 42.9% of the cost centers. Note that, while the efficiency of the analysis significantly degrades when going from $k = 1$ to $k = 2$, the performance is not significantly affected by the increment in the precision from $k = 2$ to $k = 3$. Regarding Chat, we can see that the analysis obtains its best precision with $k = 4$. Despite that, the improvement from $k = 3$ to $k = 4$ is not significant, as only 2.4% of the cost centers improve their precision, and the improvement is not relevant, 0.5%.

All in all, we argue that our experimental evaluation shows that object-sensitive cost analysis is feasible and accurate. As expected, the more precise the points-to analysis is, the more precise the upper-bounds obtained are. According to the experiments, the best value for k is between 2 and 3, but we point out that this is quite dependent on the benchmarks. In our benchmarks, the application of the points-to analysis with $k = 2$ is precise enough to obtain a good balance between precision and performance. Another interesting conclusion is that incrementing the value of k does not degrade the performance when there is no accuracy to be gained. This can be observed in MailServer, BookShop and PeerToPeer as, when the increment in the value of k does not produce new cost centers, the analysis is almost the same as for the previous value of k.

DOI: 10.1002/stvr

7.3. Case Study: Trading System

In this section, we aim at evaluating object sensitive cost analysis on a larger application, namely on the TradingSystem case study. The TradingSystem is an ABS implementation of the system described in [27], the *Common Component Modelling Example* (CoCoME). Our objective is to use our analysis results to identify potential bottlenecks related to the high resource consumption in some components (objects) of the distributed system and be able to give some hints on the deployment of the application (i.e., how to allocate objects to machines for the actual deployment). The TradingSystem models a supermarket cash desk line: it includes the processes at a single cash desk (e.g., scanning products using a bar code scanner, paying by cash or by credit card); it also handles bill printing, as well as other administrative tasks. A store consists of an arbitrary number of cash desks. Each of them is connected to the store server, holding store-local product data such as inventory stock, prices, etc. The system is divided into two main parts, the CashDeskInstallation and the CashDeskEnvironment. The CashDeskInstallation contains those classes that are in charge of modeling the hardware behaviour and the CashDeskEnvironment, which models at higher level the behaviour of the system. Furthermore, the TradingSystem includes a class for modeling a bank implementation and another one that models an inventory system. The program has 1350 lines of code. Some minor modifications have been done on the source code of the program in order to handle some loops whose number of iterations could not be bounded, such as loops whose number of iterations depends on one particular keystroke, or loops that terminate when the credit card is read properly. Besides, we have added some class invariants to specify that fields are unchanged at release points (these invariants could be automatically obtained by using [10]).

One interesting aspect of the TradingSystem is that the program points where the objects that compose the system are created are always placed in object initializations and our points-to analysis identifies all of them with $k = 1$. Thus, no gain is obtained for greater values of k. Points-to analysis identifies 23 cost centers and all of them correspond to a different class which models a different element that composes the system, e.g., a printer, a light display, a bar code scanners or a card reader. The application of the object-insensitive cost analysis takes 689 seconds and the time taken by the object-sensitive analysis with $k = 1$ is 875 seconds. This is an expected time due to the complexity and the size of the application. The number of equations goes from 343 for the object insensitive CRS to 413 equations in the object-sensitive approach. This slight increment is due to the fact that each object created in the program is responsible of modeling a concrete part of the system, and we only have one instance of each object and, in most cases, its methods are invoked only once. In order to see which objects can be overloaded and which objects do not have much computation (and thus can be grouped together and share the processor), we have applied the object-sensitive cost analysis and evaluated the upper bound for concrete values of the input arguments. In this case, we have obtained the percentage of instructions attributed to each cost center with respect to the total number of instructions. According to the results obtained we see three main groups of objects: (1) the objects responsible of creating the system structure, the bank and the inventory, accumulate a low number of instructions, namely 9 different objects that accumulate less than 0.7% of the instructions per object; (2) the objects that belong to the CashDeskEnvironment which accumulate, on average, around 9% of the total number of instructions per object; and, (3) those objects that correspond to the CashDeskInstallation which accumulate a quite significant part of the total number of instructions, because they include the most complex parts of the system. Such information can be useful for determining the number of processors that would lead to a better performance. As the objects that belong to (1) execute a very low number of instructions, their tasks can be executed in only one processor. For the set of objects (2), the objects that model the CashDeskEnvironment can run in a single server. Such server should have more capacity than the one used for the first type of objects. For the objects in (3) that interact with the hardware in CashDeskInstallation, they would better have their own processor to avoid contention in this part of the system.

DOI: 10.1002/stvr

8. RELATED WORK

Our work is closely related to other resource usage analysis frameworks [26, 28]. Most of such frameworks assume a sequential execution model and thus do not deal with the main challenges addressed in this paper. Notable exceptions are [32, 22]. In particular, a live heap space analysis for a concurrent language is proposed in [32]. This analysis is proposed for a simple model of shared memory and besides only considers a particular type of resource (memory) while we use a generic notion of cost. The approach in [22] is completely different to ours, and thus not directly comparable. It is based on the use of *dynamic matrices* for modeling cost analysis of concurrent programs. The use of cost centers has been proposed in the context of profiling, but to our knowledge, its use in the context of static analysis is new.

The termination of multi-threaded programs presented in [17] is based on inferring conditions on the global state which are sufficient to guarantee termination and are similar to our class invariants. Observe that such conditions are only one of the components within our cost analysis framework, which additionally requires the generation of a new form of recurrence relations and the definition of cost models for the concurrent setting. The particular case of occurrence counting analysis in mobile systems of processes, which in our proposal can be obtained using a particular cost model, has been addressed by several contributions in the literature, although they focus on high-level models, such as the π-calculus and BioAmbients [21, 25].

When considering cumulative cost models, as we do in this paper, asynchronous calls can be handled exactly as synchronous calls without sacrifying precision. This is because, in such cost models, what is important is to approximate the number of times a method is executed (i.e., called), and not how many of them might be running in parallel. In contrast, when considering noncumulative cost models, information on the lifetime of each task is important, since it might directly affect the peak consumption of the corresponding resource. As future work, we plan to integrate in our framework cost models that are noncumulative [8].

There exist other analyses for ABS programs that infer *liveness* properties (namely deadlock freeness), and they are thus complementary to ours. Recent work [24, 23] studies the problem of inferring deadlock freeness, i.e., there is no state in which a non-empty set of tasks cannot progress because all tasks are waiting for the termination of other tasks in the set, or otherwise we show the tasks involved in a potential deadlock set. In this case, the analysis tries to infer dependencies among instructions which may lead to deadlocks. As the goal of this analysis is different from ours, the basic techniques used in deadlock analysis are unrelated to those used in resource and termination analysis. However, both deadlock and resource analyses can benefit from the same underlying analysis. In both cases, *points-to analysis* and *may-happen-in-parallel* are auxiliary analysis that can greatly improve their efficiency. As we have seen, points-to analysis allows us to approximate to which objects a reference variable might be pointing. We can have an object-sensitive deadlock analysis which uses the information inferred by the points-to analysis in a similar way as we do. In the cases of may-happen-in-parallel, the deadlock analysis in [23] shows how it can greatly increase the accuracy of the analysis. For termination and resource consumption, a recent extension of our framework [10] proposes to rely on may-happen-in-parallel relations in order to automatically infer class invariants (like those defined in Section 5.3) which allow us to reason on the values of fields at processor release points.

An alternative approach to *static* cost analysis is the *measurement-based* approach [46], where the program is first executed on a set of input values in order to measure the cost of interest (e.g., execution time) for some code fragments, and then the results are combined to generate an estimation of the overall cost. This measurement can also be used in a probabilistic model to infer properties such average cost or its distribution [39]. The measurement-based approach is particularly useful when the cost of interest depends on external factors other than the program instructions. This is the case, for example, of timing analysis or energy consumption, where the time or energy required for executing an instruction depends on the current state of the underlying machine (e.g., the state of the cache). It is clearly less effective when analyzing a modeling language like ours where the underlying execution architecture is not know. In such setting we typically

DOI: 10.1002/stvr

concentrate on cost models that depend only on the program instructions and data, independently from the environment on which they will be deployed. In a recent work [11] we have explored the combination of static analysis of ABS with a simulation based approach, where the upper-bounds inferred for the (sequential) functional part of ABS where used to estimate the imperative part of the model by means of simulation, under some higher-level assumptions on the resources available in the underlying deployment component.

9. CONCLUSIONS, CURRENT AND FUTURE WORK

We have presented a novel cost analysis framework for concurrent and distributed programs based on actor-based concurrency. In summary, our main results are: (1) a sound size analysis for concurrent execution which is *field-sensitive*, i.e., it tracks data stored in the heap whenever it is sound to do so. The size analysis can be used in combination with *class invariants* which contain information on the shared memory; (2) an extension of the notion of cost used in sequential programming to the distributed setting by relying on the notion of *cost centers*, which represent the (distributed) components and allow separating their costs; (3) a flow-sensitive object-sensitive points-to analysis for the concurrent objects setting; (4) a novel form of *object-sensitive* recurrence relations which relies on information gathered by the object-sensitive points-to analysis in order to generate the cost equations; (5) a prototype implementation of a cost analyzer for programs written in the ABS language.

To develop the analysis, we have considered an object-oriented language based on the notion of concurrent objects which live in a distributed environment with asynchronous communication. The basics of our techniques could be adapted to other concurrent programming languages. In particular, the idea of having equations parametric on the cost centers is of general applicability in the context of resource analysis of distributed systems, due to the association of cost center and distributed component. The size analysis is tailored for the concurrency primitives of our language, but similar abstractions could be developed for other languages which use monitors, and an analogous abstraction would be directly applicable to other actor-based languages (e.g., Scala or Erlang).

Current work is focused on the automatic inference of class invariants which can be used to know the values of fields at processor release points. The challenge is on being able to infer the resource consumption and prove termination even in cases in which fields involved in the loop conditions are modified by several methods. For instance, consider the following two methods which belong to the same class (elems is a field):

```
Int m() {
  Int i = 0;
  while (i < elems) {
    f = o ! remoteCall();          void inc() {
    await f?;                        elems = elems + 1;
    i = i + 1;                     }
  }
  return i;
}
```

An interleaved execution of them would not allow us to prove termination. Recent work [10] proposes the use of a may-happen-in-parallel analysis to detect whether at the *await* instruction in the body of the loop, we might have an instance of method inc pending to be executed. If this is not the case, we can safely prove termination of the loop in m. Even more, even if there might be an instance of inc in the object queue when the processor is released, we would be able to prove termination, as the value of elems will be incremented once, but it will then remain stable. The general reasoning proposed in [10] is to prove that we have a finite number of instructions which update the field of interest in the queue. If this is the case, the value of the field will eventually not be modified any longer. Thus, the required invariants should establish the boundness of the field values.

DOI: 10.1002/stvr

The analysis described in this article infers bounds in terms of the explicit input parameters and fields. In some cases, however, the cost of programs might depend on implicit input, such as the length of a file. In principle, our analysis would fail to infer bounds for such parts of the program. However, if we allow users to annotate the corresponding code with some information, e.g., symbolic bounds for loops that traverse files, we will be able to handle such programs since such annotations basically turn implicit to explicit one. This is left for future work. It is important to note, however, that even if the analysis fails to infer bounds for some parts of the program, this does not necessarily mean that it will fail to infer bounds for all other parts. In fact, the final upper bound that we infer might include some symbolic values that refer to the (unknown) cost of the parts for which the analysis failed.

ACKNOWLEDGEMENTS

We are grateful to the anonymous referees for their comments and suggestions that have greatly improved the quality of the article.

