

Alma Mater Studiorum Università di Bologna
Archivio istituzionale della ricerca

Static analysis of cloud elasticity

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Abel, G., Laneve, C., Lienhardt, M. (2017). Static analysis of cloud elasticity. SCIENCE OF COMPUTER PROGRAMMING, 147, 27-53 [10.1016/j.scico.2017.03.008].

Availability:

This version is available at: <https://hdl.handle.net/11585/619228> since: 2020-12-23

Published:

DOI: <http://doi.org/10.1016/j.scico.2017.03.008>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Garcia, A., Laneve, C., & Lienhardt, M. (2017). Static analysis of cloud elasticity. *Science of Computer Programming*, 147, 27-53.

The final published version is available online at:
<http://dx.doi.org/10.1016/j.scico.2017.03.008>

Rights / License:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

Static analysis of cloud elasticity

Abel Garcia^a, Cosimo Laneve^a, Michael Lienhardt^b

^aDepartment of Computer Science and Engineering, University of Bologna – INRIA Focus

^bComputer Science Department, University of Torino, Italy

Abstract

We propose a *static analysis technique* that computes upper bounds of virtual machine usages in a *concurrent* language with explicit acquire and release operations of virtual machines. In our language it is possible to delegate other (ad-hoc or third party) concurrent code to release virtual machines (by passing them as arguments of invocations). Our technique is modular and consists of (i) a type system associating programs with behavioural types that record relevant information for resource usage (creations, releases, and concurrent operations), (ii) a translation function that takes behavioural types and returns cost equations, and (iii) an automatic off-the-shelf solver for the cost equations.

A soundness proof of the type system establishes the correctness of our technique with respect to the cost equations. We have experimentally evaluated our technique using a cost analysis solver and we report some results.

Keywords: Resource consumption analysis, cloud computing, behavioural type system, subject reduction, static analysis, concurrent programming

1. Introduction

The analysis of resource usage in a program is of great interest because an accurate assessment could reduce, for example, energy consumption and allocation costs. These two criteria are even more important today, in modern architectures like mobile devices or cloud computing, where resources, such as virtual machines, have hourly or monthly rates. In fact, cloud computing introduces the concept of *elasticity*, namely the possibility for virtual machines to scale according to the software needs. In order to support elasticity, cloud providers, including Amazon, Google, and Microsoft Azure, (i) have pricing models that allow one to hire on demand virtual machine instances and paying them for the time they are in use, and (ii) have APIs that include instructions for requesting and releasing virtual machine instances.

While it is relatively easy to estimate worst-case costs for simple code examples, extrapolating this information for fully real-life complex programs could be cumbersome and highly error-sensitive. The first attempts about the analysis of resource usage dates back to Wegbreit’s pioneering work in 1975 [28], which develops a technique for deriving closed-form expressions out of programs. The evaluation of these expressions would return upper-bound costs that are parametrised by programs’ inputs.

Wegbreit’s contribution has two limitations: it addresses a simple functional language and it does not formalize the connection between the language and the closed-form expressions. A number of techniques have been developed afterwards to cope with more expressive languages (see for instance [3, 10]) and to make the connection between programs and closed-form expressions precise (see for instance [12, 20]). A more detailed discussion of the related work in the literature is presented in Section 7.

To the best of our knowledge, current cost analysis techniques always address (concurrent) languages featuring only addition of resources. When removal of resources is considered, the techniques are used in a very constrained way [5]. On the other hand, cloud computing elasticity needs powerful acquire operations *as well as* release ones (see for instance the *Amazon Elastic Compute Cloud* [26] or the *Docker Fiware*). Let us consider the following

Email addresses: abel.garcia2@unibo.it (Abel Garcia), cosimo.laneve@unibo.it (Cosimo Laneve), michael.lienhardt@di.unito.it (Michael Lienhardt)

problem: given a pool of virtual machine instances and a program that acquires and releases these instances, what is the minimal cardinality of the pool guaranteeing the execution of the program without interruptions caused by lack of virtual machines? A solution to this problem, under the assumption that one can acquire a virtual machine that has been previously released, is useful both for cloud providers and for cloud customers. For the formers, it represents the possibility to estimate *in advance* the resources to allocate to a specific service. For the latter ones, it represents the possibility to pay *exactly* for the resources that are needed.

It is worth to notice that, without a full-fledged release operation, the cost of a concurrent program may be modeled by simply aggregating the sets of operations that can occur in parallel, as in [4]. By full-fledged release operation we mean that it is possible to delegate other (ad-hoc or third party) methods to release resources (by passing them as arguments of invocations). For example, consider the following method

```
Int double_release(Vm x, Vm y) {
    release x; release y; return 0 ;
}
```

that takes two machines and simply releases them. The cost of this method depends on the machines in input:

- it may be -2 when x and y are *different* and active;
- it may be -1 when one machine is active and the other is released, or when x and y are *equal* and active – consider the invocation `double_release(x, x)`;
- it may be 0 when the two machines have been already released.

In this case, one might over-approximate the cost of `double_release` to 0. However this leads to disregard releases and makes the analysis (too) imprecise.

In order to compute a precise cost of methods like `double_release`, in Section 4 we associate methods with abstract descriptions that carry information about resource usages. These descriptions are called *behavioural types* and are formally connected to the programs by means of a type system. The analysis of behavioural types is defined in Section 5 by translating them in a code that is adequate for an off-the-shelf solver – the CoFloCo solver [13]. As discussed in [7], in order to compute tight upper bounds, we have two functions per method: a function computing the *peak cost* – i.e. the worst case cost for the method to complete – and a function computing the *net cost* – i.e. the cost of the method after its completion. In fact, the functions that we associate to a method are much more than two. The point is that, if a method has two arguments – see `double_release` – and it is invoked with two *equal* arguments then its cost cannot be computed by a function taking two arguments, but it must be computed by a function with one argument only. This means that, for every method and *every partition of its arguments*, we define two cost functions: one for the peak cost and the other for the net cost.

The translation of behavioural types into CoFloCo input code has been prototyped and we are therefore able to automatically compute the cost of programs. In Section 6 we examine our prototype implementation and we report the results of some of our experiments. In particular, our technique excels in analyzing conditional branches, while it is less precise than multivariate amortized analysis [19] in presence of multi-variate arguments. We address this topic in Sections 6 and 7.

Our technique targets a simple concurrent language with explicit operations of creation and release of *virtual machines*. The language is defined in Section 2 and we discuss simplifications that ease the development of our technique in Section 3. In Section 8 we deliver concluding remarks. The correctness proof of the type system and with respect to the cost equations is reported in Appendix A. In the language we use the metaphor of cloud computing and virtual machines to illustrate our approach. However it is worth to observe that our technique may be also used for resource analysis of concurrent languages that bear operations of acquire (or creation) and release (such as heaps) as long as these operations are “simple”. For example, our technique might cover functions such as `malloc` and `free` as long as the sizes of memories that are allocated/freed are expressed in presburger arithmetics [9].

This article is a revised and enhanced version of [15] that includes a new semantics of the language, the full proofs of all the results and a discussion about our prototype implementation that includes the type inference system and the experimental assessments we did.

2. The language `vm1`

The language `vm1` is a future-based concurrent object-oriented language¹ with explicit acquire and release operations of virtual machines. In future-based languages, function/method invocations are *executed asynchronously* to the caller and are bound to variables called *futures*. The caller synchronises with the callee by using the future variables when the return value is strictly needed. The syntax and the semantics of `vm1` are defined in the following two subsections; the third subsection discusses a number of examples.

Syntax. In `vm1` we distinguish between *simple types* T , which are either integers `Int` or virtual machines `VM`, and *future types* $\text{Fut}\langle T \rangle$, which type asynchronous invocations. The future argument T is instantiated with the return type of the invoked function. We use F to range over simple and future types. The notation $\overline{T x}$ denotes any finite sequence of *name declaration* $T x$ separated by commas. Similarly we write $\overline{T x}$; for a finite sequence of name declaration separated by semicolons.

A `vm1` program is a sequence of method definitions $T \text{ m}(\overline{T x})\{\overline{F y}; s\}$, ranged over by M , plus a main body $\{\overline{F z}; s'\}$. The syntax of statements s , expressions z and pure expressions e of `vm1` is defined by the following grammar:

s	$::=$	$x = z \mid \text{if } e \{s\} \text{ else } \{s\} \mid \text{return } e \mid s ; s \mid \text{release } e$	statements
z	$::=$	$e \mid e!m(\overline{e}) \mid e.\text{get} \mid \text{new VM}()$	expressions
e	$::=$	$\text{this} \mid se \mid nse$	pure expressions

A statement s may be either one of the standard operations of an imperative language plus the `release` e operation that disposes the virtual machine e .

An expression z may change the state of the system. In particular, it may be an *asynchronous* method call $e!m(\overline{e})$ where e is the virtual machine that will execute the call, and \overline{e} are the arguments of the call. This invocation does not suspend caller's execution: when the value computed by the invocation is needed then the caller performs a *non blocking* `get` operation, if the value needed by a process is not available then an awaiting process is scheduled and executed. Expressions z also include `new VM()` that creates a new virtual machine. The intended meaning of operations taking place on different virtual machines is that they may execute in parallel, while operations in the same virtual machine interleave their evaluation (even if in the following operational semantics the parallelism is not explicit).

A (*pure*) expression e may be the reserved identifier `this`, a virtual machines identifier or an integer expression. Since our analysis will be parametric with respect to the inputs, we parse integer expressions in a careful way. In particular we split them into *size expressions* se , which are expressions in Presburger arithmetics [9] (this is a decidable fragment of Peano arithmetics that only contains addition), and *non-size expressions* nse , which are the other type of expressions. The syntax of size and non-size expressions is the following:

nse	$::=$	$k \mid x \mid nse \leq nse \mid \neg nse \mid nse \text{ and } nse \mid nse \text{ or } nse \mid nse + nse \mid nse - nse \mid nse \times nse \mid nse / nse$	non-size expressions
se	$::=$	$ve \mid ve \leq ve \mid \neg se \mid se \text{ and } se \mid se \text{ or } se$	size expressions
ve	$::=$	$k \mid x \mid ve + ve \mid k \times ve$	integer size expressions
k	$::=$	$\text{integer constants} \mid \text{err}$	

We will use the term $nse = nse'$ as an abbreviation of “ $(nse \leq nse')$ and $(nse' \leq nse)$ ”, and similarly for se . We notice that (non-size and size) expressions also contain the value `err` – see below the definition of $\llbracket e \rrbracket$, for the semantics of arithmetics expressions that contain `err`. In the whole paper, we assume that sequences of declarations $\overline{T x}$ and method declarations \overline{M} do not contain duplicate names. We also assume that `return` statements never have a continuation.

Semantics. `vm1` semantics is defined as a transition relation between *configurations*, noted cn and defined below

cn	$::=$	$\epsilon \mid \text{fut}(f, v) \mid \text{vm}(o, a, p, q) \mid \text{invoc}(o, f, m, \overline{v}) \mid cn \text{ } cn$	configurations
p	$::=$	$\{l \mid \epsilon\} \mid \{l \mid s\}$	process
q	$::=$	$\epsilon \mid p \mid q \text{ } q$	sets of processes
v	$::=$	$\text{integer constants} \mid o \mid f \mid \perp \mid \text{err}$	run-time values
a	$::=$	$\top \mid \perp$	machine states
l	$::=$	$[\dots, x \mapsto v, \dots]$	maps

¹See [17] for more details on a similar language

$$\begin{array}{c}
\text{(ASSIGN)} \\
\frac{v = \llbracket e \rrbracket_l}{vm(o, \top, \{l \mid x = e; s\}, q) \rightarrow vm(o, \top, \{l \mid x \mapsto v\} \mid s\}, q)} \\
\\
\text{(BIND-MTD)} \\
\frac{\{l \mid s\} = \text{bind}(o, f, \bar{m}, \bar{v})}{vm(o, \top, p, q) \text{ invoc}(o, f, \bar{m}, \bar{v}) \rightarrow vm(o, \top, p, q \cup \{l \mid s\}) \text{ fut}(f, \perp)} \\
\\
\text{(NEW-VM)} \\
\frac{o' = \text{fresh}(\text{VM})}{vm(o, \top, \{l \mid x = \text{new VM}(); s\}, q) \rightarrow vm(o, \top, \{l \mid x = o'; s\}, q) \text{ vm}(o', \top, \{\emptyset \mid \varepsilon\}, \emptyset)} \\
\\
\text{(ACTIVATE)} \\
\frac{}{vm(o, a, \{l' \mid \varepsilon\}, q \cup \{l \mid s\}) \rightarrow vm(o, a, \{l' \mid s\}, q)} \\
\\
\text{(VM-ERR-RETURN)} \\
\frac{f = l(\text{destiny})}{vm(o, \perp, \{l \mid s\}, q) \text{ fut}(f, \perp) \rightarrow vm(o, \perp, \{l; \varepsilon\}, q) \text{ fut}(f, \text{err})} \\
\\
\text{(READ-FUT)} \\
\frac{f = \llbracket e \rrbracket_l \quad v \neq \perp}{vm(o, \top, \{l \mid x = e.\text{get}; s\}, q) \text{ fut}(f, v) \rightarrow vm(o, \top, \{l \mid x = v; s\}, q) \text{ fut}(f, v)} \\
\\
\text{(COND-TRUE)} \\
\frac{\llbracket e \rrbracket_l \neq 0 \quad \llbracket e \rrbracket_l \neq \text{err}}{vm(o, \top, \{l \mid \text{if } e \text{ then } \{s_1\} \text{ else } \{s_2\}; s\}, q) \rightarrow vm(o, \top, \{l \mid s_1; s\}, q)} \\
\\
\text{(COND-FALSE)} \\
\frac{\llbracket e \rrbracket_l = 0 \quad \text{or} \quad \llbracket e \rrbracket_l = \text{err}}{vm(o, \top, \{l \mid \text{if } e \text{ then } \{s_1\} \text{ else } \{s_2\}; s\}, q) \rightarrow vm(o, \top, \{l \mid s_2; s\}, q)} \\
\\
\text{(RELEASE-VM)} \\
\frac{o' = \llbracket e \rrbracket_l \quad o \neq o'}{vm(o, \top, \{l \mid \text{release } e; s\}, q) \text{ vm}(o', a', p', q') \rightarrow vm(o, \top, \{l \mid s\}, q) \text{ vm}(o', \perp, p', q')} \\
\\
\text{(RELEASE-VM-SELF)} \\
\frac{o = \llbracket e \rrbracket_l}{vm(o, \top, \{l \mid \text{release } e; s\}, q) \rightarrow vm(o, \perp, \{l \mid s\}, q)} \\
\\
\text{(RETURN)} \\
\frac{v = \llbracket e \rrbracket_l \quad f = l(\text{destiny})}{vm(o, \top, \{l \mid \text{return } e\}, q) \text{ fut}(f, \perp) \rightarrow vm(o, \top, \{l \mid \varepsilon\}, q) \text{ fut}(f, v)} \\
\\
\text{(BIND-MTD-ERR)} \\
\frac{}{vm(o, \perp, p, q) \text{ invoc}(o, f, \bar{m}, \bar{v}) \rightarrow vm(o, \perp, p, q) \text{ fut}(f, \text{err})} \\
\\
\text{(BIND-PARTIAL)} \\
\frac{}{\text{invoc}(\text{err}, f, \bar{m}, \bar{v}) \rightarrow \text{fut}(f, \text{err})} \\
\\
\text{(CONTEXT)} \\
\frac{cn \rightarrow cn'}{cn \text{ } cn'' \rightarrow cn' \text{ } cn''}
\end{array}$$

Figure 1: Semantics of `vm1`.

Configurations are sets of elements – therefore we identify configurations that are equal up-to associativity and commutativity – and are denoted by the juxtaposition of the elements $cn \text{ } cn$; the empty configuration is denoted by ϵ . The transition relation uses two infinite sets of names: *vm names*, ranged over by o, o', \dots and *future names*, ranged over by f, f', \dots . The function `fresh()` returns either a fresh vm name or a fresh future name; the context will disambiguate between the two. The elements of configurations are

- *virtual machines* $vm(o, a, p, q)$ where o is a vm name; a is either \top or \perp depending on whether the machine is alive or dead; p is either $\{l \mid \epsilon\}$, representing a terminated statement, or $\{l \mid s\}$, representing an *active process*, where l maps each local variable to its value and s is the statement to execute; and q is the set of processes to evaluate.
- *future binders* $\text{fut}(f, v)$. When the value v is \perp then the actual value of f has still to be computed.
- *method invocation messages* $\text{invoc}(o, f, \bar{m}, \bar{v})$.

Runtime values v are either integers or virtual machines and future names, or \perp , meaning an un-computed value, or an erroneous value *err*. The following auxiliary functions are used in the semantic rules (we assume a fixed `vm1` program):

- $\text{dom}(l)$ returns the domain of l .
- $l[x \mapsto v]$ is the function such that $(l[x \mapsto v])(x) = v$ and $(l[x \mapsto v])(y) = l(y)$, when $y \neq x$.
- $\llbracket e \rrbracket_l$ returns the value of e , possibly retrieving the values of the names that are stored in l . As regards boolean operations, as usual, `false` is represented by 0 and `true` is represented by a value different from 0. Arithmetic operations in `vm1` are also defined on the value *err*: when one of the arguments is *err*, every arithmetic operation returns *err*. $\llbracket \bar{e} \rrbracket_l$ returns the tuple of values of \bar{e} . When e is a future name, the function $\llbracket \cdot \rrbracket_l$ is the identity. Namely $\llbracket f \rrbracket_l = f$. It is worth to notice that $\llbracket e \rrbracket_l$ is undefined whenever e contains a name that is not defined in l .
- $\text{bind}(o, f, \bar{m}, \bar{v}) = \{[\bar{x} \mapsto \bar{v}, \text{destiny} \mapsto f] \mid s[\bar{o}/\text{this}]\}$, where $T \text{ m}(\overline{T x})\{\overline{T' z}; s\}$ is a method of the program. We observe that, because of `bind`, the map l in processes $\{l \mid s\}$ also binds the special name `destiny` to a future value. We also observe that the local names \bar{z} do not belong to $\text{dom}([\bar{x} \mapsto \bar{v}, \text{destiny} \mapsto f])$.

The transition relation rules are collected in Figure 1. The rules are almost standard, except those about the management of virtual machines and method invocations, which we are going to discuss.

Rule (NEW-VM) creates a virtual machine and makes it alive. Rules (RELEASE-VM) and (RELEASE-VM-SELF) dispose a virtual machine by means of the operation `release x`: this amounts to update its state a to \perp . Once a virtual machine has been released, no operation can be performed anymore by it, except letting the processes on the queue returning err – see rules (ACTIVATE), (VM-ERR-RETURN), and (BIND-MTD-ERR).

Rule (ASYNC-CALL) defines asynchronous method invocation $x = e!m(\vec{e})$. This rule creates a fresh future name that is assigned to the identifier x . Rule (BIND-MTD) applies an invocation message by adding to the live callee virtual machine a process corresponding to the called method. In case the callee virtual machine is not live, either (BIND-MTD-ERR) or (BIND-PARTIAL) is applied, which binds err to the future name. Rule (READ-FUT) allows the caller to retrieve the value returned by the callee. It is worth to notice that the semantics of `get` is different from that of ABS [25] or of [17] because it is not blocking with respect to other processes waiting to be executed on the same machine, see rule (ACTIVATE-GET). In fact, deadlock freedom is out of the scope of this contribution; in any case we refer to [17] for an algorithm (and a prototype) verifying deadlock freedom in ABS.

The initial configuration of a `vm1` program with main body $\overline{F x ; s}$ is

$$vm(start, \top, \{[destiny \mapsto f_{start}] | s[start/this]\}, \emptyset)$$

where $start$ is a special vm name and f_{start} is a fresh future name. As usual, let \longrightarrow^* be the reflexive and transitive closure of \longrightarrow and represents (part of) computations.

Examples. We illustrate `vm1` by discussing few examples and, for every example, we also examine the output we expect from our cost analysis. We begin with two methods computing the factorial function:

```
Int fact(Int n){
  Fut<Int> x ; Int m ;
  if (n==0) { return 1 ; }
  else { x = this!fact(n-1) ;
        m = x.get ;
        return m*n ;
  }
}
```

```
Int costly_fact(Int n){
  Fut<Int> x ; Int m ; VM z ;
  if (n==0) { return 1 ; }
  else { z = new VM() ;
        x = z!costly_fact(n-1) ; m = x.get ;
        release z ; return m*n ;
  }
}
```

The method `fact` is the standard definition of factorial with the recursive invocation `fact(n-1)` always performed on the same machine. That is, the computation of `fact(n)` only requires one virtual machine. On the contrary, the method `costly_fact` performs the recursive invocation on a new virtual machine z . The caller waits for its result, let it be m , then it releases z and delivers the value $m*n$. Since every virtual machine creation occurs before any release operation, `costly_fact` will create as many virtual machines as the argument n . That is, if the available resources are k virtual machines, then `costly_fact` can compute factorials up to k .

The analysis of `costly_fact` has been easy because the `release` operation is applied to a locally created virtual machine. Yet, in `vm1`, `release` can be also applied to method arguments and the presence of this feature in concurrent codes is a major source of difficulties for the analysis. A paradigmatic example is the `double_release` method discussed in Section 1 that may have either a cost of -2 or of -1 or of 0 .

It is worth to observe that, while over-approximations (e.g not counting releases) return (too) imprecise costs, under-approximations may return wrong costs. For example, the following method creates two virtual machines and releases the second one with `this!double_release(x,x)` before the recursive invocation.

```
Int fake_method(Int n) {
  if (n==0) return 0 ;
  else { VM x, y ; Fut<Int> f, g ; Int u, v ;
        x = new VM() ; y = new VM() ;
        f = this!double_release(x,x) ; u = f.get ;
        g = this!fake_method(n-1) ; v = g.get ;
        return 0 ; }
}
```

The cost of `fake_method(n)` should be n . However this is not the case if `double_release` is under-approximated with cost -2 : In such case, one would wrongly derive a cost 0 of `fake_method(n)`. In Section 6.2 we consider an erroneous `fake_method` that increases the argument of the recursive invocation (instead of decreasing it) and discuss the corresponding cost equations returned by our technique. We notice that, in this case, the amount of virtual machines used by the erroneous method is infinite. The aim of the following sections is to present a technique for determining the cost of method invocations. Such a technique has to be sensible to the identity, and to the state of method's arguments before and after the invocation.

Alive machines. Given a configuration cn , a term $vm(o, \tau, p, q) \in cn$ is called *alive machine* in cn . Let $\text{alive}(cn)$ be the number of different alive machines in cn . The objective of this paper is to define a technique such that, given a configuration cn , it returns a n satisfying, for every $cn \rightarrow^* cn'$, $\text{alive}(cn') \leq n$ (n is an upper bound to the alive machines in computations rooted at cn).

3. Determinacy of releases of method's arguments

Our cost analysis of virtual machines uses abstract descriptions that carry informations about concurrent method invocations and about creations and removals of virtual machines. In order to ease the compositional reasonings, method's abstract descriptions also define the arguments the method releases upon termination, called *method's effects*. In this contribution we stick to method descriptions that are as simple as possible, namely we assume that method's effects are *sets*. In turn, this requires methods' behaviours to be *deterministic* with respect to method's effects and, to enforce this determinacy, we define the following simplification properties.

The following simplifications allow us to focus on the relevant problems of the resource analysis of `vm1`. It is possible to drop most of the simplifications by extending the approach presented in this article using almost standard solutions (for instance, with guarded or nondeterministic effects, see below). However these solutions would entangle a lot the technicalities of the paper, making more difficult its reading and comprehension.

Simplification 1: *the branches in a method body always release the same set of method's arguments.* For example, methods like

```
Int foo1(VM x, Int n) {
  if (n = 0) return 0 ;
  else { release x ; return 0 ; }
}
```

does not follow this simplification because the then-branch does not release anything while the else-branch releases the argument x .

This simplification is typical in type systems with (deterministic) effects. Usually people takes the least-upper bound of effects; however we cannot take least-upper bounds of removals because it would mean to underestimate the machines effectively used. We cannot take the greatest lower bound of removals either, because it would mean to over-estimate the machines used and perform erroneous releases in the continuations. Indeed, to analyse the above program, one would just have to extend the approach presented in this article with *guarded effects*, i.e., mapping from execution branches to effects, which can precisely identify the effects of each execution branch of the method.

Simplification 2: *method invocations are always synchronized within caller's body.* This implies that method's effects occur upon method termination. For example, in

```
Int foo2(VM x, VM y) {
  this!double_release(x,y) ; return 0 ;
}
```


the method's effects of `foo2` is not deterministic because `double_release` might still be running *after* `foo2` termination. This means that the termination of `foo2` has no impact on the termination of `double_release`. Hence, the effect of `foo2` is empty – the caller cannot assume that `x` and `y` have been released upon synchronizing with `foo2`. Overall, this simplification supports a more precise analysis and can be easily dropped (at the cost of returning larger upper bounds).

Simplification 3: *machines executing methods with nonempty effects must be alive.* (This includes the carrier machine, e.g. method bodies cannot release the `this` machine.) At static time “alive” means that the machine is either the caller or has been locally created and has not been/is not being released. For example, in `foo3`

```
Int simple_release(VM x) { release x; return 0; }
Int foo3(VM x) {
  VM z ; Fut<Int> f ; Int u ;
  z = new VM() ; f = z!simple_release(x) ;
  release z ; u = f.get ; return 0;
}
```

```
Int foo4(VM x, VM y) {
  Fut<Int> f ; Int u ;
  f = x!simple_release(y) ;
  u = f.get ; return 0 ;
}
```

the machine `z` is released before the synchronisation with the `simple_release` – statement `f.get`. This means that the disposal of `x` depends on scheduler's choice and, in turn, it is not possible to determine whether `foo3` will release `x` or not. A similar non-determinism arises when the callee of a method releasing arguments is itself an argument. For example, in the above `foo4` method, it is not possible to determine whether `y` is released or not because we have no clue about `x`, being it an argument of `foo4`.

To analyse such programs, one would simply need to extend the approach presented in this article with non-deterministic effects.

Simplification 4: *if a method returns a machine, the machine must be new.* For example, consider the following code:

```
VM identity(VM x) { return x; }
{
  VM x ; VM y ; VM z ; Fut<VM> f ; Fut<Int> g ; Int m ;
  x = new VM() ; y = new VM() ;
  f = y!identity(x) ; g = this!simple_release(y) ;
  z = f.get ; m = g.get ;
  release z ;
}
```

In this case it is not possible to determine whether the value of `z` is `x` or `err` and, therefore, it is not clear whether the cost of `release z` is 0 or -1. The non-determinism is caused by `identity`, which returns the argument that is going to be released by a parallel method.

To analyse programs with such methods, one would simply need to extend the approach presented in this article with non-deterministic effects.

Simplifications 1, 3, and 4 are enforced by the type system in Section 4, in particular simplification 1 by rule (T-METHOD), simplification 3 by rules (T-INVOKE) and (T-RELEASE), and simplification 4 by rules (T-INVOKE) and (T-RETURN).

4. The behavioural type system of `vm1`

Behavioural types are abstract codes highlighting the features of `vm1` programs that are relevant for the resource cost analysis in Section 5. These types support compositional reasonings and are associated to programs by means of a type system that is defined in this section.