References

1. Gul Agha. *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, MA, USA, 1986.
2. Elvira Albert, Puri Arenas, Antonio Flores-Montoya, Samir Genaim, Miguel Gómez-Zamalloa, Enrique Martin-Martin, Germán Puebla, and Guillermo Román-Díez. SACO: Static Analyzer for Concurrent Objects. In Erika Ábrahám and Klaus Havelund, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 20th International Conference, TACAS 2014*, volume 8413 of *Lecture Notes in Computer Science*, pages 562–567. Springer, 2014.
3. Elvira Albert, Puri Arenas, Samir Genaim, Miguel Gómez-Zamalloa, and Germán Puebla. Cost Analysis of Concurrent OO programs. In Hongseok Yang, editor, *Programming Languages and Systems - 9th Asian Symposium, APLAS 2011, Kenting, Taiwan, December 5-7, 2011. Proceedings*, volume 7078 of *Lecture Notes in Computer Science*, pages 238–254. Springer, 2011.
4. Elvira Albert, Puri Arenas, Samir Genaim, and Germán Puebla. Closed-Form Upper Bounds in Static Cost Analysis. *Journal of Automated Reasoning*, 46(2):161–203, 2011.
5. Elvira Albert, Puri Arenas, Samir Genaim, Germán Puebla, and Diana Ramírez-Deantes. From Object Fields to Local Variables: A Practical Approach to Field-Sensitive Analysis. In Radhia Cousot and Matthieu Martel, editors, *Static Analysis - 17th International Symposium, SAS 2010, Perpignan, France, September 14-16, 2010. Proceedings*, volume 6337 of *Lecture Notes in Computer Science*, pages 100–116. Springer, 2010.
6. Elvira Albert, Puri Arenas, Samir Genaim, Germán Puebla, and Damiano Zanardini. Removing Useless Variables in Cost Analysis of Java Bytecode. In Roger L. Wainwright and Hisham Haddad, editors, *Proceedings of the 2008 ACM Symposium on Applied Computing (SAC), Fortaleza, Ceara, Brazil, March 16-20, 2008*, pages 368–375. ACM, 2008.
7. Elvira Albert, Puri Arenas, Samir Genaim, Germán Puebla, and Damiano Zanardini. Cost Analysis of Object-Oriented Bytecode Programs. *Theoretical Computer Science (Special Issue on Quantitative Aspects of Programming Languages)*, 413(1):142–159, 2012.
8. Elvira Albert, Puri Arenas, Samir Genaim, and Damiano Zanardini. Task-Level Analysis for a Language with Async-Finish Parallelism. *ACM SIGPLAN Notices*, 46(5):21–30, 2011.
9. Elvira Albert, Puri Arenas, Miguel Gómez-Zamalloa, and Peter Y.H. Wong. aPET: A Test Case Generation Tool for Concurrent Objects. In Bertrand Meyer, Luciano Baresi, and Mira Mezini, editors, *Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering, ESEC/FSE'13, Saint Petersburg, Russian Federation, August 18-26, 2013*, pages 595–598. ACM, 2013.
10. Elvira Albert, Antonio Flores-Montoya, Samir Genaim, and Enrique Martin-Martin. Termination and Cost Analysis of Loops with Concurrent Interleavings. In Dang Van Hung and Mizuhito Ogawa, editors, *Automated Technology for Verification and Analysis - 11th International Symposium, ATVA 2013, Hanoi, Vietnam, October 15-18, 2013. Proceedings*, volume 8172 of *Lecture Notes in Computer Science*, pages 349–364. Springer, 2013.
11. Elvira Albert, Samir Genaim, Miguel Gómez-Zamalloa, Einar Broch Johnsen, Rudolf Schlatte, and Silvia Lizeth Tapia Tarifa. Simulating Concurrent Behaviors with Worst-Case Cost Bounds. In Michael Butler and Wolfram Schulte, editors, *FM 2011: Formal Methods - 17th International Symposium on Formal Methods, Limerick, Ireland, June 20-24, 2011. Proceedings*, volume 6664 of *Lecture Notes in Computer Science*, pages 353–368, 2011.
12. Elvira Albert, Samir Genaim, and Raúl Gutiérrez. Towards a Transformational Approach to Resource Analysis with Typed-Norms (Extended Abstract). In *23rd International Symposium on Logic-based Program Synthesis and Transformation (LOPSTR'13)*, pages 85–96, September 2013.
13. Lars Ole Andersen. *Program Analysis and Specialization for the C Programming Language*. PhD thesis, DIKU, University of Copenhagen, May 1994. (DIKU report 94/19).
14. Gregory R. Andrews. *Concurrent Programming: Principles and Practice*. Benjamin/Cummings, 1991.

DOI: 10.1002/stvr

15. Laurent Baduel, Françoise Baude, Denis Caromel, Arnaud Contes, Fabrice Huet, Matthieu Morel, and Romain Quilici. *Grid Computing: Software Environments and Tools*, chapter Programming, Deploying, Composing, for the Grid. Springer-Verlag, January 2006.
16. Annalisa Bossi, Nicoletta Cocco, and Massimo Fabris. Proving Termination of Logic Programs by Exploiting Term Properties. In Samson Abramsky and T. S. E. Maibaum, editors, *APSOFT'91: Proceedings of the International Joint Conference on Theory and Practice of Software Development, Brighton, UK, April 8-12, 1991, Volume 2: Advances in Distributed Computing (ADC) and Colloquium on Combining Paradigms for Software Developmemnt (CCPSD)*, volume 494 of *Lecture Notes in Computer Science*, pages 153–180. Springer, 1991.
17. Byron Cook, Andreas Podelski, and Andrey Rybalchenko. Proving Thread Termination. In Jeanne Ferrante and Kathryn S. McKinley, editors, *Proceedings of the ACM SIGPLAN 2007 Conference on Programming Language Design and Implementation, San Diego, California, USA, June 10-13, 2007*, pages 320–330. ACM, 2007.
18. Karl Crary and Stephanie Weirich. Resource Bound Certification. In Mark N. Wegman and Thomas W. Reps, editors, *POPL 2000, Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Boston, Massachusetts, USA, January 19-21, 2000*, pages 184–198. ACM, 2000.
19. Frank S. de Boer, Dave Clarke, and Einar Broch Johnsen. A Complete Guide to the Future. In Rocco de Nicola, editor, *Programming Languages and Systems, 16th European Symposium on Programming, ESOP 2007, Held as Part of the Joint European Conferences on Theory and Practics of Software, ETAPS 2007, Braga, Portugal, March 24 - April 1, 2007, Proceedings*, volume 4421 of *Lecture Notes in Computer Science*, pages 316–330. Springer, March 2007.
20. Manuel Fähndrich. Static Verification for Code Contracts. In Radhia Cousot and Matthieu Martel, editors, *Static Analysis - 17th International Symposium, SAS 2010, Perpignan, France, September 14-16, 2010. Proceedings*, volume 6337 of *Lecture Notes in Computer Science*, pages 2–5. Springer, 2010.
21. Jérôme Feret. Occurrence Counting Analysis for the Pi-Calculus. *Electronic Notes in Theoretical Computer Science*, 39(2):1–18, 2001.
22. Gian Luigi Ferrari and Ugo Montanari. Dynamic Matrices and the Cost Analysis of Concurrent Programs. In Vangalur S. Alagar and Maurice Nivat, editors, *Algebraic Methodology and Software Technology, 4th International Conference, AMAST '95, Montreal, Canada, July 3-7, 1995, Proceedings*, volume 936 of *Lecture Notes in Computer Science*, pages 307–321. Springer, 1995.
23. Antonio Flores-Montoya, Elvira Albert, and Samir Genaim. May-Happen-in-Parallel based Deadlock Analysis for Concurrent Objects. In Dirk Beyer and Michele Boreale, editors, *Formal Techniques for Distributed Systems (FMOODS/FORTE 2013)*, volume 7892 of *Lecture Notes in Computer Science*, pages 273–288. Springer, June 2013.
24. Elena Giachino and Cosimo Laneve. Analysis of Deadlocks in Object Groups. In Roberto Bruni and Jürgen Dingel, editors, *Formal Techniques for Distributed Systems - Joint 13th IFIP WG 6.1 International Conference, FMOODS 2011, and 31st IFIP WG 6.1 International Conference, FORTE 2011, Reykjavik, Iceland, June 6-9, 2011. Proceedings*, volume 6722 of *Lecture Notes in Computer Science*, pages 168–182. Springer, 2011.
25. Roberta Gori and Francesca Levi. A New Occurrence Counting Analysis for Bioambients. In Kwangkeun Yi, editor, *Programming Languages and Systems, Third Asian Symposium, APLAS 2005, Tsukuba, Japan, November 2-5, 2005, Proceedings*, volume 3780 of *Lecture Notes in Computer Science*, pages 381–400. Springer, 2005.
26. Sumit Gulwani, Krishna K. Mehra, and Trishul M. Chilimbi. SPEED: Precise and Efficient Static Estimation of Program Computational Complexity. In Zhong Shao and Benjamin C. Pierce, editors, *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, pages 127–139. ACM, 2009.
27. S. Herold, H. Klus, Y. Welsch, C. Deiters, A. Rausch, R. Reussner, K. Krogmann, H. Koziolek, R. Mirandola, B. Hummel, M. Meisinger, and C. Pfaller. Cocome - the common component modeling example. In *The Common Component Modeling Example*, volume 5153 of *Lecture Notes in Computer Science*, pages 16–53. Springer, 2008.
28. Jan Hoffmann and Martin Hofmann. Amortized Resource Analysis with Polynomial Potential. In Andrew D. Gordon, editor, *Programming Languages and Systems, 19th European Symposium on Programming, ESOP 2010, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2010, Paphos, Cyprus, March 20-28, 2010. Proceedings*, volume 6012 of *Lecture Notes in Computer Science*, pages 287–306. Springer, 2010.
29. E. B. Johnsen and O. Owe. An Asynchronous Communication Model for Distributed Concurrent Objects. *Software and Systems Modeling*, 6(1):35–58, 2007.
30. Einar Broch Johnsen, Reiner Hähnle, Jan Schäfer, Rudolf Schlatte, and Martin Steffen. ABS: A Core Language for Abstract Behavioral Specification. In Bernhard K. Aichernig, Frank S. de Boer, and Marcello M. Bonsangue, editors, *Formal Methods for Components and Objects - 9th International Symposium, FMCO 2010, Graz, Austria, November 29 - December 1, 2010. Revised Papers*, volume 6957 of *Lecture Notes in Computer Science*, pages 142–164. Springer, 2012.
31. Einar Broch Johnsen and Olaf Owe. An Asynchronous Communication Model for Distributed Concurrent Objects. *Software and System Modeling*, 6(1):39–58, 2007.
32. Martin Kero, Pawel Pietrzak, and Johan Nordlander. Live Heap Space Bounds for Real-Time Systems. In Kazunori Ueda, editor, *Programming Languages and Systems - 8th Asian Symposium, APLAS 2010, Shanghai, China, November 28 - December 1, 2010. Proceedings*, volume 6461 of *Lecture Notes in Computer Science*, pages 287–303. Springer, 2010.
33. Steven Lauterburg, Rajesh K. Karmani, Darko Marinov, and Gul Agha. Evaluating Ordering Heuristics for Dynamic Partial-Order Reduction Techniques. In David S. Rosenblum and Gabriele Taentzer, editors, *FASE*, volume 6013 of *Lecture Notes in Computer Science*, pages 308–322. Springer, 2010.
34. B. Meyer. *Object-Oriented Software Construction*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 2nd edition, 1997.
35. Ana Milanova, Atanas Rountev, and Barbara G. Ryder. Parameterized Object Sensitivity for Points-to Analysis for Java. *ACM Transactions on Software Engineering Methodology*, 14:1–41, 2005.

DOI: 10.1002/stvr

36. Richard G. Morgan and Stephen A. Jarvis. Profiling Large-Scale Lazy Functional Programs. *Journal of Functional Programing*, 8(3):201–237, 1998.
37. Benjamin C. Pierce. Concurrent Objects in a Process Calculus. In *Theory and Practice of Parallel Programming, International Workshop TPPP'94, Sendai, Japan, November 7-9, 1994, Proceedings*, volume 907 of *Lecture Notes in Computer Science*, pages 187–215. Springer, 1994.
38. Jan Schäfer and Arnd Poetzsch-Heffter. JCobox: Generalizing Active Objects to Concurrent Components. In Theo D'Hondt, editor, *ECOOP 2010 - Object-Oriented Programming, 24th European Conference, Maribor, Slovenia, June 21-25, 2010. Proceedings*, Lecture Notes in Computer Science, pages 275–299. Springer, 2010.
39. Sanjit A. Seshia and Alexander Rakhlin. Quantitative Analysis of Systems using Game-Theoretic Learning. *ACM Transactions on Embedded Computing Systems*, 11(S2):55:1–55:27, August 2012.
40. Yannis Smaragdakis, Martin Bravenboer, and Ondrej Lhoták. Pick your Contexts Well: Understanding Object-Sensitivity. In Thomas Ball and Mooly Sagiv, editors, *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, pages 17–30. ACM, 2011.
41. Fausto Spoto, Fred Mesnard, and Étienne Payet. A Termination Analyser for Java Bytecode based on Path-Length. *ACM Transactions on Programming Languages and Systems*, 32(3), 2010.
42. Sriram Srinivasan and Alan Mycroft. Kilim: Isolation-Typed Actors for Java. In Jan Vitek, editor, *ECOOP 2008 - Object-Oriented Programming, 22nd European Conference, Paphos, Cyprus, July 7-11, 2008, Proceedings*, volume 5142 of *Lecture Notes in Computer Science*, pages 104–128. Springer, 2008.
43. Raja Vallée-Rai, Phong Co, Etienne Gagnon, Laurie J. Hendren, Patrick Lam, and Vijay Sundaresan. Soot - a Java Optimization Framework. In Stephen A. MacKay and J. Howard Johnson, editors, *Proceedings of the 1999 conference of the Centre for Advanced Studies on Collaborative Research, November 8-11, 1999, Mississauga, Ontario, Canada*, pages 125–135. IBM, 1999.
44. Vasco Thudichum Vasconcelos. *A Process-Calculus Approach to Typed Concurrent Objects*. PhD thesis, Keio University, 1994.
45. Ben Wegbreit. Mechanical Program Analysis. *Communications ACM*, 18(9):528–539, 1975.
46. Ingomar Wenzel, Raimund Kirner, Bernhard Rieder, and Peter P. Puschner. Measurement-based timing analysis. In Tiziana Margaria and Bernhard Steffen, editors, *Leveraging Applications of Formal Methods, Verification and Validation, Third International Symposium, ISoLA 2008, Porto Sani, Greece, October 13-15, 2008. Proceedings*, volume 17 of *Communications in Computer and Information Science*, pages 430–444. Springer, 2008.

DOI: 10.1002/stvr

PROOFS (added for reviewers' convenience and to appear in an electronic appendix)

A. SOUNDNESS (PROOF SKETCH FOR THEOREM 1)

In this section, we sketch the proof of the soundness of the points-to analysis, Theorem 1. The proof sketch follows the next steps:

1. We first define a collecting semantics for the concrete setting. Such collecting semantics gathers all transitions that can be produced by the execution of a program P;
2. we continue by defining the points-to property as a property of this concrete collecting semantics; and
3. then we prove that the least solution of the set of constraints generated as described in Section 4.3 is a safe approximation of the concrete collecting semantics w.r.t. the points-to property.

Definition 14 (concrete collecting semantics)**.** *Given a program P, the concrete collecting semantics operator $\mathcal{C}_P$ is defined as follows:*

$$\mathcal{C}_P(X) = \{\langle \mathcal{S}, \mathcal{S}' \rangle \mid \langle \_, \mathcal{S} \rangle \in X \wedge \mathcal{S} \rightsquigarrow \mathcal{S}'\}$$

The concrete semantics is defined as $\xi_P = \bigcup_{n>0} \mathcal{C}_P^n(X_0)$, where $X_0 \equiv \{\langle \emptyset, \{\mathsf{ob}(\mathsf{main}, \bot, \bot, \langle tv, \mathit{call}(b, \mathsf{main}(\mathit{this}, \langle\rangle, \langle\rangle))\rangle, \emptyset)\}\rangle\}$ is the initial configuration.

Definition 15 (points-to property)**.** *Let P be an IR program, $\xi_P = \bigcup_{n>0} \mathcal{C}_P^n(X_0)$, and b an instruction at program point $\langle t, j\rangle$, then:*

(a) We say that $\xi_P \vDash_{\langle t,j\rangle} \mathit{this} = T$ if $T = \{o \mid \langle \mathcal{S}, \mathcal{S}' \rangle \in \xi_P \wedge \mathsf{ob}(o, \_, \_, \langle \_, b \cdot \bar{b}\rangle, \_) \in \mathcal{S}\}$.
(b) For a reference variable z, we say that $\xi_P \vDash_{\langle t,j\rangle} z = Z$ if $Z = \{\langle o, o' \rangle \mid \langle \mathcal{S}, \mathcal{S}' \rangle \in \xi_P \wedge \mathsf{ob}(o, \_, \_, \langle \_, b \cdot \bar{b}\rangle, \_) \in \mathcal{S} \wedge \mathsf{ob}(o, \_, \_, \langle tv, \bar{b}\rangle, \_) \in \mathcal{S}' \wedge tv(z) = o' \wedge o' \neq \mathit{null}\}$.
(c) For a reference field f, we say that $\xi_P \vDash_{\langle t,j\rangle} f = F$ if $F = \{\langle o, o' \rangle \mid \langle \mathcal{S}, \mathcal{S}' \rangle \in \xi_P \wedge \mathsf{ob}(o, \_, \_, \langle \_, b \cdot \bar{b}\rangle, \_) \in \mathcal{S} \wedge \mathsf{ob}(o, \_, h, \langle \_, \bar{b}\rangle, \_) \in \mathcal{S}' \wedge h(f) = o' \wedge o' \neq \mathit{null}\}$.

The following lemma states that the least solution of the constraint equation system defined in Section 4.3 is a safe approximation of ξ_P. We will use $\mathit{name}(o)$ to refer to the object name of an object identifier o, as described in Section 4. As in Section 4.3, we use $\mathcal{X}_{\langle t,j\rangle}$ to refer to the value of the constraint variable that corresponds to the abstract state after program point $\langle t, j\rangle$, and $\mathcal{X}_{\langle t,j-1\rangle}$ to refer to the value computed for the abstract state before program point $\langle t, j\rangle$. Similarly, the value of $\mathcal{X}_{\langle t,0\rangle}$ is the abstract state before the first instruction of rule r_t. Thus, for $j = 1$, it is guaranteed that the abstract state for $\langle t, 0\rangle$ always exists.

Lemma 1. *Let P be a program, $\langle t, j\rangle$ a program point and $\mathcal{X}_{\langle t,j\rangle}, \mathcal{X}_{\mathcal{G}}$ the least solution of the constraints equation system $\mathcal{L}_P$ as defined in Section 4.3. Then, for any reference variable z in r_t and any reference field f in the class of r_t, the following holds:*

(a) If $\xi_P \vDash_{\langle t,j\rangle} \mathit{this} = T$, then $\forall o \in T, \mathit{name}(o) \in \mathcal{X}_{\langle t,j\rangle}(\mathit{this})$.
(b) If $\xi_P \vDash_{\langle t,j\rangle} z = Z$, then $\forall \langle o, o' \rangle \in Z, \mathit{name}(o') \in \mathcal{X}_{\langle t,j\rangle}(z^{\mathit{name}(o)})$.
(c) If $\xi_P \vDash_{\langle t,j\rangle} f = F$, then $\forall \langle o, o' \rangle \in F, \mathit{name}(o') \in \mathcal{X}_{\langle t,j\rangle}(\mathit{name}(o).f)$.
(d) If $\xi_P \vDash_{\langle t,j\rangle} f = F$, then $\forall \langle o, o' \rangle \in F, \mathit{name}(o') \in \mathcal{X}_{\mathcal{G}}(\mathit{name}(o).f)$.

We say that $\mathcal{X}_{\langle t,j\rangle}$ *covers* z (resp. f) in $\mathcal{C}_P^n(X_0)$ at program point $\langle t, j\rangle$ when this lemma holds for the result of computing $\mathcal{C}_P^n(X_0)$. In order to prove this lemma, we can reason by induction on the value of n, the length of the traces $\mathcal{S}_0 \rightsquigarrow^n \mathcal{S}_n$ considered in $\mathcal{C}_P^n(X_0)$.

Base Case: If $n = 0$ then

$$\mathcal{C}_P^0(X_0) = \{\langle \epsilon, \{\mathsf{ob}(\langle id_0, \epsilon\rangle, \bot, \bot, \langle tv_0, \mathit{call}(b, \mathsf{main}(\mathit{this}, \langle\rangle, \langle\rangle))\rangle, \{\})\}\rangle\}$$

and Lemma 1 trivially holds.

Induction Hypothesis: We assume that Lemma 1 holds for all $\rightsquigarrow$-traces of length $n \geq 0$ and after n iterations of the algorithm that computes the least solution of the constraint equation system.

For any trace $\mathcal{S}_0 \rightsquigarrow^m \mathcal{S}_m$, $m \leq n$, where $\mathcal{S}_0$ is the initial state, the algorithm depicted in Section 4.3, after n iterations, generates a value for the constraint variables $\mathcal{X}_{\langle t,j\rangle}$ such that

DOI: 10.1002/stvr

(a) $\mathcal{X}_{\langle t,j\rangle}$ covers *this* in $\mathcal{C}_P^n(X_0)$ at program point $\langle t, j\rangle$.
(b) $\mathcal{X}_{\langle t,j\rangle}$ covers any program variable in rule r_t in $\mathcal{C}_P^n(X_0)$ at program point $\langle t, j\rangle$.
(c) $\mathcal{X}_{\langle t,j\rangle}$ covers any reference field in the class of r_t in $\mathcal{C}_P^n(X_0)$ at program point $\langle t, j\rangle$.
(d) $\mathcal{X}_{\mathcal{G}}$ covers any reference field in $\mathcal{C}_P^n(X_0)$.

Inductive Case: Let us consider traces of length $n+1 > 0$. Let $\langle \mathcal{S}_n, \mathcal{S}_{n+1}\rangle \in \mathcal{C}_P^{n+1}(X_0)$ be an element that corresponds to a transition $\mathcal{S}_n \rightsquigarrow \mathcal{S}_{n+1}$. We assume that transition $\mathcal{S}_n \rightsquigarrow \mathcal{S}_{n+1}$ executes the instruction at program point $\langle t, j\rangle$, which is the instruction b_j of a rule $r_t \equiv p(this, \bar{x}, \bar{y}) \leftarrow g, b_1, \ldots, b_k$, except for Rules (6), (7) and (11). Abstract states stored in constraint variables $\mathcal{X}_{\langle t,j\rangle}$ and $\mathcal{X}_{\langle t,j-1\rangle}$ have been obtained after computing n iterations of the algorithm shown in Section 4.3. To extend the Lemma to traces of length $n+1$ we reason for all possible cases in Figure 3. We assume that transition $\mathcal{S}_n \rightsquigarrow \mathcal{S}_{n+1}$ executes b_j from an object $a \in \mathcal{S}_n$ as follows.