The syntax of behavioural types uses *vm names* $\alpha, \beta, \gamma, \dots$, and *future names* f, f', \dots . Sets of `vm names` will be ranged over by S, S', R, \dots , and sets of future names will be ranged over by F, F', \dots . We assume that *err* is a *special*

\circ	$::=$	$_ \mid \alpha$	basic value
\mathfrak{t}	$::=$	$\alpha \downarrow \mid \partial \mid \perp \mid \top$	vm value
op'	$::=$	$+ \mid - \mid = \mid \leq \mid \geq \mid \wedge \mid \vee$	linear operation
\mathbb{T}, \mathbb{S}	$::=$	$\circ \mid se$	typing value
\mathbb{Z}	$::=$	$(\circ, \alpha, \text{Ft}, \text{R}) \mid \circ$	future value
\mathbb{X}	$::=$	$_ \mid \text{Ft} \mid f \mid \mathbb{Z}$	extended value
\mathfrak{a}	$::=$	$0 \mid \nu\alpha \mid \nu f : \mathfrak{m} \alpha(\overline{\mathbb{S}}) \mid \alpha^\checkmark \mid f_\circ^\checkmark$	atom
\mathbb{C}	$::=$	$\mathfrak{a} \triangleright \Gamma \mid \mathfrak{a} \S \mathbb{C} \mid (se)\{\mathbb{C}\} \mid \mathbb{C} + \mathbb{C}$	behavioural type

Figure 2: Behavioural Types Syntax

(T-VAR)	(T-PRIMITIVE)	(T-OP)	T-UNIT	(T-OP-UNIT)
$\frac{x \in \text{dom}(\Gamma)}{\Gamma \vdash x : \Gamma(x)}$	$\Gamma \vdash k : k$	$\frac{\Gamma \vdash e_1 : se_1 \quad \Gamma \vdash e_2 : se_2}{\Gamma \vdash e_1 \text{ op}' e_2 : se_1 \text{ op}' se_2}$	$\frac{\Gamma \vdash e : se}{\Gamma \vdash e : _}$	$\frac{\Gamma \vdash e_1 : _ \text{ or } \Gamma \vdash e_2 : _ \text{ or } \text{op} \in \{*, /\}}{\Gamma \vdash e_1 \text{ op } e_2 : _}$
	(T-PURE)	(T-METHOD-SIG)		
	$\frac{\Gamma \vdash e : \mathbb{X}}{\Gamma \vdash e : \mathbb{X}, 0 \triangleright \Gamma}$	$\frac{\Gamma(\mathfrak{m}) = \alpha(\overline{\mathbb{T}}) : \circ, \text{R} \quad \overline{\beta} \subseteq \text{fv}(\alpha, \overline{\mathbb{T}}, \circ) \quad \sigma \text{ is a vm renaming such that } \circ \notin \text{fv}(\alpha, \overline{\mathbb{T}}) \text{ implies } \sigma(\circ) \text{ fresh}}{\Gamma \vdash \mathfrak{m} \sigma(\alpha)(\sigma(\overline{\mathbb{T}})) : \sigma(\circ), \sigma(\text{R})}$		

Figure 3: Typing rules for pure expressions

vm name representing the erroneous machine; therefore vm names will also range over *err*. The syntactic rules are presented in Figure 2.

Behavioural types $\nu\alpha$ and α^\checkmark express creations of virtual machines and their removal, respectively. The type $\nu f : \mathfrak{m} \alpha(\overline{\mathbb{S}})$ defines method invocations and f_\circ^\checkmark defines the corresponding synchronisation with the computation of the future f , in this case \circ may be either α , which acts like a binder to the name of the returned virtual machine, or $_$ if the target of f returns a different type. The conditional type $(se)\{\mathbb{C}\}$ behaves like \mathbb{C} whenever se is true; the type $\mathbb{C} + \mathbb{C}'$ models nondeterminism. We will always shorten the type f_\circ^\checkmark into f^\checkmark .

In order to have a more precise type of continuations, the leaves of behavioural types are labelled with *environments*, ranged over by Γ, Γ', \dots . Environments are maps from method names \mathfrak{m} to terms $\alpha(\overline{\mathbb{T}}) : \circ, \text{R}$, from names to extended values \mathbb{X} , from future names to future values, and from vm names to extended values Ft , which are called *vm states* in the following. We assume that $\Gamma(\text{err}) = \emptyset \perp$. These environments occurring in the leaves are only used in the typing proofs and are dropped in the final types (method types and the main statement type).

Vm states Ft are a collection F of future names plus the value \mathfrak{t} of the virtual machines. This F specifies the set of parallel methods that are going to release the virtual machine; \mathfrak{t} defines whether the virtual machine is alive (\top), or it has been already released (\perp) or, according to scheduler's choices, it may be either alive or released (∂). Vm values also include terms α and $\alpha\downarrow$. The value α is given to the argument machines of methods (they will be instantiated by the invocations – see the cost analysis in Section 5), the value $\alpha\downarrow$ is given to argument values that are returned by methods *and* can be released by parallel methods ($\alpha\downarrow$ will be also evaluated in the cost analysis). Vm values are partially ordered by the relation \leq defined by

$$\partial \leq \top \quad \partial \leq \perp \quad \alpha\downarrow \leq \perp \quad \alpha\downarrow \leq \alpha .$$

In the following we will use the partial operation $\mathfrak{t} \sqcap \mathfrak{t}'$ returning, whenever it exists, the greatest lower bound between \mathfrak{t} and \mathfrak{t}' . For example $\top \sqcap \perp = \partial$, but $\partial \sqcap \alpha\downarrow$ is not defined.

The type system uses judgments of the following form:

- $\Gamma \vdash e : \mathbb{X}$ for pure expressions e , $\Gamma \vdash f : \mathbb{Z}$ for future names f , and $\Gamma \vdash \mathfrak{m} \alpha(\overline{\mathbb{T}}) : \circ, \text{R}$ for methods.
- $\Gamma \vdash_S z : \mathbb{X}, \mathbb{C} \triangleright \Gamma'$ for expressions z , where \mathbb{X} is the value, \mathbb{C} is the behavioural type for z and Γ' is the environment Γ *with updates* of names and future names.

- $\Gamma \vdash_S s : \mathbb{C}$, in this case the updated environments are inside the behavioural type Γ' , in correspondence of every branch of its.

The index S in the judgments for expressions and statements defines the set of method's arguments – see rule (T-METHOD) – and is used in the rule (T-RETURN) in order to constrain that the returned machine, if any, does not belong to method's arguments (*cf.* Restriction 4 in Section 3).

Since Γ is a function, we use the standard predicates $x \in \text{dom}(\Gamma)$ or $x \notin \text{dom}(\Gamma)$ and the environment update

$$\Gamma[x \mapsto \mathbb{X}](y) \stackrel{\text{def}}{=} \begin{cases} \mathbb{X} & \text{if } y = x \\ \Gamma(y) & \text{otherwise} \end{cases}$$

With an abuse of notation (see rule (T-RETURN)), we let $\Gamma[_ \mapsto \mathbb{X}] \stackrel{\text{def}}{=} \Gamma$ (because $_$ does not belong to any environment).

We will also use the operation and notation below:

- $F \Downarrow$ is defined as follows:

$$F \Downarrow \stackrel{\text{def}}{=} \begin{cases} \Downarrow & \text{if } F = \emptyset \\ \partial & \text{if } F \neq \emptyset \text{ and } \Downarrow = \top \\ \alpha \downarrow & \text{if } F \neq \emptyset \text{ and } \Downarrow = \alpha \end{cases}$$

and, in Section 5, we write $(F_1 \Downarrow_1, \dots, F_n \Downarrow_n) \Downarrow$ for $(F_1 \Downarrow_1 \Downarrow, \dots, F_n \Downarrow_n \Downarrow)$.

- the *multihole contexts* $C[]$ defined by the following syntax:

$$C[] ::= [] \mid \mathbb{A} \circ C[] \mid C[] + C[] \mid (se)\{C[]\}$$

and, whenever $\mathbb{C} = C[\mathbb{A}_1 \triangleright \Gamma_1] \dots [\mathbb{A}_n \triangleright \Gamma_n]$, then $\mathbb{C}[x \mapsto \mathbb{X}]$ is defined as $C[\mathbb{A}_1 \triangleright \Gamma_1[x \mapsto \mathbb{X}]] \dots [\mathbb{A}_n \triangleright \Gamma_n[x \mapsto \mathbb{X}]]$.

The type system for expressions is reported in Figure 3. It is worth to notice that this type system is not standard because (size) expressions containing method's arguments are typed with the expressions themselves. This is crucial in the cost analysis of Section 5. It is also worth to observe that, by rule (T-PRIMITIVE), $\Gamma \vdash \text{err} : \text{err}$ (while $\Gamma(\text{err}) = \emptyset \perp$). This expedient allows us to save one rule when typing method invocations with either erroneous or released carriers (*cf.* rule (T-INVOK-BOT)).

The type system for expressions with side effects and statements is reported in Figure 4. We discuss rules (T-INVOK), (T-GET), (T-RELEASE), and (T-RETURN). Rule (T-INVOK) types method invocations $e!m(\bar{e})$ by using a fresh future name. In particular, in the behavioural type $\nu f : \mathbb{M} \alpha(\mathbb{S}) \rightarrow \beta$, the (fresh) future name f is associated to the the method, the vm name of the callee, the arguments *and to the returned value*. This last value is fundamental when the invocation returns a new machine because, in this case, the type acts as a binder of β . Another important remark is about the value of f in the updated environment. This value contains the returned value, the vm name of the callee and its state, and the set of the arguments that the method is going to remove. The vm state of the callee will be used when the method is synchronized to update the state of the returned object, if any (see rule (T-GET)). It is also important to observe that the environment returned by (T-INVOK) is updated with information about vm names released by the method: every such name will contain f in its state. Let us now discuss the constraints in the second and third lines of the premise of (T-INVOK). As regards the second line, assuming that the callee has not been already released ($\Gamma(\alpha) \neq F \perp$), there are two cases:

- (i) either $\Gamma(\alpha) = \emptyset \top$ or α is the caller object α' : namely the callee is alive because it has been created by the caller or it is the caller itself,
- (ii) or $\Gamma(\alpha) \neq \emptyset \top$: this case has two subcases, namely either (ii.a) the callee is being released by a parallel method or (ii.b) it is an argument of the caller method – see rule (T-METHOD).

While in (i) we admit that the invoked method releases vm names, in case (ii) we forbid any release, as we discussed in Restriction 3 in Section 3. We observe that, in case (ii.b), being α an argument of the method, it may retain any state when the method is invoked and, for reasons similar to (ii.a), it is not possible to determine at static time the exact subset of \mathbb{R} that will be released. The constraint in the third line of the premise of (T-INVOK) enforces Restriction 3 to the other invocations in parallel and to the object executing $e!m(\bar{e})$.

$$\begin{array}{c}
\text{(T-Assign-Var)} \\
\frac{\Gamma(x) = \mathbb{X} \quad \Gamma \vdash_S z : \mathbb{X}', \mathbb{C}}{\Gamma \vdash_S x = z : \mathbb{C}[x \mapsto \mathbb{X}']} \\
\\
\text{(T-Invoke)} \\
\frac{\Gamma \vdash e : \alpha \quad \Gamma \vdash \bar{e} : \bar{\mathbb{S}} \quad \Gamma \vdash m(\bar{\mathbb{S}}) : \mathbb{O}, R \quad \Gamma \vdash \text{this} : \alpha' \quad \Gamma(\alpha) \neq F\perp \\
(\Gamma(\alpha) \neq \emptyset\top \text{ and } \alpha \neq \alpha') \implies R = \emptyset \\
R \cap (\{\alpha'\} \cup \{\beta \mid f' \in \text{dom}(\Gamma) \text{ and } \Gamma(f') = (\mathbb{O}', \beta, \mathbb{X}, R') \text{ and } R' \neq \emptyset\}) = \emptyset \\
f \text{ fresh} \quad \Gamma' = \Gamma[\beta \mapsto (\{f\} \cup F')\mathbb{t}]^{\beta \in R, \Gamma(\beta) = F'\mathbb{t}}}{\Gamma \vdash_S e!m(\bar{e}) : f, \nu f : m(\bar{\mathbb{S}}) \triangleright \Gamma'[f \mapsto (\mathbb{O}, \alpha, \Gamma(\alpha), R)]} \\
\\
\text{(T-Get)} \\
\frac{\Gamma \vdash x : f \quad \Gamma \vdash f : (\mathbb{O}, \alpha, F\mathbb{t}, R) \\
R' = f\nu(\mathbb{O}) \setminus R \quad \mathbb{t}' = \mathbb{t} \sqcap (\Gamma(\alpha) \Downarrow) \\
\Gamma' = \Gamma[\beta \mapsto \emptyset\perp]^{\beta \in R}[\beta' \mapsto \emptyset\mathbb{t}']^{\beta' \in R'}}{\Gamma \vdash_S x.\text{get} : \mathbb{O}, f_{\mathbb{O}} \checkmark \triangleright \Gamma'[f \mapsto \mathbb{O}]} \\
\\
\text{(T-Get-Done)} \\
\frac{\Gamma \vdash x : f \quad \Gamma \vdash f : \mathbb{O}}{\Gamma \vdash_S x.\text{get} : \mathbb{O}, 0 \triangleright \Gamma} \\
\\
\text{(T-Invoke-Bot)} \\
\frac{\Gamma \vdash e : \alpha \quad \Gamma \vdash \bar{e} : \bar{\mathbb{S}} \quad \Gamma(\alpha) = F\perp \quad f \text{ fresh}}{\Gamma \vdash_S e!m(\bar{e}) : f, \nu f : m(\bar{\mathbb{S}}) \triangleright \Gamma[f \mapsto (err, \alpha, F\perp, \emptyset)]} \\
\\
\text{(T-New)} \\
\frac{\beta \text{ fresh}}{\Gamma \vdash_S \text{new VM}() : \beta, \nu \beta \triangleright \Gamma[\beta \mapsto \emptyset\top]} \\
\\
\text{(T-Release)} \\
\frac{\Gamma \vdash x : \alpha \quad \Gamma(\alpha) \neq F\perp \\
\alpha \notin \{\beta \mid f' \in \text{dom}(\Gamma) \text{ and } \Gamma(f') = (\mathbb{O}', \beta, \mathbb{X}, R') \text{ and } R' \neq \emptyset\}}{\Gamma \vdash_S \text{release } x : \alpha' \triangleright \Gamma[\alpha \mapsto \emptyset\perp]} \\
\\
\text{(T-Release-Bot)} \\
\frac{\Gamma \vdash x : \alpha \quad \Gamma(\alpha) = F\perp}{\Gamma \vdash_S \text{release } x : 0 \triangleright \Gamma} \\
\\
\text{(T-If)} \\
\frac{\Gamma \vdash e : se \quad \Gamma \vdash_S s_1 : \mathbb{C}_1 \quad \Gamma \vdash_S s_2 : \mathbb{C}_2}{\Gamma \vdash_S \text{if } e \{ s_1 \} \text{ else } \{ s_2 \} : (se)(\mathbb{C}_1) + (\neg se)(\mathbb{C}_2)} \\
\\
\text{(T-If-ND)} \\
\frac{\Gamma \vdash e : _ \quad \Gamma \vdash_S s_1 : \mathbb{C}_1 \quad \Gamma \vdash_S s_2 : \mathbb{C}_2}{\Gamma \vdash_S \text{if } e \{ s_1 \} \text{ else } \{ s_2 \} : \mathbb{C}_1 + \mathbb{C}_2} \\
\\
\text{(T-Seq)} \\
\frac{\Gamma \vdash_S s_1 : C[\mathbb{a}_1 \triangleright \Gamma_1] \cdots [\mathbb{a}_n \triangleright \Gamma_n] \quad (\Gamma_i \vdash_S s_2 : \mathbb{C}'_i)^{i \in 1..n}}{\Gamma \vdash_S s_1; s_2 : C[\mathbb{a}_1 \mathbin{\&} \mathbb{C}'_1] \cdots [\mathbb{a}_n \mathbin{\&} \mathbb{C}'_n]} \\
\\
\text{(T-Return)} \\
\frac{\Gamma \vdash e : \mathbb{O} \quad \Gamma \vdash \text{destiny} : \mathbb{O}' \\
\mathbb{O} \notin S}{\Gamma \vdash_S \text{return } e : 0 \triangleright \Gamma[\mathbb{O}' \mapsto \Gamma(\mathbb{O})]}
\end{array}$$

Figure 4: Type rules for expressions and statements.

Rule (T-GET) defines the synchronisation with a method invocation that corresponds to a future f . Let $(\mathbb{O}, \alpha, F\mathbb{t}, R)$ be the value of f in the environment. There are two cases: either (i) $R \neq \emptyset$ or (ii) $R = \emptyset$. In case (i), by Restriction 3, $F\mathbb{t}$, which is the value of the caller α when the method has been invoked, is equal to the value $\Gamma(\alpha)$ (and $F = \emptyset$). In this case, if the method returns a new machine (cf. Restriction 4), its state must be $\emptyset\mathbb{t}$ (notice that $\mathbb{t} \sqcap \Gamma(\alpha) \Downarrow = \mathbb{t}$). The case (ii) is more problematic because the caller may be released by a method running in parallel and, therefore, the possible returned virtual machine must record this information in its state. Let $F\mathbb{t}$ be the state of the caller (which is recorded in $\Gamma(f)$). We use the operation $\mathbb{t} \sqcap \Gamma(\alpha) \Downarrow$ to this purpose, which means that the returned machine gets the same state of the carrier α if no method is releasing α , otherwise its state is either $\emptyset\partial$ or $\emptyset\alpha \downarrow$, according to the state of α was $F\top$ or $F\alpha$.

Rule (T-RELEASE) models the removal of a vm name α . The premise in the second line verifies that the disposal do not address machines that are executing methods, as discussed in Restriction 3 of Section 3.

Rule (T-RETURN) is a bit cryptic. First of all it applies provided $\mathbb{O} \notin S$. In fact, by Restriction 4, the type of e in `return e` can be either $_$ or a virtual machine that has been created by the method. In both cases, these values do not belong to S , which is the set of virtual machines in method's arguments – see rule (T-METHOD). In particular, when the type of e is $_$, by (T-METHOD), $\Gamma(\text{destiny}) = _$ and, by definition, $\Gamma[\mathbb{O}' \mapsto \Gamma(\mathbb{O})] = \Gamma$.

The type system of `vm1` is completed with the rules for method declarations and programs, given in Figure 5. Without loss of generality, rule (T-METHOD) assumes that formal parameters of methods are ordered: those of `Int` type occur before those of `Vm` type. We observe that the environment typing the method body binds integer parameters to their same name, while the other ones are bound to fresh vm names (this lets us to have a more precise cost analysis in Section 5). We also observe that the returned value \mathbb{O} may be either $_$ or a fresh vm name ($\mathbb{O} \notin \{\alpha\} \cup \beta$) as discussed in Restriction 4 of Section 3. The constraints in the third line of the premises of (T-METHOD) implement Restriction 1 of Section 3. We also observe that $(\Gamma_i(\gamma) = \Gamma_j(\gamma))^{i,j \in 1..n, \gamma \in S \cup f\nu(\mathbb{O})}$ guarantees that every branch of the behavioural type creates a new vm name and, by rule (T-RETURN), the state of the chosen vm name must be always the same.

We display behavioural types examples by using codes from Sections 1 and 2. Actually, the following types do not abstract a lot from codes because the programs of the previous sections have been designed for highlighting the

$$\begin{array}{c}
\text{(T-METHOD)} \\
\frac{\Gamma(\mathbf{m}) = \alpha(\bar{x}, \bar{\beta}) : \mathbb{O}, \mathbf{R} \quad S = \{\alpha\} \cup \bar{\beta} \quad \mathbb{O} \notin S \quad \Gamma[\text{this} \mapsto \alpha][\text{destiny} \mapsto \mathbb{O}][\bar{x} \mapsto \bar{x}][\bar{z} \mapsto \bar{\beta}][\alpha \mapsto \emptyset \alpha][\bar{\beta} \mapsto \emptyset \bar{\beta}] \vdash_S s : C[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n] \quad \left(\Gamma_i(\gamma) = \Gamma_j(\gamma) \right)_{i,j \in 1..n, \gamma \in \text{Sufv}(\mathbb{O})} \quad \mathbf{R} = (S \cup \text{fv}(\mathbb{O})) \cap \{\gamma \mid \Gamma_1(\gamma) = \mathbf{F}\perp\}}{\Gamma \vdash T \mathbf{m} (\overline{\text{Int } x}, \overline{\text{Vm } z})(\overline{F y} ; s) : \mathbf{m} \alpha(\bar{x}, \bar{\beta}) \{C[\mathbf{a}_1 \triangleright \emptyset] \cdots [\mathbf{a}_n \triangleright \emptyset]\} : \mathbb{O}, \mathbf{R}} \\
\\
\text{(T-PROGRAM)} \\
\frac{\Gamma \vdash \overline{M} : \overline{\mathbb{C}} \quad \Gamma[\text{this} \mapsto \text{start}] \vdash_{\text{start}} s : C[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n]}{\Gamma \vdash \overline{M} (\overline{F x} ; s) : \overline{\mathbb{C}}, C[\mathbf{a}_1 \triangleright \emptyset] \cdots [\mathbf{a}_n \triangleright \emptyset]}
\end{array}$$

Figure 5: Behavioural typing rules of method and programs.

issues of our technique. The following listing shows the behavioural type of `double_release` and the step by step rule application for its calculation.

<pre>Int double_release(Vm x, Vm y) { release x; release y; return 0; }</pre>	<pre>double_release α(β, γ) { // instance of T-Method β✓ § // instance of T-Release γ✓ § // instance of T-Release and T-Seq 0 // instance of T-Return and T-Seq } - , {β, γ}</pre>
---	--

The behavioural types of (`fact` and `costly_fact` examples from Section 2 are the following (we remove trailing 0 when irrelevant):

<pre>fact α(n) { (n==0){ 0 } +(n>0){ vx: fact α(n-1) § x✓ } } - , { }</pre>	<pre>costly_fact α(n) { (n==0){ 0 } +(n>0){ vβ § vx: costly_fact β(n-1) § x✓ § β✓ } } - , { }</pre>
--	--

we notice that the type of `costly_fact` records the order between the recursive invocation and the release of the machine. In the case of the behavioural type of `double_release` the key point is that the releases β^\checkmark and γ^\checkmark in `double_release` are conditioned by the values of β and γ when the method is invoked, we keep track of this with the “effects set” $\{\beta, \gamma\}$.

5. The analysis of behavioural types

The types returned by the system in Section 4 are used to compute the resource cost of a `vm1` program. This computation is performed by a solver of *cost equations*. These cost equations are terms

$$m(\bar{x}) = \text{exp} \quad [se]$$

where m is a (cost) function symbol, exp is an expression that may contain (cost) function symbols applications, and se is a size expression whose variables are contained in \bar{x} (the syntax of exp and se is given in full detail in [13]).

Basically, our translation maps method types into cost equations, where

- method invocations are translated into function applications,
- virtual machine creations are translated into a +1 cost,
- virtual machine releases are translated into a -1 cost,

There are two function calls for every method invocation: one returns the maximal number of resources needed to execute a method m , called *peak cost* of m and noted m_{peak} , and the other returns the number of resources the method m creates without releasing, called *net cost* of m and noted m_{net} . These functions are used to define the cost of sequential execution and parallel execution of methods. For example, omitting arguments of methods, the peak cost of the sequential composition of two methods m and m' is the maximal value between m_{peak} and $m_{\text{net}} + m'_{\text{peak}}$; while the cost of the parallel execution of m and m' is $m_{\text{peak}} + m'_{\text{peak}}$.

There are two difficulties that entangle our translation, both related to method invocations: the management of arguments' identities and the management of arguments' values.

Arguments' identities. Consider the following code and the corresponding behavioural types

```
Int simple_release(VM x) { release x ; return 0 ; }
```

```
Int m(VM x, VM y) {  
  Fut<Int> f; f = this!simple_release(x); release y; f.get;  
  return 0;  
}
```

```
simple_release  $\alpha(\beta) \{ \beta^\vee \} \_ , \{ \beta \}$ 
```

```
 $m \alpha(\beta, \gamma) \{ \nu f : \text{simple\_release } \alpha(\beta) \ ; \ \gamma^\vee \ ; \ f^\vee \} \_ , \{ \beta, \gamma \}$ 
```

We notice that, in the type of m , there is not enough information to determine whether γ^\vee will have a cost equal to -1 or 0 . In fact, in the rule (T-METHOD), we assumed that the arguments were pairwise different. However, this is not the case for invocations. For instance, if m is invoked with two arguments that are equal $-\beta = \gamma$ – then γ is going to be released by the invocation $\text{free}(\beta)$ and therefore it counts 0 . We solve this problem of arguments' identity in the analysis of behavioural types, in particular in the translation of method types. Namely, the above method m is translated in four cost functions: $m_{\text{peak}}^{\{1\},\{2\}}(x, y)$ and $m_{\text{net}}^{\{1\},\{2\}}(x, y)$, which correspond to the invocations where $x \neq y$, and $m_{\text{peak}}^{\{1,2\}}(x)$ and $m_{\text{net}}^{\{1,2\}}(x)$, which correspond to the invocations where $x = y$. (The equivalence relation in the superscript never mention this , which is also an argument, because, in this case this cannot be identified with the other arguments, see below.) Then the translation of an invocation to a method m redirects the invocation to m^Ξ , where Ξ is the equivalence relation expressing the identity of the arguments.

The function computing the equivalence relation of the arguments of an invocation is EqRel . EqRel takes a tuple of vm names and returns a partition of the indices of the tuple:

$$\text{EqRel}(\alpha_0, \dots, \alpha_n) = \bigcup_{i \in 0..n} \{ \{j \mid \alpha_j = \alpha_i\} \}$$

That is, if two indices are in the same set then the corresponding elements in the tuple are equal. So, for instance, $\text{EqRel}(\alpha, \beta, \alpha) = \{\{0, 2\}, \{1\}\}$. We notice that, the possible outputs of $\text{EqRel}(\alpha_0, \alpha_1, \alpha_2)$, without knowing the identities of α_0 , α_1 , and α_2 , are

$\{\{0\}\{1\}, \{2\}\}$	the three arguments are pairwise different
$\{\{0, 1\}, \{2\}\}$	the first and second argument are equal, the third is different
$\{\{0, 2\}, \{1\}\}$	the first and third argument are equal, the second is different
$\{\{0\}, \{1, 2\}\}$	the second and third argument are equal, the first is different
$\{\{0, 1, 2\}\}$	all the arguments are equal

Henceforth, 10 cost functions will correspond to a method with three arguments, 5 for its peak cost and 5 for its net cost.

Next, we also notice that the actual arguments of a $m^{\{0,1\},\{2\}}$ are not three but two: in this case the second argument is useless. Therefore we need a notation for selecting the two relevant arguments in these cases. We use the notation (perhaps awkward) $\text{EqRel}(\alpha, \beta, \alpha, \gamma)(\alpha, \beta, \alpha, \gamma)$ that actually returns the triple (α, β, γ) .

Without loosing in generality, we will always assume that the canonical representative of a set containing 0 is always 0 . This index represents the `this` object and we remind that, by Simplification 3 in Section 3, such an object cannot be released. This is the reason why, in the foregoing discussion about the method m , we have not mentioned `this`. Additionally, in order to simplify the translation of method invocations, we also assume that the argument `this` is always different from other arguments (the general case just requires more details).

(Re)computing argument's states. In the rules of Figure 4, in order to enforce the restrictions in Section 3, we have already computed the state of every machine. In this section we recompute them for a different reason: obtaining a (more) precise cost analysis. Of course one might record the computation of vm states in behavioural types. However, this solution has the drawback that behavioural types become unintelligible because they carry information that is needed at a later stage by the analyser.