[Rule (1)] $a \equiv \mathtt{ob}(o, C, h, \langle tv, x := e \cdot \bar{b}\rangle, \mathcal{Q})$. If x is a reference variable, e can be one of the following:

- $e \equiv y$, where y is a local variable (different from x). At state $\mathcal{S}_{n+1}$ Rule (1) uses tv from $\mathcal{S}_n$ and sets $tv[x \mapsto tv(y)]$. By the induction hypothesis, at iteration n of the algorithm $\mathcal{X}_{\langle t,j-1\rangle}$ covers any variable in $\mathcal{C}_P^n(X_0)$ before program point $\langle t, j\rangle$ and $\mathcal{S}_{n+1}$ does not change $tv(y)$. After the iteration $n+1$, the algorithm evaluates the equation

$$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathcal{X}_{\langle t,j-1\rangle}[x^l \mapsto \mathcal{X}_{\langle t,j-1\rangle}(y^l)], \forall l \in \mathcal{X}_{\langle t,j-1\rangle}(this)$$

  This equation updates $\mathcal{X}_{\langle t,j\rangle}$ so that it covers x in $\mathcal{C}_P^{n+1}(X_0)$ at $\langle t, j\rangle$.
- $e \equiv this.f$, where f is a field of class C. This case is similar to the previous case, applied to fields. At state $\mathcal{S}_{n+1}$, Rule (1) sets $tv[x \mapsto h(f)]$. By the induction hypothesis, at the iteration n, $\mathcal{X}_{\langle t,j-1\rangle}$ covers any field of the class in $\mathcal{C}_P^n(X_0)$ before program point $\langle t, j\rangle$ and $\mathcal{S}_{n+1}$ does not change $h(f)$. In the iteration $n+1$, the algorithm evaluates the equation

$$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathcal{X}_{\langle t,j-1\rangle}[x^l \mapsto \mathcal{X}_{\langle t,j-1\rangle}(l.f)], \forall l \in \mathcal{X}_{\langle t,j-1\rangle}(this)$$

  This equation updates $\mathcal{X}_{\langle t,j\rangle}$ so that it covers f in $\mathcal{C}_P^{n+1}(X_0)$ at program point $\langle t, j\rangle$.
- $e \equiv \mathit{null}$. In this case Lemma 1 is not applicable.

In all applicable cases, case (b) of Lemma 1 holds. Cases (c) and (d) also hold, as $\mathcal{X}_{\langle t,j\rangle}(\mathit{name}(o).f)$ does not change for any $f \in \bar{f}_C$. Case (a) holds, as $tv(this)$ is not modified in $\mathcal{S}_{n+1}$.

[Rule (2)] $a \equiv \mathtt{ob}(o, C, h, \langle tv, this.f := y \cdot \bar{b}\rangle, \mathcal{Q})$. At state $\mathcal{S}_{n+1}$, Rule (2) uses tv from $\mathcal{S}_n$ and sets $h[f \mapsto tv(y)]$. By the induction hypothesis, at iteration n of the algorithm, $\mathcal{X}_{\langle t,j-1\rangle}$ covers any variable in $\mathcal{C}_P^n(X_0)$ before program point $\langle t, j\rangle$ and $\mathcal{S}_{n+1}$ does not change $tv(y)$. Then, case (a) of Lemma 1 holds. After the iteration $n+1$, the algorithm evaluates the equation

$$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathcal{X}_{\langle t,j-1\rangle}[l.f \mapsto \mathcal{X}_{\langle t,j-1\rangle}(y^l)], \forall l \in \mathcal{X}_{\langle t,j-1\rangle}(this)$$

This equation updates $\mathcal{X}_{\langle t,j\rangle}$ so that it covers f for any object pointed to by *this* in $\mathcal{C}_P^{n+1}(X_0)$ at program point $\langle t, j\rangle$. In addition, by the induction hypothesis, at iteration n of the algorithm variable $\mathcal{X}_{\mathcal{G}}$ covers any reference field in $\mathcal{C}_P^n(X_0)$. The following equation updates the global state stored in $\mathcal{X}_{\mathcal{G}}$:

$$\mathcal{X}_{\mathcal{G}} \sqsupseteq \mathcal{X}_{\mathcal{G}}[o.f \mapsto \mathcal{X}_{\mathcal{G}}(o.f) \cup \mathcal{X}_{\langle t,j-1\rangle}(o.f)], \forall o \in \mathcal{X}_{\langle t,j-1\rangle}(this)$$

This equation updates $\mathcal{X}_{\mathcal{G}}$ so that it covers f in $\mathcal{C}_P^{n+1}(X_0)$. Therefore, cases (c) and (d) of Lemma 1 hold. Case (b) also holds, as variables are not changed at $\mathcal{S}_{n+1}$. Case (a) holds, as $tv(this)$ is not modified in $\mathcal{S}_{n+1}$.

[Rule (3)] $a \equiv \mathtt{ob}(o, C, h, \langle tv, x := \mathit{new}\ D \cdot \bar{b}\rangle, \mathcal{Q})$. As a result of the application of this rule, a fresh object identifier o_1 is created and stored in $tv(x)$ by means of $\mathit{newRef}(i, o)$. By the induction hypothesis, at iteration n of the algorithm $\mathcal{X}_{\langle t,j-1\rangle}$ covers *this* before program point $\langle t, j\rangle$. After the iteration $n+1$ the algorithm evaluates the equation

$$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathcal{X}_{\langle t,j-1\rangle}[x^l \mapsto \{l \oplus_k i\}], \forall l \in \mathcal{X}_{\langle t,j-1\rangle}(this)$$

This equation correctly updates $\mathcal{X}_{\langle t,j\rangle}$ so that it covers x in $\mathcal{C}_P^{n+1}(X_0)$ at program point $\langle t, j\rangle$. Therefore, case (b) of Lemma 1 holds. As fields are not changed in $\mathcal{S}_{n+1}$, cases (c) and (d) also hold. Case (a) holds, as $tv(this)$ is not modified in $\mathcal{S}_{n+1}$.

DOI: 10.1002/stvr

[Rule (4)] $a \equiv \mathsf{ob}(o, C, h, \langle tv, \mathit{call}(\mathsf{b}, q(this, \bar{z}, \bar{w})) \cdot \bar{b}\rangle, \mathcal{Q})$. This rule generates a fresh renaming of program rule p and creates a new mapping in which reference variables are initialized to *null*. Then the fresh variables $\bar{z}'$ that correspond to formal parameters are assigned the values of the actual parameters in the call, $tv(\bar{z})$.

Observe that, according to row 1 of Definition 6, each rule $r_{t'} \equiv q(this, \bar{u}, \bar{v}) \leftarrow b'_1 \dots b'_{n_{t'}}$ has an equation $\mathcal{X}_{\langle t',0\rangle} \sqsupseteq \mathcal{X}_q^{call}$, and for each call to q there is an equation $\mathcal{X}_q^{call} \sqsupseteq \mathsf{restrict\_block}(\mathcal{X}_{\langle t,j-1\rangle}, \bar{z}, \bar{u})$. Therefore,

$$\mathcal{X}_{\langle t',0\rangle} \sqsupseteq \mathcal{X}_q^{call} \sqsupseteq \mathsf{restrict\_block}(\mathcal{X}_{\langle t,j-1\rangle}, \bar{z}, \bar{u})$$

In the proof we directly use

$$\mathcal{X}_{\langle t',0\rangle} \sqsupseteq \mathsf{restrict\_block}(\mathcal{X}_{\langle t,j-1\rangle}, \bar{z}, \bar{u})$$

Analogously, upon block exit we have

$$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathsf{extend\_block}(\mathcal{X}_{\langle t,j-1\rangle}, \mathcal{X}_q^{exit}, \bar{v}, \bar{w}), \text{ and } \mathcal{X}_q^{exit} \sqsupseteq \mathcal{X}_{\langle t',n_{t'}\rangle}$$

In the proof we use

$$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathsf{extend\_block}(\mathcal{X}_{\langle t,j-1\rangle}, \mathcal{X}_{\langle t',n_{t'}\rangle}, \bar{v}, \bar{w})$$

By the induction hypothesis, at iteration n of the algorithm $\mathcal{X}_{\langle t,j-1\rangle}$ covers *this* in $\mathcal{C}_P^n(X_0)$ before program point $\langle t, j\rangle$. During iteration $n+1$ the algorithm evaluates the equations

$$\left.\begin{array}{lcl} \mathcal{X}^q_{\langle t',0\rangle} & \sqsupseteq & \mathsf{restrict\_block}(\mathcal{X}_{\langle t,j-1\rangle}, \bar{z}, \bar{u}) \\ \mathcal{X}_{\langle t,j\rangle} & \sqsupseteq & \mathsf{extend\_block}(\mathcal{X}_{\langle t,j-1\rangle}, \mathcal{X}_{\langle t',n_{t'}\rangle}, \bar{v}, \bar{w}) \end{array}\right\} \begin{array}{l} \forall t' \text{rule identifiers for } q, \\ n_{t'} \text{number of instructions in rule } t' \end{array}$$

First equation propagates fields and actual parameters in the call to formal parameters in q, according to the definition of extend_block. After iteration $n+1$, all initial constraint variables for rules defining q are updated with the points-to information regarding parameters and fields. Therefore, $\mathcal{X}_{\langle t',0\rangle}$ covers all affected variables and fields in $\mathcal{C}_P^{n+1}(X_0)$, and consequently, Lemma 1 holds.

Observe that second equation propagates with extend_block the result values and fields from the final abstract state of the rules for q back to the variables and fields of the calling rule, with the abstract state stored in $\mathcal{X}_{\langle t,j\rangle}$. This equation will be used when $b'_{n_{t'}}$ is executed, whose abstract state $\mathcal{X}_{\langle t',n_{t'}\rangle}$ is used by extend_block.

[Rule (5)] $a \equiv \mathsf{ob}(o, C, h, \langle tv, \mathit{call}(\mathsf{m}, m(rec, \bar{z}, w)) \cdot \bar{b}\rangle, \mathcal{Q})$. This rule corresponds to an asynchronous call to a method on object $o_1 = tv(rec)$ that adds a new task to object o_1 with the formal parameters $\bar{z}'$ initialized to the actual parameters $\bar{z}$ in the call, i.e.,

$$\mathsf{ob}(o_1, D, h_1, \_, \{\langle tv_3, b'_1 \cdots b'_{n_{t'}}\rangle\} \cup \mathcal{Q}')$$

where D is the class that o_1 belongs to, h_1 is the current mapping of fields to values in o_1 local heap, and tv_3 is a mapping of local variables, initialized to 0 or *null*, and parameters, that take their values from tv, namely $tv_3(\bar{x}') = tv(\bar{x})$ and $tv_3(this) = tv(rec) = o_1$.

As before, for each rule $r_{t'} \equiv m(rec, \bar{u}, v) \leftarrow b'_1 \dots b'_{n_{t'}}$, as we have

$$\mathcal{X}_{\langle t',0\rangle} \sqsupseteq \mathcal{X}_m^{call} \sqsupseteq \mathsf{include\_global}(\mathcal{X}_{\mathcal{G}}, \mathsf{restrict\_method}(\mathcal{X}_{\langle t,j-1\rangle}, \bar{z}, \bar{u}))$$

$$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathcal{X}_m^{exit} \sqsupseteq \mathsf{extend\_method}(\mathcal{X}_{\langle t,j-1\rangle}, \mathcal{X}_{\langle t',n_{t'}\rangle}, v, w),$$

we directly use

$$\mathcal{X}_{\langle t',0\rangle} \sqsupseteq \mathsf{include\_global}(\mathcal{X}_{\mathcal{G}}, \mathsf{restrict\_method}(\mathcal{X}_{\langle t,j-1\rangle}, \bar{z}, \bar{u}))$$
$$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathsf{extend\_method}(\mathcal{X}_{\langle t,j-1\rangle}, \mathcal{X}_{\langle t',n_{t'}\rangle}, v, w).$$

By the induction hypothesis, at iteration n of the algorithm $\mathcal{X}_{\langle t,j-1\rangle}$ covers all variables in $\mathcal{C}_P^n(X_0)$ before program point $\langle t, j\rangle$. During iteration $n+1$ the algorithm evaluates the equations

DOI: 10.1002/stvr

$$\left.\begin{array}{ll}(\star) & \mathcal{X}_{\langle t',0\rangle} \sqsupseteq \mathsf{include\_global}(\mathcal{X}_{\mathcal{G}}, \mathsf{restrict\_method}(\mathcal{X}_{\langle t,j-1\rangle}, rec, \bar{z}, \bar{u})) \\ (\diamond) & \mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathsf{extend\_method}(\mathcal{X}_{\langle t,j-1\rangle}, \mathcal{X}_{\langle t',n_{t'}\rangle}, rec, v, w)\end{array}\right\} \begin{array}{l}\forall t' \text{ rule id. for } m,\\ n_{t'} \text{ number of}\\ \text{instr. in rule } t'\end{array}$$

For any rule $r_{t'}$ defining method m, first equation updates $\mathcal{X}_{\langle t',0\rangle}$ to propagate actual parameters in the call to formal parameters in m, according to the definition of extend_method, and to include the global information stored in $\mathcal{X}_{\mathcal{G}}$ by means of include_global. In addition, extend_method also updates $\mathcal{X}_{\langle t',0\rangle}$ to cover the value of *this*, thus case (a) of Lemma 1 holds.

Note that $\mathcal{X}_{\langle t',0\rangle}$ is the constraint variable that keeps points-to information before the first instruction of $r_{t'}$. Therefore, after iteration $n+1$ of the algorithm $\mathcal{X}_{\langle t',0\rangle}$ covers all variables and fields of the class at the beginning of rule t', thus cases (b), (c) and (d) of Lemma 1 hold. This result will be used by Rule (11).

Second equation uses extend_method to propagate the result value from the final abstract state of the rules for m to the future variable y of the calling rule. At state $\mathcal{S}_{n+1}$ future variable y does not point to any value yet, as the method call is asynchronous. It will change its value when executing a w.*get* instruction. Therefore, after iteration $n+1$, $\mathcal{X}_{\langle t,j\rangle}$ covers all affected variables and fields in $\mathcal{C}_P^{n+1}(X_0)$.

[**Rules (6)** and **(7)**] $a \equiv \mathtt{ob}(o, C, h, \langle tv, \epsilon\rangle, \mathcal{Q})$. Both rules correspond to the end of a block or a method, respectively. By the induction hypothesis, Lemma 1 holds for $\mathcal{C}_P^n(X_0)$, and therefore for any program point $\langle t, j\rangle$, $\mathcal{X}_{\langle t,j\rangle}$ covers any variable and field in the class in $\mathcal{C}_P^n(X_0)$. As variables and fields are not changed, Lemma 1 holds at iteration $n+1$ of the algorithm. Observe that Rule (7) stores the value returned by an asynchronous execution of a method in $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}_{n+1}$. The iteration $n+1$ of the algorithm propagates this value to the abstract state of any program point $\langle t'', j''\rangle$ that corresponds to a call to this method by evaluating the following equation, which was already considered $(\diamond)$ in Rule (5):

$$\mathcal{X}_{\langle t'',j''\rangle} \sqsupseteq \mathsf{extend\_method}(\mathcal{X}_{\langle t'',j''-1\rangle}, \mathcal{X}_{\langle t,n_t\rangle}, rec, v, w) \qquad \begin{array}{l}\forall t \text{ rule identifier for } m,\\ n_t \text{ number of instructions in rule } t\end{array}$$

Observe that function extend_method updates the abstract value of w at the calling program point with the return value v at the last program point of the method called. Therefore, for any program point $\langle t'', j''\rangle$ that corresponds to a call to m, at iteration $n+1$ $\mathcal{X}_{\langle t'',j''\rangle}$ covers variable w in $\mathcal{C}_P^{n+1}(X_0)$.

[**Rule (8)**] $a \equiv \mathtt{ob}(o, C, h, \langle tv, x := y.\mathit{get} \cdot \bar{b}\rangle, \mathcal{Q})$. This rule corresponds to the execution of a *get* instruction. By the induction hypothesis, Lemma 1 holds for $\mathcal{S}_n$ and previous states. The application of this rule requires the application of Rule (7) in a previous step $\mathcal{S}_j$ for some $j \leq n$, in order to set a value for the future variable $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}_j$. As it has been already mentioned, the value of the future variable is correctly propagated to future variable y in a previous iteration of the algorithm. Therefore, $\mathcal{X}_{\langle t,j-1\rangle}$ covers variable y in $\mathcal{C}_P^n(X_0)$ before program point $\langle t, j\rangle$. Iteration $n+1$ of the algorithm evaluates equation

$$\mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathcal{X}_{\langle t,j-1\rangle}[x^l \mapsto \mathcal{X}_{\langle t,j-1\rangle}(y^l)], \forall l \in \mathcal{X}_{\langle t,j-1\rangle}(\mathit{this})$$

This equation updates $\mathcal{X}_{\langle t,j\rangle}$ so that it covers x in $\mathcal{C}_P^{n+1}(X_0)$ at program point $\langle t, j\rangle$. Therefore, case (b) of Lemma 1 holds, and cases (c) and (d) also hold, as fields are not changed. Case (a) holds, as $tv(\mathit{this})$ is not modified in $\mathcal{S}_{n+1}$.

[**Rules (9)** and **(10)**] $a \equiv \mathtt{ob}(o, C, h, \langle tv, \mathit{await}\ x? \cdot \bar{b}\rangle, \mathcal{Q})$. These rules correspond to the execution of an *await* instruction. First rule considers the case in which the future variable is ready, and the task continue its execution. In the case of rule (10), the condition of the await does not hold and the processor is released. Any other task in o may continue executing, possibly changing field values in h.

By the induction hypothesis, at iteration n of the algorithm $\mathcal{X}_{\langle t,j-1\rangle}$ covers all variables and fields and $\mathcal{X}_{\mathcal{G}}$ covers all fields in $\mathcal{C}_P^n(X_0)$ before program point $\langle t, j\rangle$. During iteration $n+1$ the algorithm evaluates the equation

$$(\blacklozenge)\quad \mathcal{X}_{\langle t,j\rangle} \sqsupseteq \mathsf{include\_global}(\mathcal{X}_{\mathcal{G}}, \mathcal{X}_{\langle t,j-1\rangle})$$

Function include_global updates $\mathcal{X}_{\langle t,j\rangle}$ with the global information stored in $\mathcal{X}_{\mathcal{G}}$. Since $\mathcal{X}_{\mathcal{G}}$ covers all fields in $\mathcal{C}_P^n(X_0)$ and the *await* instruction does not change any field, $\mathcal{X}_{\mathcal{G}}$ also covers all fields in $\mathcal{C}_P^{n+1}(X_0)$, and the equation above guarantees that $\mathcal{X}_{\langle t,j\rangle}$ also covers all fields in $\mathcal{C}_P^{n+1}(X_0)$.

DOI: 10.1002/stvr

Therefore, cases (c) and (d) of Lemma 1 hold. Case (b) also holds since variables are not changed in $\mathcal{S}_{n+1}$. Case (a) holds, as $tv(this)$ is not modified in $\mathcal{S}_{n+1}$. Observe that future variable x is updated by means of equation $(\diamond)$ at all program points $\langle t'', j'' \rangle$ which call the method awaited for by variable x.

[**Rule (11)**] $a \equiv \mathtt{ob}(o, C, h, \epsilon, \mathcal{Q})$. This rule corresponds to a context switch in an object. At state $\mathcal{S}_{n+1}$ a task from $\mathcal{Q}$ is selected and executed. The execution of the program starts with an empty $\mathcal{Q}$, and tasks are added to $\mathcal{Q}$ in two cases: when a method is invoked, using Rule (5), or when an *await* is executed and the processor is released, Rule (10). Therefore, if Rule (11) selects a task from $\mathcal{Q}$, there exists a previous step $m \leq n$ in the corresponding trace that has executed one of those rules, and by the induction hypothesis at iteration $m+1$ the corresponding constraint variable covers the variables in the calling (resp. releasing) rule and the fields in the class of the rule in $\mathcal{C}_P^{m+1}(X_0)$. The algorithm depicted in Section 4.3 obtains the least solution by means of the operation $\sqcup$, and therefore if $\mathcal{X}_{\langle t,j \rangle}$ covers a variable or field at In particular, if the task was added to $\mathcal{Q}$ by Rule (5), then there is a method rule $r_{t'}$ that was selected by (5). Equation $(\star)$ in Rule (5) guarantees that $\mathcal{X}_{\langle t',0 \rangle}$ covers all fields in the class and variables in the rule $r_{t'}$ before the first instruction of the rule. Note that equation $(\star)$ uses function include_global to include the global information stored in $\mathcal{X}_{\mathcal{G}}$ into local information of the selected task, thus local information is correctly updated regarding fields. In the case of Rule (10), the processor was released at program point $\langle t, j \rangle$ such that, at iteration m, $\mathcal{X}_{\langle t,j \rangle}$ covers all fields and variables in $\mathcal{C}_P^{m+1}(X_0)$ at that program point (Equation $(\blacklozenge)$). When the task is selected again, it is resumed at that program point, and therefore Lemma 1 also holds. As before, equation $(\blacklozenge)$ uses function include_global to include the global information stored in $\mathcal{X}_{\mathcal{G}}$ into local information of the selected task, thus flow-sensitive information regarding fields is correctly updated.

Proof of Theorem 1 is straightforward using Lemma 1.

B. SOUNDNESS (PROOF SKETCH FOR THEOREM 2)

Note that, each time we give a $\rightsquigarrow^{\alpha}$-step, of the form $\mathcal{A} \circ \phi \rightsquigarrow^{\alpha} \mathcal{A}' \circ \phi'$, it holds that either $\phi' = \phi$ or $\phi' = \phi \wedge \varphi$, where $\phi \wedge \varphi$ is satisfiable. This means that $\phi' \models \phi$ trivially. Hence it is trivial that in any α-trace of the form $\mathcal{A}_0 \circ \phi_0 \rightsquigarrow^{\alpha} \ldots \rightsquigarrow^{\alpha} \mathcal{A}_n \circ \phi_n$ it holds that $\phi_n \models \phi_i$, for all $0 \leq i \leq n$. The rest of conditions of the Theorem are proved by induction on the length n of the trace $\mathcal{S}_0 \rightsquigarrow^{n} \mathcal{S}_n$.

Base Case: If the trace is of length 0, i.e., $n = 0$, then:

$$\mathcal{S}_0 \equiv \{\mathtt{ob}(\mathsf{main}, \bot, \bot, \langle tv_0, \mathit{call}(b, \mathsf{main}(\mathit{this}, \langle\rangle, \langle\rangle))\rangle, \emptyset)\}$$

Then it is enough to consider $\mathcal{A}_0 \equiv \langle \mathit{call}(b, \mathsf{main}(\mathit{this}, \langle\rangle, \langle\rangle)), \rho \cdot \rho' \rangle$, where ρ (resp. ρ') is the identity mapping. Thus all conditions in Definition 10 trivially holds.