Let a *translation environment*, ranged over Ψ, Ψ' , be a mapping from vm names to vm states and from future names to triples $(\Psi', R, m \beta(\overline{se}, \overline{\beta}) \rightarrow \oplus)$. The translation environment Ψ' in $(\Psi', R, m \beta(\overline{se}, \overline{\beta}) \rightarrow \oplus)$ is only defined on vm names. For this reason, we call it *vm-translation environment*. We define the following auxiliary functions

- let Ψ be a vm-translation environment. Then

$$\Psi|_X(\alpha) \stackrel{def}{=} \begin{cases} \Psi(\alpha) & \text{if } \alpha \in X \\ \text{undefined} & \text{otherwise} \end{cases}$$

- the *update of a vm-translation environment* Ψ with respect to f and Ψ' , written $\Psi \searrow \Psi'$, returns a vm-translation environment defined as follows:

$$(\Psi \searrow \Psi')(\alpha) \stackrel{def}{=} \begin{cases} (F' \setminus \{f\})(\mathbb{t} \sqcap \mathbb{t}') & \text{if } \Psi(\alpha) = F\mathbb{t} \text{ and } \Psi'(\alpha) = F'\mathbb{t}' \\ \text{undefined} & \text{otherwise} \end{cases}$$

This operation $\Psi \searrow \Psi'$ updates the vm-translation environment Ψ that is stored in the future f with the translation environment at the synchronisation point. It is worth to observe that, by the definition of our type system and the following translation function, the values of $\Psi(\alpha)$ and $\Psi'(\alpha)$ are related. In particular, if $\mathbb{t} = \alpha$ then \mathbb{t}' can be either α or $\alpha\downarrow$ (the machine is released by a method that has been invoked in parallel) or \perp (the machine has been released before the `get` operation on the future f); if $\mathbb{t} = \top$ then \mathbb{t}' can be either \top or ∂ (the machine is released by a method that has been invoked in parallel) or \perp (the machine has been released before the `get` operation).

- the *merge operation*, noted $\Psi(\Delta)$, where Ψ is a vm-translation environment and Δ is an equivalence relation, returns a *substitution* defined as follows. Let

$$\mathbb{t} \otimes^\alpha \mathbb{t}' \stackrel{def}{=} \begin{cases} \perp & \text{if } \mathbb{t} = \perp \text{ or } \mathbb{t}' = \perp \\ \alpha & \text{if } \mathbb{t} \text{ and } \mathbb{t}' \text{ are variables} \\ \alpha\downarrow & \text{otherwise} \end{cases} \quad F_1\mathbb{t}_1 \otimes^\alpha F_2\mathbb{t}_2 \stackrel{def}{=} (F_1 \cup F_2)(\mathbb{t}_1 \otimes^\alpha \mathbb{t}_2)$$

Then

$$\Psi(\Delta) : \alpha \mapsto \bigotimes^{\Delta(\alpha)} \{\Psi(\beta) \mid \beta \in \text{dom}(\Psi) \text{ and } \Delta(\beta) = \Delta(\alpha)\}$$

for every $\alpha \in \text{dom}(\Psi)$.

The operator \otimes^α has not been defined on vm values as ∂ or \top because we merge vm names whose image by Ψ are either $F\beta$ or $F\beta\downarrow$ or $F\perp$. As a notational remark, we observe that $\Psi(\Delta)$ is noted as a map $[\alpha_1 \mapsto F_1\mathbb{t}_1, \dots, \alpha_n \mapsto F_n\mathbb{t}_n]$ instead of the standard notation $[F_1\mathbb{t}_1, \dots, F_n\mathbb{t}_n / \alpha_1, \dots, \alpha_n]$. These two notations are clearly equivalent: we prefer the former one because it will let us to write $\Psi(\Delta)(\alpha)$ or even $\Psi(\Delta)(\alpha_1, \dots, \alpha_n)$ with the obvious meanings.

To clarify the reason for a merge operator, consider the atom f^\vee within a behavioural type that binds f to `foo` $\alpha(\beta, \gamma)$. Assume to evaluate this type with $\Delta = \{\beta, \gamma\}$. That is, the two arguments are actually identical. Which are the values of β and γ for evaluating `foo`_{peak} $^\Delta$ and `foo`_{net} $^\Delta$? Well, we have

1. to select the representative between β and γ : it will be $\Delta(\beta)$ – which is equal to $\Delta(\gamma)$;
2. to take a value that is smaller than $\Psi(\beta)$ and $\Psi(\gamma)$ (but greater than any other value that is smaller);
3. to substitute β and γ with the result of 2.

For instance, let $\Psi = [\alpha \mapsto \emptyset\alpha, \beta \mapsto \emptyset\beta\downarrow, \gamma \mapsto \emptyset\gamma]$ and $\Delta(\beta) = \Delta(\gamma) = \beta$. We expect that a value for the item 2 above is $\emptyset\beta\downarrow$ and the substitution of the item 3 is $[\emptyset\beta\downarrow, \emptyset\beta\downarrow / \beta, \gamma]$. Formally, the operation returning the value for 2 is \otimes^β and the substitution of item 3 is the output of the merge operation.

The translation function. The translation function, called `translate`, is structured in three parts that respectively correspond to simple atoms, behavioural types, and method types and full programs. This function carries five arguments:

1. Δ is the *equivalence relation on formal parameters* identifying those that are equal. We assume that $\Delta(x)$ returns the unique representative of the equivalence class of x . For simplicity we also let $\Delta(x) = x$ for every x that belongs to the local variables. Therefore we can use Δ also as a substitution operation.
2. Ψ is the *translation environment* which stores temporary information about futures that are active (unsynchronised) and about the state of vm names;
3. α is the name of the virtual machine of the current behavioural type;
4. \bar{e} is the sequence of (over-approximated) costs of the current execution branch;
5. the behavioural type being translated; it may be either \mathbf{a} , \mathbf{c} or $\overline{\mathbf{C}}$.

In the definition of `translate` we use the two functions

$$\text{CNEW}(\alpha) = \begin{cases} 0 & \alpha = \perp \\ 1 & \text{otherwise} \end{cases} \quad \text{CREL}(\alpha) = \begin{cases} -1 & \alpha = \top \\ 0 & \text{otherwise} \end{cases}$$

The left-hand side function is used when a virtual machine is created. It returns 1 or 0 according to the virtual machine that is executing the code can be alive ($\alpha \neq \perp$) or not, respectively. The right-hand side function is used when a virtual machine is released (in correspondence of atoms β^\vee). The release is effectively computed – value -1 – only when the virtual machine that is executing the code is alive ($\alpha = \top$).

We will assume the presence of a lookup function `lookup` that takes method invocations $\mathbf{m} \alpha(\bar{\tau}, \bar{\beta})$ and returns tuples $\mathbf{c} : \mathbf{O}, \mathbf{R}$. This function is left unspecified. We also write $\mathbf{R}[\mathbf{O}'/\mathbf{O}]$ to denote a set that is equal to \mathbf{R} if $\mathbf{O} \notin \mathbf{R}$, that is equal to $(\mathbf{R} \setminus \{\mathbf{O}\}) \cup \{\mathbf{O}'\}$, otherwise.

The definition of `translate` follows. We begin with the translation of atoms.

$$\text{translate}[\Delta, \Psi, \alpha](\bar{e}; e)(\mathbf{a}) = \begin{cases} (\Psi, \bar{e}; e) & \text{when } \mathbf{a} = 0 \\ (\Psi[\beta \mapsto \top], \bar{e}; e; e + \text{CNEW}(\mathbf{t})) & \text{when } \mathbf{a} = \nu\beta \text{ and } \Psi(\alpha) = \mathbf{Ft} \\ (\Psi[\Delta(\beta) \mapsto \top], \bar{e}; e; e + \text{CREL}(\mathbf{t})) & \text{when } \mathbf{a} = \beta^\vee \text{ and } \Psi(\Delta(\beta)) = \mathbf{Ft} \\ (\Psi'[f \mapsto (\Psi|_{\Delta(\beta, \bar{\beta})}, \mathbf{R}, \mathbf{m} \Delta(\beta)(\bar{\tau}, \Delta(\bar{\beta}) \rightarrow \mathbf{O}))], \bar{e}; e; e + f) & \begin{array}{l} \text{when } \mathbf{a} = \nu f : \mathbf{m} \beta(\bar{\tau}, \bar{\beta}) \\ \text{and } \Psi' = \Psi[\beta \mapsto (\mathbf{F} \cup \{f\})] \text{ and } \Theta(\mathbf{EqRel}(\beta, \bar{\beta})) = (\mathbf{F}_1 \mathbf{t}_1, \dots, \mathbf{F}_n \mathbf{t}_n) \\ \text{and } \text{lookup}(\mathbf{m} \Delta(\beta)(\bar{\tau}, \Delta(\bar{\beta}))) = \mathbf{c} : \mathbf{O}, \mathbf{R} \end{array} \\ (\Psi'' \setminus f, (\bar{e}; e)\sigma'; (e)\sigma' + \sum_{\gamma \in \Delta(\mathbf{R}), \Theta(\gamma) = \mathbf{Ft}, \mathbf{F} \neq \emptyset} \text{CREL}(\mathbf{t})) & \begin{array}{l} \text{when } \mathbf{a} = f_{\mathbf{O}'}^\vee \text{ and } \Psi(f) = (\Psi', \mathbf{R}, \mathbf{m} \beta(\bar{\tau}, \bar{\beta}) \rightarrow \mathbf{O}) \\ \text{and } \Theta = \Psi' \searrow \Psi \text{ and } \Theta(\mathbf{EqRel}(\beta, \bar{\beta})) = (\mathbf{F}_1 \mathbf{t}_1, \dots, \mathbf{F}_n \mathbf{t}_n) \\ \text{and } \sigma = [\mathbf{m}_{\text{peak}}^{\bar{\tau}}(\bar{\tau}, \mathbf{F}_1 \mathbf{t}_1 \Downarrow, \dots, \mathbf{F}_n \mathbf{t}_n \Downarrow) / f] \\ \text{and } \sigma' = [\mathbf{m}_{\text{net}}^{\bar{\tau}}(\bar{\tau}, \mathbf{F}_1 \mathbf{t}_1 \Downarrow, \dots, \mathbf{F}_n \mathbf{t}_n \Downarrow) / f] \\ \text{and } \mathbf{R}' = \mathbf{R}[\mathbf{O}'/\mathbf{O}] \\ \text{and } \Psi'' = \Psi[\gamma \mapsto \top] \text{ and } \Psi'[\gamma' \mapsto \Theta(\beta)] \text{ and } \gamma' \in \text{fv}(\mathbf{O}') \setminus \mathbf{R}' \end{array} \end{cases}$$

In the definition of `translate` we always highlight the last expression in the sequence of costs of the current execution branch (the fourth input). This is because the cost of the parsed atom applies to it, except for the case of f^\vee . In this last case, let $\bar{e}; e$ be the expression. Since the atom expresses the synchronisation of f , $\bar{e}; e$ will have occurrences of f . In this case, the function `translate` has to compute two values: the maximum number of resources used by (the method corresponding to) f during its execution – the *peak cost* used in the substitution σ – and the resources used upon the termination of (the method corresponding to) f – the *net cost* used in the substitution σ' . In particular, this last value has to be decreased by the number of resources released by the method. This is the purpose of the addend $\sum_{\gamma \in \Delta(\mathbf{R}), \Theta(\gamma) = \mathbf{Ft}, \mathbf{F} \neq \emptyset} \text{CREL}(\mathbf{t})$ that removes machines that are going to be removed by parallel methods (the constraint $\mathbf{F} \neq \emptyset$) because the other ones have been already counted both in the peak cost and in the net cost. We observe that the instances of the method \mathbf{m}_{peak} and \mathbf{m}_{net} that are invoked are those corresponding to the equivalence relation of the tuple $(\beta, \bar{\beta})$.

The translation of behavioural types is given by composing the definitions of the atoms. In this case, the output of `translate` is a set of cost equations.

$$\text{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbb{C}) = \begin{cases} \{(se')\{\bar{e}'\}\} & \text{when } \mathbb{C} = \mathbb{a} \triangleright \emptyset \text{ and } \text{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbb{a}) = (\Psi', (se')\{\bar{e}'\}) \\ C'' & \text{when } \mathbb{C} = \mathbb{a} \mathbin{\&} \mathbb{C}' \text{ and } \text{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbb{a}) = (\Psi', \{(se')\{\bar{e}'\}\}) \\ & \text{and } \text{dom}(\Psi') \setminus \text{dom}(\Psi) = S \text{ and } \text{translate}(\Delta \cup \{S\}, \Psi', \alpha, (se')\{\bar{e}'\}, \mathbb{C}') = (\Psi'', C'') \\ C' \cup C'' & \text{when } \mathbb{C} = \mathbb{C}_1 + \mathbb{C}_2 \text{ and } \text{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbb{C}_1) = (\Psi', C') \\ & \text{and } \text{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbb{C}_2) = (\Psi'', C'') \\ C' & \text{when } \mathbb{C} = (se')\{\mathbb{C}'\} \text{ and } \text{translate}(\Delta, \Psi, \alpha, (se \wedge se')\{\bar{e}\}, \mathbb{C}') = (\Psi', C') \\ C' & \text{when } \mathbb{C} = (e')\{\mathbb{C}'\} \text{ and } e' \text{ contains } _ \text{ and } \text{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbb{C}') = (\Psi', C') \end{cases}$$

The translation of method types and behavioural type programs is given below. Let \mathcal{P} be the set of partitions of $1..n$. Then

$$\text{translate}(\mathbb{M} \alpha_1(\bar{x}, \alpha_2, \dots, \alpha_n) \{\mathbb{C}\} : \mathbb{O}, \mathbb{R}) = \bigcup_{\Xi \in \mathcal{P}} \text{translate}(\Xi, \mathbb{M} \alpha_1(\bar{x}, \alpha_2, \dots, \alpha_n) \{\mathbb{C}\} : \mathbb{O}, \mathbb{R})$$

where $\text{translate}(\Xi, \mathbb{M} \alpha_1(\bar{x}, \alpha_2, \dots, \alpha_n) \{\mathbb{C}\} : \mathbb{O}, \mathbb{R})$ is defined as follows. Let

$$\Delta = \{ \{ \alpha_{i_1}, \dots, \alpha_{i_m} \} \mid \{ i_1, \dots, i_m \} \in \Xi \} \quad \text{and} \quad [\Delta] = \{ \alpha \mapsto \emptyset \alpha \mid \alpha = \Delta(\alpha) \}$$

$$\text{and} \quad \text{translate}(\Delta, [\Delta], \alpha_1)(0)(\mathbb{C}) = \bigcup_{i=1}^n (se_i)\{e_{1,i}; \dots; e_{h_i,i}\}.$$

Then

$$\text{translate}(\Xi, \mathbb{M} \alpha_1(\bar{x}, \alpha_2, \dots, \alpha_k) \{\mathbb{C}\} : \mathbb{O}, \mathbb{R}) = \begin{cases} \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = 0 & [\alpha_1 = \perp] \\ \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{1,1} & [se_1 \wedge \alpha_1 \neq \perp] \\ \vdots & \\ \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_1,1} & [se_1 \wedge \alpha_1 \neq \perp] \\ \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{1,2} & [se_2 \wedge \alpha_1 \neq \perp] \\ \vdots & \\ \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_n,n} & [se_n \wedge \alpha_1 \neq \perp] \\ \mathbb{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = 0 & [\alpha_1 = \perp] \\ \mathbb{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \dots, \alpha_n]) & [\alpha_1 = \partial] \\ \mathbb{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_1,1} & [se_1 \wedge \alpha_1 = \top] \\ \vdots & \\ \mathbb{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_n,n} & [se_n \wedge \alpha_1 = \top] \end{cases}$$

Let $(\mathbb{C}_1 \dots \mathbb{C}_n, \mathbb{C})$ be a behavioural type program and let $\text{translate}(\emptyset, \emptyset, \alpha, (\text{true})\{0\}, \mathbb{C}) = \bigcup_{j=1}^m (se_j)\{e_{1,j}; \dots; e_{h_j,j}\}$. Then

$$\text{translate}(\mathbb{C}_1 \dots \mathbb{C}_n, \mathbb{C}) = \begin{cases} \text{translate}(\mathbb{C}_1) \dots \text{translate}(\mathbb{C}_n) \\ \text{main}() = 1 + e_{1,1} & [se_1] \\ \vdots & \\ \text{main}() = 1 + e_{h_1,1} & [se_1] \\ \text{main}() = 1 + e_{1,2} & [se_2] \\ \vdots & \\ \text{main}() = 1 + e_{h_m,m} & [se_m] \end{cases}$$

As an example, we show the output of `translate` when applied to the behavioural type of `double_release` computed in Section 4. Since `double_release` has two arguments, we generate two sets of equations, as discussed above. In order to ease the reading, we omit the equivalence classes of arguments that label function names: the reader can grasp them from the number of arguments. For the same reason, we represent a partition $\{\{1\}, \{2\}, \{3\}\}$ corresponding to vm names α_1, α_2 and α_3 by $[\alpha_1, \alpha_2, \alpha_3]$ and $\{\{1\}, \{2, 3\}\}$ by $[\alpha_1, \alpha_2]$ (we write the canonical representatives). For simplicity we do not add the partition to the name of the method.

$$\begin{aligned}
\text{translate}([\alpha_1, \alpha_2, \alpha_3], \text{double_release } \alpha_1(\alpha_2, \alpha_3) \{ \mathbb{C} \} : _, \{ \alpha_2, \alpha_3 \}) = & \quad \text{translate}([\alpha_1, \alpha_2], \text{double_release } \alpha_1(\alpha_2) \{ \mathbb{C} \} : _, \{ \alpha_2 \}) = \\
\left\{ \begin{array}{ll} \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2, \alpha_3) = 0 & [\alpha_1 = \perp] \\ \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2, \alpha_3) = 0 & [\alpha_1 \neq \perp] \\ \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2, \alpha_3) = \text{CREL}(\alpha_2) & [\alpha_1 \neq \perp] \\ \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2, \alpha_3) = \text{CREL}(\alpha_2) + \text{CREL}(\alpha_3) & [\alpha_1 \neq \perp] \\ \\ \text{double_release}_{\text{net}}(\alpha_1, \alpha_2, \alpha_3) = 0 & [\alpha_1 = \perp] \\ \text{double_release}_{\text{net}}(\alpha_1, \alpha_2, \alpha_3) = \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2, \alpha_3) & [\alpha_1 = \partial] \\ \text{double_release}_{\text{net}}(\alpha_1, \alpha_2, \alpha_3) = \text{CREL}(\alpha_2) + \text{CREL}(\alpha_3) & [\alpha_1 = \top] \end{array} \right. & \quad \left\{ \begin{array}{ll} \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2) = 0 & [\alpha_1 = \perp] \\ \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2) = 0 & [\alpha_1 \neq \perp] \\ \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2) = \text{CREL}(\alpha_2) & [\alpha_1 \neq \perp] \\ \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2) = \text{CREL}(\alpha_2) + \text{CREL}(\perp) & [\alpha_1 \neq \perp] \\ \\ \text{double_release}_{\text{net}}(\alpha_1, \alpha_2) = 0 & [\alpha_1 = \perp] \\ \text{double_release}_{\text{net}}(\alpha_1, \alpha_2) = \text{double_release}_{\text{peak}}(\alpha_1, \alpha_2) & [\alpha_1 = \partial] \\ \text{double_release}_{\text{net}}(\alpha_1, \alpha_2) = \text{CREL}(\alpha_2) + \text{CREL}(\perp) & [\alpha_1 = \top] \end{array} \right.
\end{aligned}$$

To highlight a cost computation concerning `double_release`, consider the following two potential users and the corresponding behavioural types

```

Int user1() {
  Vm x ; Vm y ; Fut<Int> f ;
  x = new Vm ; y = new Vm ;
  f = this!double_release(x, y) ;
  f.get ; return 0 ;
}

```

```

Int user2() {
  Vm x ; Fut<Int> f ;
  Vm x = new Vm ;
  f = this!double_release(x, x) ;
  f.get ; return 0 ;
}

```

```

user1 α( ) {
  νβ § νγ § νf : double_release α(β, γ) § f✓
} - , { }

user2 α( ) {
  νβ § νf : double_release α(β, β) § f✓
} - , { }

```

The translations of the foregoing types give the following set of equations

$$\begin{aligned}
\text{translate}([\alpha_1], \text{user1 } \alpha_1() \{ \mathbb{C}_{\text{user1}} \} : _, \{ \}) = & \quad \text{translate}([\alpha_1], \text{user2 } \alpha_1() \{ \mathbb{C}_{\text{user2}} \} : _, \{ \}) = \\
\left\{ \begin{array}{ll} \text{user1}_{\text{peak}}(\alpha_1) = 0 & [\alpha_1 = \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = 0 & [\alpha_1 \neq \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) & [\alpha_1 \neq \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{CNEW}(\alpha_1) & [\alpha_1 \neq \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{CNEW}(\alpha_1) & + \text{double_release}_{\text{peak}}(\alpha_1, \top, \top) \quad [\alpha_1 \neq \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{CNEW}(\alpha_1) & + \text{double_release}_{\text{net}}(\alpha_1, \top, \top) \quad [\alpha_1 \neq \perp] \\ \\ \text{user1}_{\text{net}}(\alpha_1) = 0 & [\alpha_1 = \perp] \\ \text{user1}_{\text{net}}(\alpha_1) = \text{user1}_{\text{peak}}(\alpha_1) & [\alpha_1 = \partial] \\ \text{user1}_{\text{net}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{CNEW}(\alpha_1) & + \text{double_release}_{\text{net}}(\alpha_1, \top, \top) \quad [\alpha_1 = \top] \end{array} \right. & \quad \left\{ \begin{array}{ll} \text{user2}_{\text{peak}}(\alpha_1) = 0 & [\alpha_1 = \perp] \\ \text{user2}_{\text{peak}}(\alpha_1) = 0 & [\alpha_1 \neq \perp] \\ \text{user2}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) & [\alpha_1 \neq \perp] \\ \text{user2}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{double_release}_{\text{peak}}(\alpha_1, \top) & [\alpha_1 \neq \perp] \\ \text{user2}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{double_release}_{\text{net}}(\alpha_1, \top) & [\alpha_1 \neq \perp] \\ \\ \text{user2}_{\text{net}}(\alpha_1) = 0 & [\alpha_1 = \perp] \\ \text{user2}_{\text{net}}(\alpha_1) = \text{user2}_{\text{peak}}(\alpha_1) & [\alpha_1 = \partial] \\ \text{user2}_{\text{net}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{double_release}_{\text{net}}(\alpha_1, \top) & [\alpha_1 = \top] \end{array} \right.
\end{aligned}$$

If we compute the cost of $\text{user1}_{\text{peak}}(\alpha)$ and $\text{user1}_{\text{net}}(\alpha)$ we obtain 2 and 0, respectively. That is, in this case, `double_release` being invoked with two different arguments has cost -2. On the contrary, the cost of $\text{user2}_{\text{peak}}(\alpha)$ and $\text{user2}_{\text{net}}(\alpha)$ is 1 and 0, respectively. That is, in this case, `double_release` being invoked with two equal arguments has cost -1.

Theorem 5.1 (Correctness). *Let $\overline{M} \{ \overline{Fz} ; s' \}$ be a well-typed program and let $\overline{\mathbb{C}}, \mathbb{C}$ be its behavioural type and cn be its initial configuration. Let also n be a solution of the function $\text{translate}(\overline{\mathbb{C}}, \mathbb{C})$. Then, for every $cn \xrightarrow{*} cn'$, $\text{alive}(cn') \leq n$.*

The proof of correctness of our technique is long, even if almost standard (see [17] for a similar proof). It consists of two parts:

$$\begin{array}{c}
\text{(INF-METHOD)} \\
\frac{\Gamma(\mathbf{m}) = \alpha(\bar{x}, \bar{\beta}) : \emptyset, \mathbf{R} \quad S = \{\alpha\} \cup \bar{\beta} \quad \emptyset \notin S \quad \Gamma[\text{this} \mapsto \alpha][\text{destiny} \mapsto \emptyset][\bar{x} \mapsto \bar{x}][\bar{z} \mapsto \bar{\beta}][\alpha \mapsto \emptyset][\bar{\beta} \mapsto \emptyset] \vdash_S s : C[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n]}{\Gamma \vdash T \mathbf{m} (\overline{\text{Int } x, \forall \mathbf{m} z}) \{F y ; s\} : \mathbf{m} \alpha(\bar{x}, \bar{\beta}) \{C[\mathbf{a}_1 \triangleright \emptyset] \cdots [\mathbf{a}_n \triangleright \emptyset]\} : \emptyset, \mathbf{R}' ; \Gamma[\mathbf{m} \mapsto \alpha(\bar{x}, \bar{\beta}) : \emptyset, \mathbf{R}']} \\
\text{(INF-PROGRAM)} \\
\frac{\Gamma \vdash \overline{M} : \overline{C}, \Gamma \quad \Gamma[\text{this} \mapsto \text{start}] \vdash_{\text{start}} s : C[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n]}{\Gamma \vdash \overline{M} \{F x ; s\} : \overline{C}, C[\mathbf{a}_1 \triangleright \emptyset] \cdots [\mathbf{a}_n \triangleright \emptyset]}
\end{array}$$

Figure 6: Behavioural inference rules of method and programs.

1. the first part addresses the correctness of the type system in Section 4. As usual with type systems, the correctness is represented by a subject reduction theorem expressing that if a configuration cn of the operational semantics is well typed and $cn \rightarrow cn'$ then cn' is well-typed as well.
2. the second part uses a notion of *direct cost of a behavioural type* (of a configuration), which is the number of virtual machines occurring in the type plus an over-approximation of those that will be created, and demonstrate that the direct cost never increases during the evolution of a system.

The details are reported in Appendix A.

6. The SRA tool

The analysis technique presented in this article has been prototyped [14]. In this section we discuss a number of technical details about the type system and about the translation of the cost equations in Section 5 into an input that is adequate to an off-the-shelf solver – the CoFloCo solver [13]. We also deliver a preliminary assessment of our prototype by discussing a number of examples.

6.1. Type inference

The system defined in Section 4 is a type-checking process, that is it compels the user to write the behavioural types of methods. In order to relieve users from writing verbose annotations, it is possible to turn the type checking system into a type inference. In fact, in the case of `vm1`, this adaptation is not difficult because the only problematic issue is the inference of method's effects \mathbf{R} . This issue is easily solved by changing the rules (T-METHOD) and (T-PROGRAM) in Figure 5 into those of Figure 6. In particular, rule (INF-METHOD) uses a new judgment $\Gamma \vdash T \mathbf{m} (\overline{\text{Int } x, \forall \mathbf{m} z}) \{F y ; s\} : \mathbf{m} \alpha(\bar{x}, \bar{\beta}) \{C\} : \emptyset, \mathbf{R}' ; \Gamma'$, which also returns a new environment Γ' and where \mathbf{R}' may be different from the method's effects stored in $\Gamma(\mathbf{m})$. The main difference between (T-METHOD) and (INF-METHOD) is that the latter one lets Γ' record method's effects computed in the premise. The key constraint in Figure 6 is the left premise of (INF-PROGRAM). This premise imposes method definitions to be typed with the *same* input and output environment. That is, the environment Γ is a *fixpoint* of the rule.