Inductive Case: Let us consider traces of length $n+1 > 0$. Assuming that the theorem holds for all $\rightsquigarrow$-traces of length $n \geq 0$ (the induction hypothesis), we show that it also holds for traces that consist of $n+1$ steps. Consider a $\rightsquigarrow$-trace of length n:

$$\mathcal{S}_0 \equiv \{\mathtt{ob}(\mathsf{main}, \bot, \bot, \langle tv_0, \mathit{call}(b, \mathsf{main}(\mathit{this}, \langle\rangle, \langle\rangle))\rangle, \emptyset)\} \rightsquigarrow^{n} \mathcal{S}_n$$

By the induction hypothesis, there exists an abstract trace:

$$\mathcal{A}_0 \equiv \langle \mathit{call}(b, \mathsf{main}(\mathit{this}, \langle\rangle, \langle\rangle)), \rho \cdot \rho' \rangle \circ \mathit{true} \quad \rightsquigarrow_{\alpha}^{n} \mathcal{A}_n \circ \phi_n$$

such that $S_i \approx \mathcal{A}_i \circ \phi_i$, $1 \leq i \leq n$. Let us analyze how the theorem extends to all possible $\rightsquigarrow$-traces of length $n+1$ generated from the above concrete and abstract traces. We reason for all possible cases in Figure 3, by assuming we select non-deterministically one $a \in S_n$ as follows:

[**Rule (1)**] $a \equiv \mathtt{ob}(o, C, h, \langle tv, x := e \cdot \bar{b} \rangle, \mathcal{Q})$. Then $\mathcal{S}_{n+1}$ is equals to $\mathcal{S}_n$ by replacing object a by the new one $\mathtt{ob}(o, C, h, \langle tv', \bar{b} \rangle, \mathcal{Q})$, where $v = eval_e(e, h, tv)$ and $tv' = tv[x \mapsto v]$. Now by the induction hypothesis, it holds that $a^{\alpha} \equiv \langle \varphi \cdot \bar{b}^{\alpha}, \rho_1 \cdot \rho_2 \cdot \bar{\rho} \rangle \in \mathcal{A}_n$ and the conditions of the Theorem holds. Let σ be an assignment satisfying point 2 of Definition 10, and let us analyze all possible forms of expression e:

- $e \equiv y$. Then $\varphi \equiv \rho_2(x) = \rho_1(y)$ and by the induction hypothesis ϕ_n is satisfiable, and thus, since $\rho_2(x)$ is a fresh variable, then $\phi_{n+1} \equiv \phi_n \wedge \rho_2(x) = \rho_1(y)$ is also satisfiable and we can apply the $\rightsquigarrow_{\alpha}$ rule $(3)_{\alpha}$ to $\mathcal{A}_n \circ \phi_n$ and compute the abstract state $\mathcal{A}_{n+1} \circ \phi_{n+1}$ in which $\langle \bar{b}^{\alpha}, \rho_2 \cdot \bar{\rho} \rangle \in \mathcal{A}_{n+1}$. Let us define a new assignment σ' as σ together with $\sigma'(\rho_2(x)) =$

DOI: 10.1002/stvr

$\alpha(x, \mathcal{S}_{n+1})$. Then the only point to prove is that $\sigma' \models \phi_{n+1}$, i.e., that $\sigma'(\rho_2(x)) \equiv \sigma'(\rho_1(y)) \equiv \sigma(\rho_1(y))$. But from the definition of σ' together with the induction hypothesis, this is equivalent to prove that $\alpha(x, \mathcal{S}_{n+1}) \equiv \alpha(y, \mathcal{S}_n)$. Now it is enough to distinguish the possible values of $tv(y)$:

  - If $tv(y) \in \mathbb{Z}$ then $tv'(x) \in \mathbb{Z}$. Thus $\alpha(x, \mathcal{S}_{n+1}) = tv'(x) = v = tv(y) = \alpha(y, \mathcal{S}_n)$ and the result holds.
  - If $tv(y) \in \mathit{Objects}$ then $tv'(x) \in \mathit{Objects}$. Hence $\alpha(y, \mathcal{S}_n) = 1 = \alpha(x, \mathcal{S}_{n+1})$.
  - If $tv(y) = \mathit{null}$ then $tv'(x) = \mathit{null}$. Then $\alpha(y, \mathcal{S}_n) = 1 = \alpha(x, \mathcal{S}_{n+1})$ and the result holds.
  - If $tv(y) = \mathtt{fn}$ and $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}_n$, then $tv'(y) = \mathtt{fn}$ and $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}_{n+1}$. Thus, in order to prove the result, it is enough to reason similarly as done in the items above, by considering the different possible types of v.

- $e \equiv \mathit{null}$. Then $\varphi \equiv \rho_2(x) = 0$ and since $\rho_2(x)$ is a fresh variable, then $\phi_{n+1} \equiv \phi_n \wedge \varphi$ is clearly satisfiable. Thus we can apply rule $(3)_\alpha$ and compute $\mathcal{A}_{n+1} \circ \phi_{n+1}$. Let us select σ' defined as σ together with $\sigma'(\rho_2(x)) = 0$. Clearly $\sigma' \models \phi_{n+1}$ and, since $tv'(x) = \mathit{null} = tv(y)$, then $\alpha(x, \mathcal{S}_{n+1}) = 0 = \alpha(y, \mathcal{S}_n)$. Then the result holds.
- $e \equiv \mathit{this}.f$, $eval_e(e, h, tv) = h(f) = v$, $\varphi \equiv \rho_2(x) = \rho_1(f)$ and $\phi_{n+1} \equiv \phi_n \wedge \rho_2(x) = \rho_1(f)$. Again, since by the induction hypothesis ϕ_n is satisfiable and $\rho_2(x)$ is a fresh variable, then ϕ_{n+1} is also satisfiable and rule $(3)_\alpha$ can be applied. Let us define σ' as σ extended with $\sigma'(\rho_2(x)) = \alpha(x, \mathcal{S}_{n+1})$. Then it is enough to prove that $\sigma' \models \rho_2(x) = \rho_1(f)$, i.e., $\sigma'(\rho_2(x)) = \sigma'(\rho_1(f))$. But since $\sigma'(\rho_1(f)) = \sigma(\rho_1(f))$ and by induction hypothesis $\sigma(\rho_1(f)) = \alpha(f, \mathcal{S}_n)$, then it is enough to prove that $\alpha(x, \mathcal{S}_{n+1}) = \alpha(f, \mathcal{S}_n)$. We distinguish all possible cases depending on the type of $h(f)$:
  - If $h(f) = v \in \mathbb{Z}$, then $\alpha(f, \mathcal{S}_n) = v$. Furthermore, since $tv'(x) = v \in \mathbb{Z}$ then $\alpha(x, \mathcal{S}_{n+1}) = v$ and the result holds.
  - If $h(f) = v \in \mathit{Objects}$ then $\alpha(f, \mathcal{S}_n) = 1$. Similarly to the previous cases, then $tv'(x) \in \mathit{Objects}$, i.e., $\alpha(x, \mathcal{S}_{n+1}) = 1$, getting the result.
  - If $h(f) = v = \mathit{null}$, then $\alpha(f, \mathcal{S}_n) = 0$. But since $tv'(x) = \mathit{null}$, then $\alpha(x, \mathcal{S}_{n+1}) = 0$ and the result holds.
  - If $h(f) = \mathtt{fn}$ and $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}_n$, then $tv'(y) = \mathtt{fn}$ and $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}_{n+1}$. Thus, in order to prove the result, it is enough to reason similarly as done in the items above, by considering the different possible types of v.
- $e \equiv n$ and $n \in \mathbb{Z}$. Then $\varphi \equiv \rho_2(x) = n$. We take σ' which extends σ with $\sigma'(\rho_2(x)) = n$. Thus $\sigma' \models \varphi \wedge \phi_n$, i.e., rule $(3)_\alpha$ can be applied. Now it is enough to prove that $\sigma'(\rho_2(x)) = \alpha(x, \mathcal{S}_{n+1})$. This is trivial since by definition $\alpha(x, \mathcal{S}_{n+1}) = tv'(x) = n$.
- $e \equiv y + w$. Then $tv'(x) = tv(y) + tv(w)$ and $\varphi \equiv \rho_2(x) = \rho_1(y) + \rho_1(w)$. Consider σ' extending σ with $\sigma'(\rho_2(x)) = \sigma(\rho_1(y)) + \sigma(\rho_1(w))$. Clearly $\phi_{n+1} \equiv \phi_n \wedge \varphi$ is satisfiable and rule $(3)_\alpha$ can be applied. Thus we only need to ensure that $\sigma'(\rho_2(x)) = \alpha(x, \mathcal{S}_{n+1})$. But by definition of α it holds that $\alpha(x, \mathcal{S}_{n+1}) = tv'(x) = tv(y) + tv(w)$ and $tv(y) + tv(w) = \alpha(y, \mathcal{S}_n) + \alpha(w, \mathcal{S}_n)$. Finally by the induction hypothesis we get that $\alpha(y, \mathcal{S}_n) + \alpha(w, \mathcal{S}_n) = \sigma(\rho_1(y)) + \sigma(\rho_1(w))$ and the result holds.

[Rule (2)] $a \equiv \mathtt{ob}(o, C, h, \langle tv, \mathit{this}.f := y \cdot \bar{b}\rangle, \mathcal{Q})$ and $\mathtt{ob}(o, C, h', \langle tv, \bar{b}\rangle, \mathcal{Q}) \in \mathcal{S}_{n+1}$, where $tv(y) = v$ and $h' = h[f \mapsto v]$. After applying the induction hypothesis we compute $\langle \varphi \cdot \bar{b}^\alpha, \rho_1 \cdot \rho_2 \cdot \bar{\rho}\rangle \in \mathcal{A}_n$, where $\varphi \equiv \rho_2(f) = \rho_1(y)$ and the conditions of the theorem hold. Since $\rho_2(f)$ is a fresh variable and ϕ_n is satisfiable, then $\phi_{n+1} \equiv \phi_n \wedge \rho_2(f) = \rho_1(y)$ is satisfiable and rule $(3)_\alpha$ can be applied to $\mathcal{A}_n \circ \phi_n$. Now, let σ be the valuation satisfying the conditions of the theorem and let us define a new valuation σ' which behaves similarly to σ but extended with $\sigma'(\rho_2(f)) = \alpha(f, \mathcal{S}_{n+1})$. Thus to get the result, we only need to prove that $\sigma'(\rho_2(f)) = \sigma'(\rho_1(y))$, what by definition of σ' and by the induction hypothesis, it is equivalent to prove $\alpha(f, \mathcal{S}_{n+1}) = \alpha(y, \mathcal{S}_n)$. We distinguish several cases depending on the type of v.

- If $tv(y) = v \in \mathbb{Z}$, then by definition of α it holds that $\alpha(y, \mathcal{S}_n) = v$. Now $h'(f) = v$ and thus $\alpha(f, \mathcal{S}_{n+1}) = v$ and the result holds.
- If $tv(y) = v = \mathit{null}$ then $\alpha(y, \mathcal{S}_n) = 0$. Since $h'(f) = \mathit{null}$, then $\alpha(f, \mathcal{S}_{n+1}) = 0$ and the result holds.
- If $tv(y) = v \in \mathit{Objects}$ then $\alpha(y, \mathcal{S}_n) = 1$. Since $h'(f) = v \in \mathit{Objects}$, then $\alpha(f, \mathcal{S}_{n+1}) = 1$ and the result holds.
- Finally, if $tv(y) = \mathtt{fn}$ and $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}_n$, then $h'(f) = \mathtt{fn}$ and $\mathtt{fut}(\mathtt{fn}, v) \in \mathcal{S}_{n+1}$. Thus, in order to prove the result, it is enough to reason similarly as done in the items above, by considering the different possible types of v.

DOI: 10.1002/stvr

[Rule (3)] $a \equiv \texttt{ob}(o, C, h, \langle tv, x := \textit{new}\ D \cdot \bar{b}\rangle, \mathcal{Q})$ and $\texttt{ob}(o, C.h, \langle tv', \bar{b}\rangle, \mathcal{Q})$ $\texttt{ob}(o_1, D, h_1, \epsilon, \emptyset)$ belongs to $\mathcal{S}_{n+1}$, where $tv' = tv[x \mapsto o_1]$. By the induction hypothesis $\langle \rho_2(x) = 1 \cdot \bar{b}^{\alpha}, \rho_1 \cdot \rho_2 \cdot \bar{\rho}\rangle \in \mathcal{A}_n$, and all conditions in Definition 10 hold. Note that since $\rho_2(x)$ is a fresh variable and ϕ_n is satisfiable (by the induction hypothesis), then $\phi_{n+1} \equiv \phi_n \wedge \rho_2(x) = 1$ is also satisfiable. Let σ be an assignment ensuring $\mathcal{S}_n \approx \mathcal{A}_n \circ \phi_n$. Let us consider σ' defined as σ together with $\sigma'(\rho_2(x)) = 1$. Then $\sigma' \models \phi_{n+1}$, and we only need to prove that $\sigma'(\rho_2(x)) = 1 = \alpha(x, \mathcal{S}_{n+1})$. But this follows from the definition of α, since $tv'(x) \in \mathit{Objects}$. Note that since $\mathit{tasks}(\texttt{ob}(o_1, D, h_1, \epsilon, \emptyset)) = \emptyset$, this object has not to be considered.

[Rule (4)] In this case it holds that $a \equiv \texttt{ob}(o, C, h, \langle tv, \textit{call}(\mathsf{b}, p(\mathit{this}, \bar{x}, \bar{y})) \cdot \bar{b}\rangle, \mathcal{Q}) \in \mathcal{S}_n$, $\texttt{ob}(o, C, h, \langle tv \cup tv_2, b'_1 \cdots b'_n \cdot \bar{b}\rangle, \mathcal{Q}) \in \mathcal{S}_{n+1}$ and $r \equiv p(\mathit{this}', \bar{x}', \bar{y}) \leftarrow g',\ \ b'_1, \ldots, b'_n \ll_p^{\langle \bar{y}\rangle} P$, $tv_1 = \mathit{newEnv}(\mathit{vars}(r) - \{\bar{y}\})$, $tv_2 = tv_1[\mathit{this}' \mapsto o, \bar{x}' \mapsto tv(\bar{x})]$, $\mathit{eval}_{gd}(g', tv_2) = \mathit{true}$. By the induction hypothesis $\langle \textit{call}(\mathsf{b}, p(\mathit{this}_1, \bar{x}_1, \bar{y}_1)) \cdot \bar{b}^{\alpha}, \rho_0 \cdot \rho_1 \cdot \bar{\rho}\rangle \in \mathcal{A}_n$ and the conditions of the theorem hold. Concretely,

$$(*)_1\ \rho_0(\mathit{this}) = \mathit{this}_1, \rho_0(\bar{x}) = \bar{x}_1, \rho_1(\bar{y}) = \bar{y}_1$$

Let us do the abstract compilation of r by starting from the identity renaming ρ'_0, what results in r^{α}, where $r^{\alpha} \equiv p(\mathit{this}', \bar{x}', \bar{y}) \rightarrow g'^{\alpha}, b'^{\alpha}_1, \ldots, b'^{\alpha}_n \circ \rho'_0 \cdots \rho'_{n+1}$. Let σ be the assignment satisfying the conditions of the theorem, and concretely that $\sigma \models \phi_n$. We build a new assignment σ' defined as σ but extended with:

$$(*)_2\ \sigma'(\bar{x}') = \sigma(\bar{x}_1), \sigma'(\mathit{this}') = \sigma(\mathit{this}_1)$$

Then $\sigma' \models \phi_n \wedge \bar{x}' = \bar{x}_1 \wedge \mathit{this}' = \mathit{this}$. Let us prove also that $\sigma' \models g'^{\alpha}$. To this end, let us suppose that $g'^{\alpha} \equiv z' \diamond w'$, where $z', w' \subseteq \bar{x}'$. Then:

$\mathit{eval}_{gd}(g', tv_2) = \mathit{true}$	$\Rightarrow$	(By definition)
$tv_2(z') \diamond tv_2(w')$ is true	$\Rightarrow$	(By definition)
$tv(z) \diamond tv(w)$ is true	$\Rightarrow$	(By definition)
$\alpha(z, \mathcal{S}_n) \diamond \alpha(w, \mathcal{S}_n)$ is true	$\Rightarrow$	(By the induction hypothesis)
$\sigma(\rho_0(z)) \diamond \sigma(\rho_0(w))$ is true	$\Rightarrow$	(By $(*)_1$)
$\sigma(z_1) \diamond \sigma(w_1)$ is true	$\Rightarrow$	(By $(*)_2$)
$\sigma'(z') \diamond \sigma(w')$ is true		

This means that $\phi_{n+1} \equiv \phi_n \wedge \bar{x}' = \bar{x}_1 \wedge \mathit{this}' = \mathit{this} \wedge g'^{\alpha}$ is satisfiable and rule $(1)_{\alpha}$ can be applied on $\mathcal{A}_n \circ \phi_n$ to compute $\mathcal{A}_{n+1} \circ \phi_{n+1}$, where $\langle b'^{\alpha}_1 \cdot \cdots \cdot b'^{\alpha}_n \cdot \bar{b}^{\alpha} \circ \rho'_1 \cdots \rho'_{n+1} \cdot \rho_1 \cdot \bar{\rho}\rangle \in \mathcal{A}_{n+1}$. All conditions required on $\mathcal{S}_{n+1}$ and $\mathcal{A}_{n+1} \circ \phi_{n+1}$ to prove the equivalence trivially hold except for the following both statements:

$$\alpha(\mathit{this}', \mathcal{S}_{n+1}) = \sigma'(\rho'_1(\mathit{this}'))$$
$$\alpha(\bar{x}', \mathcal{S}_{n+1}) = \sigma'(\rho'_1(\bar{x}'))$$

However observe that $\alpha(\mathit{this}', \mathcal{S}_{n+1})$ is defined using $tv_2(\mathit{this}')$ which is equals to $tv(\mathit{this})$. Hence $\alpha(\mathit{this}', \mathcal{S}_{n+1}) = \alpha(\mathit{this}, \mathcal{S}_n)$. On the other hand, note that by the induction hypothesis:

$\alpha(\mathit{this}, \mathcal{S}_n)$	$=$	(By the induction hypothesis)
$\sigma(\rho_0(\mathit{this}))$	$=$	(By $(*)_1$)
$\sigma(\mathit{this}_1)$	$=$	(By $(*)_2$)
$\sigma'(\mathit{this}')$	$=$	(ρ'_1 is the identity renaming)
$\sigma'(\rho'_1(\mathit{this}'))$		

i.e., $\alpha(\mathit{this}', \mathcal{S}_{n+1}) = \sigma'(\rho'_1(\mathit{this}'))$. Similarly we can prove also $\alpha(\bar{x}', \mathcal{S}_{n+1}) = \sigma'(\rho'_1(\bar{x}'))$ to get the result.

[Rule (5)] It can be reasoned similarly to rule (4) but applying rule $(2)_{\alpha}$. Note that at this point the new future variable introduced in $\mathcal{S}_{n+1}$ is undefined and thus nothing new w.r.t. the above case must be proved.

DOI: 10.1002/stvr

[**Rule (6)**] The result trivially holds by the induction hypothesis and the application of rule $(6)_\alpha$.

[**Rule (7)**] It holds that $\texttt{ob}(o, C, h, \langle tv, \epsilon\rangle, \mathcal{Q}) \in \mathcal{S}_n$, $\texttt{fut}(\texttt{fn}, \bot) \in \mathcal{S}_n$, $\texttt{ret} \in dom(tv)$, $(y,\ \texttt{fn}) = tv(\texttt{ret}), v = tv(y)$ and $\texttt{ob}(o, C, h, \epsilon, \mathcal{Q}) \in \mathcal{S}_{n+1}$ and $\texttt{fut}(\texttt{fn}, v) \in \mathcal{S}_{n+1}$. By the induction hypothesis $\langle \epsilon, \_\rangle \in \mathcal{A}_n$. Thus rule $(6)_\alpha$ can be applied on such an abstract configuration.

Let us assume that there exists an object[†] o_1 such that $\texttt{ob}(o_1, C_1, h_1, \langle tv_1, \bar{b}_1\rangle, \mathcal{Q}_1) \in \mathcal{S}_n$ and $tv_1(y_1) = \texttt{fn}$. Thus $\texttt{ob}(o_1, C_1, h_1, \langle tv_1, \bar{b}_1\rangle, \mathcal{Q}_1) \in \mathcal{S}_{n+1}$. By the induction hypothesis $\langle \bar{b}_1^\alpha, \rho \cdot \bar{\rho}\rangle \in \mathcal{A}_n$ and also to $\mathcal{A}_{n+1}$. Suppose that σ is the valuation satisfying the Theorem. Then what it is required to prove is that:

$$\sigma(\rho(y_1)) = \alpha(y_1, \mathcal{S}_{n+1})$$

Suppose that $v \in \mathbb{Z}$ (the rest of cases can be reasoned similarly). Then, by definition of α it holds that $\alpha(y_1, \mathcal{S}_{n+1}) = tv_1(y_1) = v$, i.e., it is enough to prove:

$$\sigma(\rho(y_1)) = v$$

Because of the form of o_1 in $\mathcal{S}_n$, it is for sure that there exists a configuration $\mathcal{S}_k$, with $k < n$, such that $\texttt{ob}(o_1, C_1, \_, \langle tv_2, \mathit{call}(m, p(rec_1, \bar{x}_1, y_1) \cdot \_\rangle, \_) \in \mathcal{S}_k$, $tv_2(rec_1) = o$ and, after applying rule (5) in Figure 3 with $r \equiv p(this', \bar{x}', y) \leftarrow \_ \ll_p^{\langle y\rangle} P$, it holds that $\texttt{ob}(o_1, C_1, \_, \langle tv_3, \_\rangle, \_) \in \mathcal{S}_{k+1}$, where $tv_3(y_1) = \texttt{fn}$. By the induction hypothesis, then $\langle \mathit{call}(m(p, rec_2, \bar{x}_2, y_2)) \cdot \_, \rho_1' \cdot \rho_2' \cdot \bar{\rho}'\rangle \in \mathcal{A}_k$, where $\rho_2'(y_1) = y_2$, and after applying rule $(2)_\alpha$ with the abstract compilation of r, let us say $r^\alpha \equiv p(this', \bar{x}', y') \leftarrow \_ \circ \rho_1'' \cdots \bar{\rho}_{n+1}'' \ll_p^\alpha P^\alpha$, with ρ_1'' being the identity mapping and $\rho_{n+1}''(y) = y'$, we get that $y' = y_2 \in \phi_{k+1}$, and thus $y' = y_2 \in \phi_j$, $k+1 \leq j \leq n+1$.

Now since $\sigma \models \phi_n$, then $\sigma(y_2) = \sigma(y')$, where $\sigma(y') = \sigma(\rho_{n+1}''(y))$. By the induction hypothesis, $\sigma(\rho_{n+1}''(y)) = \alpha(y, \mathcal{S}_n)$, where $\alpha(y, \mathcal{S}_n) = v$. Since $y_2 = \rho_2'(y_1)$, then we have proved that $\sigma(\rho_2'(y_1)) = v$. The only point to observe now is that since y_1 is an output variable, $\rho(y_1)$ will be for sure equals to $\rho_2'(y_1)$ and hence we have proved that $\sigma(\rho(y_1)) = v$.