In order to compute the fixpoint Γ in the rule (INF-PROGRAM), we notice that the differences between Γ and Γ' in $\Gamma \vdash \overline{M} : \overline{C}, \Gamma'$, if any, are in methods' effects. In addition, these method's effects are all *finite* because they are always a subset of method's arguments. Therefore, there is a standard technique for computing the *least fixpoint* environment:

1. verify that the transformation $\Gamma \vdash \overline{M} : \overline{C}, \Gamma'$ is monotone with respect to the subset ordering;
2. let Γ_0 be such that $\Gamma_0(\mathbf{m}) = \alpha(\bar{x}, \bar{\beta}) : \emptyset, \emptyset$ and let $\Gamma_i \vdash \overline{M} : \overline{C}, \Gamma_{i+1}$. Then there exists a least n such that $\Gamma_n \vdash \overline{M} : \overline{C}, \Gamma_n$.

It is easy to verify that 1 holds. Henceforth, the correctness of [14]. For example, let

$\text{m1 } \alpha(\beta, \gamma, \delta) \{ \begin{array}{l} \nu f : \text{m2 } \alpha(\beta) \ ; \ f^\vee \ ; \ \gamma^\vee \\ \} - , \{ \} \end{array}$	$\text{m2 } \alpha(\beta) \{ \begin{array}{l} \nu f : \text{m3 } \alpha(\beta) \ ; \ f^\vee \\ \} - , \{ \} \end{array}$	$\text{m3 } \alpha(\beta) \{ \begin{array}{l} \beta^\vee \\ \} - , \{ \} \end{array}$
--	--	---

Then, the reader can verify that the computation of the least fixpoint environment gives the following sequence of environments:

- step 0: $\Gamma_0 = [m1 \mapsto \alpha(\beta, \gamma, \delta) : -, \emptyset, m2 \mapsto \alpha(\beta) : -, \emptyset, m3 \mapsto \alpha(\beta) : -, \emptyset]$;
- step 1: $\Gamma_1 = [m1 \mapsto \alpha(\beta, \gamma, \delta) : -, \{\gamma\}, m2 \mapsto \alpha(\beta) : -, \emptyset, m3 \mapsto \alpha(\beta) : -, \{\beta\}]$;
- step 2: $\Gamma_2 = [m1 \mapsto \alpha(\beta, \gamma, \delta) : -, \{\gamma\}, m2 \mapsto \alpha(\beta) : -, \{\beta\}, m3 \mapsto \alpha(\beta) : -, \{\beta\}]$;
- step 3: $\Gamma_3 = [m1 \mapsto \alpha(\beta, \gamma, \delta) : -, \{\beta, \gamma\}, m2 \mapsto \alpha(\beta) : -, \{\beta\}, m3 \mapsto \alpha(\beta) : -, \{\beta\}]$.

6.2. The translation of cost equations

To comply with CoFloCo input formats, we need to encode the vm values and the functions CNEW and CREL into integer values and cost equations, respectively. Let

- \top be encoded by 1, ∂ be encoded by 2, and \perp be encoded by 3. The vm value $\alpha \downarrow$ is encoded by the conditional value $[\alpha = 3]3 + [1 \leq \alpha \leq 2]2$;
- CNEW and CREL equations are encoded as follows and are fixed for all programs:

```
eq(CNEW(A), 0, [], [A = 3]).
eq(CNEW(A), 1, [], [A < 3]).
```

```
eq(CREL(A), -1, [], [A = 1]).
eq(CREL(A), 0, [], [A > 1]).
```

In order to illustrate the encoding in CoFloCo of the cost equations in Section 5, we discuss the equations of the two factorial programs `fact` and `cheap_fact` in Section 2. The CoFloCo equations are reported in Figure 7.

```
1 eq(fact_peak(A,B), 0, [], [A = 3]).
2 eq(fact_peak(A,B), 0, [], [B = 0]).
3 eq(fact_peak(A,B), 0, [], [B > 0]).
4 eq(fact_peak(A,B), 0, [fact_peak(1,B)], [B > 0]).
5 eq(fact_peak(A,B), 0, [fact_net(1,B)], [B > 0]).
6
7 eq(fact_net(A,B), 0, [], [A = 3]).
8 eq(fact_net(A,B), 0, [fact_peak(A,B)], [A = 2]).
9 eq(fact_net(A,B), 0, [], [A = 1, B = 0]).
10 eq(fact_net(A,B), 0, [fact_net(1, B-1)], [A = 1, B > 0]).
11
12 eq(costly_fact_peak(A,B), 0, [], [A = 3]).
13 eq(costly_fact_peak(A,B), 0, [], [B = 0]).
14 eq(costly_fact_peak(A,B), 0, [], [B > 0]).
15 eq(costly_fact_peak(A,B), 0, [cnew(A)], [B > 0]).
16 eq(costly_fact_peak(A,B), 0, [cnew(A), costly_fact_peak(1, B-1)], [B > 0]).
17 eq(costly_fact_peak(A,B), 0, [cnew(A), costly_fact_net(1, B-1)], [B > 0]).
18 eq(costly_fact_peak(A,B), 0, [cnew(A), costly_fact_net(1, B-1), crel(1)], [B > 0]).
19
20 eq(costly_fact_net(A,B), 0, [], [A = 3]).
21 eq(costly_fact_net(A,B), 0, [costly_fact_peak(A,B)], [A = 2]).
22 eq(costly_fact_net(A,B), 0, [], [A = 1, B = 0]).
23 eq(costly_fact_net(A,B), 0, [cnew(A), costly_fact_net(1, B-1), crel(1)], [A = 1, B > 0]).
```

Figure 7: Cost equations of `fact` and `costly_fact` in CoFloCo format

These equations do not require any comments because they exactly correspond to the output of the `translate` function in Section 5. We just observe that, in addition to these equations, one needs to specify a so-called *entry point*, which corresponds to the main body. This entry point has the format

`entry(METHOD_NAME(LIST_OF_ARGUMENTS) : [CONDITIONS])`

that is instantiated into `entry(main(MAINVM): [])` for the main body. It is also possible to add ad-hoc entries for particular methods, if the user wants to highlight their cost (in particular, our translator adds entries for peak and net costs of every method). For example, the following table reports the outputs (of CoFloCo) corresponding to the entries in the left column.

Entry Point	Cost
<code>entry(fact_net(1,B): [B>=0]) .</code>	0
<code>entry(fact_peak(1,B): [B>=0]) .</code>	0
<code>entry(costly_fact_net(1,B): [B>=0]) .</code>	0
<code>entry(costly_fact_peak(1,B): [B>=0]) .</code>	$\max([B, 1])$

The *net cost* of `fact` and `costly_fact` is equal to 0 because, in both cases, every created virtual machine is released before the end of the program. For *peak cost*, the number of virtual machines in `fact` is 0. On the other hand, `costly_fact` creates at each step a virtual machine and releases it after the recursive call. This management of virtual machines gives a *peak cost* equal to maximum between 1 and B.

Unbounded recursive method with release operations. Consider the following wrong version of the `fake_method` example from Section 2, where the programmer has erroneously written “ $n + 1$ ” instead of “ $n - 1$ ” in the recursive invocation (henceforth the recursion is unbounded because the base case is never met). As a result one can expect that the number of created virtual machines is infinite because, at every step, only one of the two created machines is released.

```
Int fake_method(Int n) {
  if (n=0) return 0 ;
  else { VM x, y ; Fut<Int> f, g; Int u, v;
    x = new VM() ; y = new VM() ;
    f = this!double_release(x,x) ; u = f.get ;
    g = this!fake_method(n+1) ; v = g.get ;
    return 0 ; }
}
```

The equations generated by the analysis are shown in Figure 8. The results of the analysis of these equations produce a cost of “infinity” for the main method, as suspected. However, let’s suppose the situation in which both vms are released (change `double_release(x,x)` to `double_release(x,y)`). In this case the set of equations of `fake_method` remains almost identical with the only difference that the element `crel(1)` in equations 10, 12, 14 and 18, now appears two times). If we re-run the cost analysis with the new equations we obtain, instead of infinity, a maximum cost of two vms for the main method. Again this is the expected behavior because even though the program runs indefinitely there are no more than two vms (in addition to the main one) co-existing at the same time.

There is an issue about Figure 7 that deserves to be commented. CoFloCo outputs an error when it is fed with that set of equations. This is because instructions 8 and 21 are mutually recursive with instructions 5 and 17-18, respectively, and mutual recursion is banned by the analyser. It is worth to notice that 8 and 21 correspond to invocations of `fact_net` and `costly_fact_net` when the carrier is δ . Intuitively, the idea behind these equations is that when the vm state is unknown the analyzer should assume the worst and consider the highest possible cost (the peak cost) as the net cost. The workaround to solve this problem is to remove the mutual recursion by replacing *in-place* the internal function invocation by its body.

6.3. Assessments – a comparison with C4B and SACO

There are very few tools targeting resource analysis in programs that feature concurrency and negative costs due to explicit release operations. To the best of our knowledge, the most relevant tool in the literature is SACO², which is based on the results of [4] and is able to address a setting similar to ours. We also compare SRA with C4B³ [11], a tool based on amortized analysis and using the results in [19]. Notwithstanding C4B only targets sequential programs, it may be considered as the state-of-the-art of *certified tools* for the automatic static analysis of resource usage bounds.

Some remarks on the tool comparison are in order:

²Web site at <http://ei. abs-models.org:8082/clients/web/>

³Web site at <http://www.cs.yale.edu/homes/qcar/aaa/>

```

1 eq(fakeMethod01peak(THISVM,N), 0, [], [THISVM = 3]).
2 eq(fakeMethod01peak(THISVM,N), 0, [], [THISVM < 3]).
3
4 eq(fakeMethod01net(THISVM,N), 0, [], [THISVM = 3]).
5 eq(fakeMethod01net(THISVM,N), 0, [], [N = 0, THISVM = 1]).
6
7 eq(fakeMethod01peak(THISVM,N), 0, [cnew(THISVM)], [N > 0, THISVM < 3]).
8 eq(fakeMethod01peak(THISVM,N), 0, [cnew(THISVM),cnew(THISVM)], [N > 0, THISVM < 3]).
9 eq(fakeMethod01peak(THISVM,N), 0, [cnew(THISVM),cnew(THISVM),doubleRelease011peak(THISVM, 2)], [N > 0, THISVM < 3]).
10 eq(fakeMethod01peak(THISVM,N), 0, [cnew(THISVM),cnew(THISVM),doubleRelease011net(THISVM, 2),
11   crel(1)], [N > 0, THISVM < 3]).
12 eq(fakeMethod01peak(THISVM,N), 0, [cnew(THISVM),cnew(THISVM),doubleRelease011net(THISVM, 2),
13   crel(1),fakeMethod01peak(THISVM, N + 1)], [N > 0, THISVM < 3]).
14 eq(fakeMethod01peak(THISVM,N), 0, [cnew(THISVM),cnew(THISVM),doubleRelease011net(THISVM, 2),
15   crel(1),fakeMethod01net(THISVM, N + 1)], [N > 0, THISVM < 3]).
16
17 eq(fakeMethod01net(THISVM,N), 0, [], [THISVM = 3]).
18 eq(fakeMethod01net(THISVM,N), 0, [cnew(THISVM),cnew(THISVM),doubleRelease011net(THISVM, 2),
19   crel(1),fakeMethod01net(THISVM, N + 1)], [N > 0, THISVM = 1]).
20
21 eq(doubleRelease011peak(THISVM,X), 0, [], [THISVM = 3]).
22 eq(doubleRelease011peak(THISVM,X), 0, [], [THISVM < 3]).
23 eq(doubleRelease011peak(THISVM,X), 0, [crel(X)], [THISVM < 3]).
24 eq(doubleRelease011peak(THISVM,X), 0, [crel(X),crel(3)], [THISVM < 3]).
25
26 eq(doubleRelease011net(THISVM,X), 0, [], [THISVM = 3]).
27 eq(doubleRelease011net(THISVM,X), 0, [crel(X),crel(3)], [THISVM = 1]).

```

Figure 8: Cost equations of fake_method and double_release in CoFloCo format

- The three tools targets three different programs: C4B targets C, SACO targets ABS, and SRA targets vm1. Therefore, in order to compare the three tools, every test program has been rewritten according to the corresponding language of the tool. The sources of these tests are available in github ⁴.
- Every tool uses different metrics for quantitative analysis. SRA quantifies the resource usage by means of the operations `new VM()` and `release`. C4B has two metrics: the (i) back-edge metrics that assigns a cost of 1 to every back edge in the control-flow graph, thus not allowing negative costs, and (ii) the tick metric that uses a special operation `tick(n)` that has a cost of n where n can be either a positive or negative number. SACO, on the other hand, considers both of these kinds of metrics and, in addition, it also support other ones that are specific to the domain of the targeted language.
- Both SACO and SRA separate *peak costs* and *net costs*. This distinction is not made in C4B because it does not address concurrency. Therefore we restrict the results of Table 9 only to the *net costs* of the corresponding programs.

Table 9 reports the results of our analysis. It is worth to observe that SRA returns better results than both C4B and SACO for the program `betterThanAmortized`. This program, reported in Figure 10, defines a method `betterThanAmortised(n, m, x)` that recursively invokes itself and, when ($n > m$) every invocation has cost 1, otherwise it has cost 2. Therefore, the cost of `betterThanAmortised($2 * n, n, x$)` is $3 * n$ and this is the result of SRA. However, the other techniques do not recognize that the most costly branch is executed only half of the times and return $2 * (2 * n) = 4 * n$. Such precision is achieved in SRA by means of the treatment given to numerical expressions, and by the use of a solver (as [13]) capable to deal with conditional branches accurately. This precision is lost in techniques based on type systems because of the unifications that are necessary to enforce on the branches of conditionals, in order to type the continuations. We also observe that `betterThanAmortized` is not an exotic, ad-hoc program but a rather common programming pattern (a sequence of nested conditional statements).

SRA also returns remarkable results (i) in presence of unbounded recursion – see the example `producerConsumer` –, (ii) in cases where the identity of the arguments is relevant for the cost analysis – see the example `doubleRelease` –, and (iii) for standard patterns of concurrency – see the examples `doWorkAsync` and `producerConsumer`.