[**Rule (8)**] $\texttt{ob}(o, C, h, \langle tv, x := y.\mathit{get} \cdot \bar{b}\rangle, \mathcal{Q}) \in \mathcal{S}_n$, $\texttt{fut}(\texttt{fn}, v) \in \mathcal{S}_n$, where $\texttt{fn} = tv(y)$, $v \neq \bot$. By the induction hypothesis, $\langle \rho_2(x) = \rho_1(y) \cdot b^\alpha, \rho_1 \cdot \rho_2 \cdot \bar{\rho}\rangle \in \mathcal{A}_n$, where ϕ_n is satisfiable, and there exists σ such that $\sigma \models \phi_n$ and $\sigma(\rho_1(y)) = v$. Now we apply rule (8) on $\mathcal{S}_n$, what transforms $\mathcal{S}_n$ as $\texttt{ob}(o, C, h, \langle tv[x \mapsto v], \bar{b}\rangle, \mathcal{Q}), \texttt{fut}(\texttt{fn}, v) \in \mathcal{S}_{n+1}$. Since ϕ_n is satisfiable and $\rho_2(x)$ is a fresh variable, then $\phi_{n+1} \equiv \phi_n \wedge \rho_2(x) = \rho_1(y)$ is satisfiable and we can apply rule $(3)_\alpha$ to compute $\langle b^\alpha, \rho_2 \cdot \bar{\rho}\rangle \in \mathcal{A}_{n+1}$. Suppose that $v \in \mathbb{Z}$ (the rest of cases can be reasoned similarly). Let us choose an assignment σ' defined as σ but extended with $\sigma'(\rho_2(x)) = v$. This assignment satisfies that $\sigma' \models \phi_{n+1}$. Furthermore $\sigma'(\rho_2(x)) = tv[x \mapsto v](x) = v = \alpha(x, \mathcal{S}_{n+1})$. Thus the result holds.

[**Rules (9) and (10)**] In both cases the proof follows immediately from the induction hypothesis by applying respectively, rules $(4)_\alpha$ and $(5)_\alpha$. Note that for rule (10) we have that $tasks(\mathcal{S}_n) = tasks(\mathcal{S}_{n+1})$ since the task is only introduced in the queue of the current object.

[**Rule (11)**] Since after the application of rule (11) it holds that $tasks(\mathcal{S}_n) = tasks(\mathcal{S}_{n+1})$, then the result trivially follows from the induction hypothesis and the application of rule $(7)_\alpha$.

C. SOUNDNESS (PROOF SKETCH FOR THEOREM 3)

We sketch the main ideas of the proof for the object-insensitive analysis, and then we comment on the straightforward changes required to handle the object-sensitive case. The proof sketch consists of two parts:

- In the first one, we instrument the abstract states in the abstract operational semantics of Figure 9 with a cost component that measures the cost of abstract executions; and show that each concrete trace has a corresponding abstract one with the same cost.
- Then, in a second part, we show that the cost relations generated from the abstract program indeed approximate the resource consumption behaviour of the abstract program.

DOI: 10.1002/stvr

$$(1)_\alpha \quad \frac{p(this', \bar{x}', \bar{y}') \leftarrow g^\alpha, {b'}^\alpha_1, \ldots, {b'}^\alpha_n \circ \rho_0 \cdots \rho_{n+1} \ll^\alpha_p P^\alpha, g^\alpha \wedge \phi \not\models false, e' = \mathcal{M}(p(this, \bar{x}, \bar{y})) + e}{\begin{array}{c}\{\langle \mathit{call}(b, p(this, \bar{x}, \bar{y})) \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho}\rangle | \mathcal{A}\} \circ \phi \circ e \rightsquigarrow_\alpha \{\langle {b'}^\alpha_1 \cdots {b'}^\alpha_n \cdot \bar{b}^\alpha, \rho_1 \cdots \rho_{n+1} \cdot \bar{\rho}\rangle | \mathcal{A}\} \circ \\ \phi \wedge g^\alpha \wedge this = this' \wedge \bar{x}' = \bar{x} \circ e'\end{array}}$$

$$(2)_\alpha \quad \frac{p(rec', \bar{x}', y') \leftarrow {b'}^\alpha_1, \ldots, {b'}^\alpha_n \circ \rho_1 \cdots \rho_{n+1} \ll^\alpha_p P^\alpha, e' = \mathcal{M}(p(rec, \bar{x}, \bar{y})) + e}{\begin{array}{c}\{\langle \mathit{call}(m, p(rec, \bar{x}, y)) \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho}\rangle | \mathcal{A}\} \circ \phi \circ e \rightsquigarrow_\alpha \{\langle \bar{b}^\alpha, \bar{\rho}\rangle, \langle {b'}^\alpha_1 \cdots {b'}^\alpha_n, \rho_1 \cdots \rho_{n+1}\rangle | \mathcal{A}\} \circ \\ \phi \wedge rec' = rec \wedge \bar{x}' = \bar{x} \wedge y' = y \circ e'\end{array}}$$

$$(3)_\alpha \quad \frac{\varphi \wedge \phi \not\models false, e' = \mathcal{M}(\varphi) + e}{\{\langle \varphi \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho}\rangle | \mathcal{A}\} \circ \phi \circ e \rightsquigarrow_\alpha \{\langle \bar{b}^\alpha, \bar{\rho}\rangle | \mathcal{A}\} \circ \phi \wedge \varphi \circ e'}$$

$$(4)_\alpha \quad \frac{e' = \mathcal{M}(\bot) + e}{\{\langle \bot \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho}\rangle | \mathcal{A}\} \circ \phi \circ e \rightsquigarrow_\alpha \{\langle \bar{b}^\alpha, \bar{\rho}\rangle | \mathcal{A}\} \circ \phi \circ e'}$$

$$(5)_\alpha \quad \frac{}{\{\langle \bot \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho}\rangle | \mathcal{A}\} \circ \phi \circ e \rightsquigarrow_\alpha \{\langle \bot \cdot \bar{b}^\alpha, \rho \cdot \bar{\rho}\rangle | \mathcal{A}\} \circ \phi \circ e}$$

$$(6)_\alpha \quad \frac{}{\{\langle \epsilon, \rho\rangle | \mathcal{A}\} \circ \phi \circ e \rightsquigarrow_\alpha \{\epsilon | \mathcal{A}\} \circ \phi \circ e}$$

$$(7)_\alpha \quad \frac{}{\mathcal{A} \circ \phi \circ e \rightsquigarrow_\alpha \mathcal{A} \circ \phi \circ e}$$

Figure 13. Semantics of abstract programs with cost annotations

Figure 13 depicts an abstract operational semantics derived from the one of Figure 9 by instrumenting the abstract states with a component that accumulates cost. Namely, an abstract state now has the form $\mathcal{A} \circ \phi \circ e$ where e is the amount of resources consumed so far. The instrumentation is straightforward: when executing b^α we simply accumulate the cost $\mathcal{M}(b^\alpha)$ that, by abusing of notation, we assume to be equivalent to $\mathcal{M}(b)$, i.e., to the cost of the original instruction from which b^α originate. Given an abstract trace $\mathcal{T}^\alpha \equiv \mathcal{A}_0 \circ true \circ 0 \rightsquigarrow^n_\alpha \mathcal{A}_n \circ \phi_n \circ e_n$, its cost is defined as $\mathcal{M}(\mathcal{T}^\alpha) = e_n$. Note that this instrumentation has no effect on the abstract executions, i.e., we still have the same abstract traces as those generated using the abstract semantics of Figure 9, and, moreover, Theorem 2 holds for the instrumented abstract semantics of Figure 13.

Now let $\mathcal{S}_0$ be an initial state, and let $\mathcal{T} \equiv \mathcal{S}_0 \rightsquigarrow^n \mathcal{S}_n$ be a concrete trace that starts from $\mathcal{S}_0$. Recall that the cost of $\mathcal{T}$, denoted $\mathcal{M}(\mathcal{T})$, is defined as the sum of all $\mathcal{M}(b)$ for each instruction b used in an execution step of $\mathcal{T}$. From Theorem 2, it immediately follows that there is a corresponding abstract trace $\mathcal{T}^\alpha \equiv \mathcal{A}_0 \circ true \circ \rightsquigarrow^n_\alpha \mathcal{A}_n \circ \phi_n \circ e_n$ such that $\mathcal{M}(\mathcal{T}^\alpha) = e_n = \mathcal{M}(\mathcal{T})$. This is because according to the proof in Appendix B, the instructions of $\mathcal{T}$ and $\mathcal{T}^\alpha$ coincide, i.e., whenever we make a concrete step that uses an instruction b, we can also make an abstract step that uses b^α. This means that any concrete trace $\mathcal{T}$, has a corresponding abstract one with the same cost.

Next we briefly explain why the cost relations generated from the abstract rules approximate the resource consumption of the abstract program, and thus the resource consumption of the original program. We do this by starting from the abstract program and the abstract semantics of Figure 13, and then modify them several times until we obtain the corresponding cost relations and the corresponding semantics [4].

In the first step, we consider a program that is obtained from the abstract program by removing all output variables, we refer to this program as output-free program. Removing output variables means syntactically removing them from the abstract rule heads and calls. Namely, a head $p(this, \bar{x}, \bar{y})$ is rewritten into $p(this, \bar{x})$, and a call $\mathit{call}(ct, q(rec, \bar{z}, \bar{w}))$ is rewritten into $\mathit{call}(ct, q(rec, \bar{z}))$. Clearly, any trace obtained using the abstract program has a corresponding trace that is obtained using the output-free program with the same resource consumption. This is true since the only difference is that in each step we might add less constraints to the store (we do not add those that match the formal and actual output parameters).

[†] Otherwise nothing must be proved.

DOI: 10.1002/stvr

In the second step, we change the abstract semantics such that instead of accumulating the resource consumption of each execution step, it accumulates the resource consumption of all abstract instructions immediately when they are added to the abstract state in rules $(1)_\alpha$ and $(2)_\alpha$. This change amounts to: (i) changing rules $(3)_\alpha$ and $(4)_\alpha$ such that they do not accumulate any cost, and (ii) changing rules $(1)_\alpha$ and $(2)_\alpha$ to accumulate also $c = \mathcal{M}(b_1^\alpha) + \cdots + \mathcal{M}(b_n^\alpha)$. Clearly, this change only anticipates the consumption of resources, and thus for any abstract trace that is obtained using the output-free program and the abstract semantics of Figure 13, we can generate a corresponding abstract trace using the same program and the modified abstract semantics such that it consumes at least the same amount of resources.

In the third step, we eliminate Rules $(3)_\alpha$-$(7)_\alpha$ from the abstract semantics and we modify rules $(1)_\alpha$ and $(2)_\alpha$ such that (i) they add all constraints that appear in the body of the selected rules (let us call them $\varphi = \varphi_1 \wedge \varphi_k$) to the store, and the rest, which are calls, are added as usual to the corresponding task. It is still guaranteed that using this abstract semantics we can reproduce the resource consumption of any trace generated in the above step. This is because the constraints in the body are obtained by applying a single static assignment transformation, thus for any $i > j$ the constraint φ_i does not restrict the values of the variables in φ_j.

Now let us consider an equation $p(\bar{x}) = c + \Sigma q_i(\bar{w}_i), \varphi$ in the cost relation. Here c and φ are the total resource consumption and the constraints of a given rule respectively (as above). It is easy to see that this equation is just a denotational form of the resource consumption as developed in the third step above. Thus, any upper-bound of the cost relation is also an upper bound in the resource consumption of the corresponding abstract traces.

The correctness for the object-sensitive case is straightforward given the soundness of the points-to analysis. The above proof can be adapted to the object-sensitive case by: (i) modifying the abstract program such that it includes corresponding points-to annotations; and (ii) change the abstract semantics in order to accumulate expressions of the form $c(o) * \mathcal{M}(b)$. The correctness of the points-to analysis guarantees that if in the concrete setting we accumulate $\mathcal{M}(b)$ when executing within object o', then in the abstract setting we accumulate $c(o) * \mathcal{M}(b)$ where o is the approximation of the o' inferred by the points-to analysis. Finally, cloning the equation as done in Definition 13 just makes the points-to information explicit in the rules names.

DOI: 10.1002/stvr