⁴<https://github.com/abelunibo/SRA-Tests-Files>

Program	C4B	SACO	SRA
betterThanAmortised	$2 * n$	$2 * n$	$n + m$
greatestCommonDivisor	$Max(n, m)$	<i>failed</i>	$n + m$
costlyFactorial	n	n	n
doWorkSync	2	2	2
quicksort	$2 * n$	<i>not analysable</i>	<i>not analysable</i>
SPEED1	$2 * n$	<i>not analysable</i>	<i>infinite</i>
doWorkASync	<i>not analysable</i>	n	n
producerConsumer	<i>not analysable</i>	<i>not analysable</i>	3
doubleRelease	<i>failed</i>	-2	$[x = y]-1; [x \neq y]-2$

Figure 9: Comparison of C4B, SACO, and SRA for some simple programs

```

Int betterThanAmortized(Int n, Int m, VM x) {
  Fut<Int> f; Int z;
  if (n==0) return 0;
  else if (n>m) {
    VM v = new VM();
    f = x!betterThanAmortized(n-1,m,v);
    z = f.get; return 0; }
  } else {
    VM v = new VM(); VM w = new VM();
    f = v!betterThanAmortized(n-1,m,w);
    z = f.get; return 0;
  }
}

{
  VM x = new VM(); Int z;
  Fut<Int> f = this!betterThanAmortized(2*n, n, x);
  f.get;
}

```

Figure 10: The program `betterThanAmortized`

As a main downside of SRA, it is unsatisfactory for computing costs in presence of multi-variate arguments, as in `greatestCommonDivisor` and `SPEED1`. However, it is worth to observe that this loss of precision is actually a limitation of the cost equation solver that cannot handle multi-variate arguments in its current version. For a similar reason, SRA fails in computing the cost of the `quicksort` program because of a restriction in the CoFloCo solver that does not handle non-linear recursion (while SRA correctly derives the behavioural types of `quicksort`).

7. Related Work

After the pioneering work by Wegbreit in 1975 [28] that discussed a method for deriving upper-bounds costs of functional programs, a number of cost analysis techniques have been developed. Those ones that are closely related to this contribution are based either on cost equations (solvers) or on amortized analysis.

The techniques based on cost equations address cost analysis in three steps by: (i) extracting relevant information out of the original programs by abstracting data structures to their size and assigning a cost to every program expression, (ii) converting the abstract program into cost equations, and (iii) solving the cost equations with an automatic tool. Recent advances have been done for improving the accuracy of upper-bounds for cost equations [2, 3, 10, 13, 18] and we refer to [13] for a comparison of these tools. The main advantage of these techniques is that cost equations may carry Presburger arithmetic conditions thus supporting a precise cost analysis of conditional statements. As a downside, techniques that extract control flow graphs and use control flow refinement techniques (such as [3, 4]) have a less precise alias analysis and name identity management than our technique (Section 4). These two features are

essential for function or procedure abstraction; in fact, a weak approach would jeopardise compositional reasoning when large programs are considered.

The techniques based on amortized analysis [27] associate so-called potentials to program expressions by means of type systems (these potentials determine the resources needed for each expression to be evaluated). The connection between the original program and the cost equations can be indeed demonstrated by a standard subject-reduction theorem [12, 20, 22–24]. While the techniques based on types are intrinsically compositional and, more importantly, type derivations can be seen as certificates of abstract descriptions of functions, type based methods do not model the interaction of integer arithmetic with resource usage, thus being less accurate in some cases (such as the program `betterThanAmortised` in Figure 10).

Our technique combines the advantages of the two approaches discussed above. It is modular, like the techniques based on cost equations, our one also consists of three steps, and it extracts the relevant information of programs by means of a *behavioural* type system, like the technique based on amortized analysis. Therefore, our technique is compositional and can be proved sound by means of a standard subject-reduction theorem. At the same time it is accurate in modelling the interaction of integer arithmetic with resource usage.

A common feature of cost analysis techniques in the literature is that they analyze *cumulative resources*. That is, resources that *do not decrease* during the execution of the programs, such as execution time, number of operations, memory (without an explicit `free` operation). As already discussed, this assumption eases the analysis because it permits to compute over-approximated cost. On the contrary, the presence of an *explicit or implicit* release operation entangles the analysis. In [6], a memory cost analysis is proposed for languages with garbage collection. It is worth to say that the setting of [6] is not difficult because, by definition of garbage collection, released memory is always inactive. The impact of the release operation in the cost analysis is thoroughly discussed in [7] by means of the notions of *peak cost* and *net cost* that we have also used in Section 5. It is worth to notice that, for cumulative analysis, this two notions coincide while, in non-cumulative analysis (in presence of a release operation), they are different and the *net cost* is key for computing tight upper bounds.

Recently [5] has analysed the cost of a language with explicit releases. We observe that the release operation studied in [5] is used in a very restrictive way: only locally created resources can be released. This constraint guarantees that costs of functions are always not negative, thus permitting the (re)use of non-negative cost models of cumulative analysis.

We conclude by discussing cost analysis techniques for concurrent systems, which are indeed very few [1, 4, 21]. In order to reduce the imprecision of the analysis caused by the nondeterminism, [1, 4] use a clever technique for isolating sequential code from parallel code, called *may-happen-in-parallel* [8]. We notice that no one of these contributions consider a concurrent language with a powerful `release` operation that allows one remove the resources taken in input. In fact, without this operation, one can model the cost by simply aggregating the sets of operations that can occur in parallel, as in [4], and all the theoretical development is much easier.

8. Conclusion

This paper presents the first (to the best of our knowledge) static analysis technique that computes upper bound of virtual machines usages in concurrent programs that may create and, more importantly, may release such machines. Our analysis consists of a type system that extracts relevant information about resource usages in programs, called behavioural types; an automatic translation that transforms these types into cost expressions; the application of solvers, like CoFloCo [13], on these expressions that compute upper bounds of the usage of virtual machines in the original program. A relevant property of our technique is its modularity. For the sake of simplicity, we have applied the technique to a small language. However, by either extending or changing the type system, the analysis can be applied to many other languages with primitives for creating and releasing resources. In addition, by changing the translation algorithm, it is possible to target other solvers that may compute better upper bounds, such as those based on multivariate amortized analysis [19].

For the future, we consider at least two lines of work. First, we intend to alleviate the restrictions introduced in Section 3 on the programs we can analyse. This may be pursued by retaining more expressive notations for the effect of a method, *i.e.* by considering R as a set of sets instead of a simple set. Such a notation is more suited for modelling nondeterministic behaviours and it might be made even more expressive by tagging all the different effects

in R with a condition specifying when such effect is yielded. Clearly, the management of these domains becomes more complex and the trade-off between simplicity and expressiveness must be carefully evaluated. Second, we intend to implement our analysis targeting a programming language with a formal model as ABS [25], of which `vm1` is a very basic sub-calculus. The current prototype translates a behavioural type program into cost equations.

Our behavioural types are abstract specifications of programs that highlight resource usages. In our mind, these abstract specifications must bridge the gap between programs and performance constraints in Service Level Agreement (SLA, in short) documents. As discussed in [16], `vm1` (actually, full ABS) is intended to be a language using Amazon Elastic Cloud Computing features as library functions. Therefore, `vm1` might be used to model cloud services, where a primary issue is the compatibility with the SLA document. The technique conveyed by the present article is to (i) extract the abstract specifications by means of the type (inference) system, (ii) compute the cost by means of the solver, and (iii) verify whether the computed cost complies with the SLA or not. This paper demonstrates that these steps are all correct. In this respect, a relevant future issue is the integration of our technique with other solvers, such as theorem provers, to be used in those cases where the cost equation solver is too imprecise.

Acknowledgements. We thank Elena Giachino for the valuable discussions and for the help in the assessment of our technique with respect to those based on amortized analysis. We also thank Antonio Flores-Montoya and Samir Genaim for the discussions about the cost analysers `CoFloCo` and `PUBS`.

References

- [1] Albert, E., Arenas, P., Correias, J., Genaim, S., Gómez-Zamalloa, M., Puebla, G., Román-Díez, G., 2015. Object-sensitive cost analysis for concurrent objects. *Software Testing, Verification and Reliability* 25 (3), 218–271. *stvr.1569*.
URL <http://dx.doi.org/10.1002/stvr.1569>
- [2] Albert, E., Arenas, P., Genaim, S., Puebla, G., 2008. Automatic inference of upper bounds for recurrence relations in cost analysis. In: *Proceedings SAS 2008*. Vol. 5079 of *Lecture Notes in Computer Science*. Springer, pp. 221–237.
- [3] Albert, E., Arenas, P., Genaim, S., Puebla, G., Zanardini, D., 2012. Cost analysis of object-oriented bytecode programs. *Theoretical Computer Science* 413 (1), 142–159.
- [4] Albert, E., Correias, J., Román-Díez, G., 2014. Peak cost analysis of distributed systems. In: *Proceedings of SAS 2014*. Vol. 8723 of *Lecture Notes in Computer Science*. Springer, pp. 18–33.
- [5] Albert, E., Correias, J., Román-Díez, G., 2015. Non-Cumulative Resource Analysis. In: *Proceedings of TACAS 2015*. *Lecture Notes in Computer Science*. Springer, pp. 85–100.
- [6] Albert, E., Genaim, S., Gómez-Zamalloa, M., 2010. Parametric inference of memory requirements for garbage collected languages. *SIGPLAN Not.* 45 (8), 121–130.
- [7] Alonso-Blas, D. E., Genaim, S., 2012. On the limits of the classical approach to cost analysis. In: *Proceedings of SAS 2012*. Vol. 7460 of *Lecture Notes in Computer Science*. Springer, pp. 405–421.
- [8] Barik, R., 2006. Efficient computation of may-happen-in-parallel information for concurrent java programs. In: *Proceedings of LCPC 2005*. Vol. 4339 of *Lecture Notes in Computer Science*. Springer, pp. 152–169.
- [9] Berman, L., 1980. The complexity of logical theories. *Theoretical Computer Science* 11 (1), 71 – 77.
URL <http://www.sciencedirect.com/science/article/pii/0304397580900377>
- [10] Brockschmidt, M., Emmes, F., Falke, S., Fuhs, C., Giesl, J., 2014. Alternating runtime and size complexity analysis of integer programs. In: *Proceedings of TACAS 2014*. Vol. 8413 of *Lecture Notes in Computer Science*. Springer, pp. 140–155.
- [11] Carbonneaux, Q., Hoffmann, J., Shao, Z., 2015. Compositional certified resource bounds. In: *Proceedings of the 36th PLDI Conference*. ACM, pp. 467–478.
- [12] Chin, W., Nguyen, H. H., Qin, S., Rinard, M., 2005. Memory usage verification for oo programs. In: *Proceedings of SAS 2005*. Vol. 3672 of *Lecture Notes in Computer Science*. Springer, pp. 70–86.
- [13] Flores Montoya, A., Hähnle, R., 2014. Resource analysis of complex programs with cost equations. In: *Proceedings of 12th Asian Symposium on Programming Languages and Systems*. Vol. 8858 of *Lecture Notes in Computer Science*. Springer, pp. 275–295.
- [14] Garcia, A., Laneve, C., 2015. Static analyzer of resource usage upper bounds.
URL sra.cs.unibo.it
- [15] Garcia, A., Laneve, C., Lienhardt, M., 2015. Static analysis of cloud elasticity. In: *Proceedings of the 17th PPDP*. ACM, pp. 125–136.
- [16] Giachino, E., de Gouw, S., Laneve, C., Nobakht, B., 2016. Statically and dynamically verifiable sla metrics. To appear on *Lecture Notes in Computer Science*.
URL cs.unibo.it/~laneve/papers/SLA-metrics.pdf
- [17] Giachino, E., Laneve, C., Lienhardt, M., 2015. A framework for deadlock detection in ABS. *Software and Systems Modeling* To appear.
URL <http://dx.doi.org/10.1007/s10270-014-0444-y>
- [18] Gulwani, S., Mehra, K. K., Chilimbi, T., 2009. Speed: precise and efficient static estimation of program computational complexity. In: *ACM SIGPLAN Notices*. Vol. 44. ACM, pp. 127–139.
- [19] Hoffmann, J., Aehlig, K., Hofmann, M., 2011. Multivariate amortized resource analysis. In: *Proceedings of POPL 2011*. ACM, pp. 357–370.
- [20] Hoffmann, J., Aehlig, K., Hofmann, M., 2012. Multivariate amortized resource analysis. *ACM Trans. Program. Lang. Syst.* 34 (3), 14.
URL <http://doi.acm.org/10.1145/2362389.2362393>

- [21] Hoffmann, J., Shao, Z., 2015. Automatic static cost analysis for parallel programs. [Online; accessed 11-February-2015]. URL <http://cs.yale.edu/homes/hoffmann/papers/parallelcost2014.pdf>
- [22] Hofmann, M., Jost, S., 2003. Static prediction of heap space usage for first-order functional programs. In: Proceedings of POPL 2003. ACM, pp. 185–197.
- [23] Hofmann, M., Jost, S., 2006. Type-based amortised heap-space analysis. In: Proceedings of ESOP 2006. Vol. 3924 of Lecture Notes in Computer Science. Springer, pp. 22–37.
- [24] Hofmann, M., Rodriguez, D., 2009. Efficient type-checking for amortised heap-space analysis. In: Proceedings of CSL 2009. Vol. 5771 of Lecture Notes in Computer Science. Springer, pp. 317–331.
- [25] Johnsen, E. B., Hähnle, R., Schäfer, J., Schlatte, R., Steffen, M., 2011. ABS: A core language for abstract behavioral specification. In: Proceedings of FMCO 2010. Vol. 6957 of Lecture Notes in Computer Science. Springer, pp. 142–164.
- [26] Mathew, S., November 2014. Overview of amazon web services, at d0.awsstatic.com/whitepapers/aws-overview.pdf.
- [27] Tarjan, R. E., 1985. Amortized computational complexity. SIAM Journal on Algebraic Discrete Methods 6 (2), 306–318.
- [28] Wegbreit, B., 1975. Mechanical program analysis. Communications of the ACM 18 (9), 528–539.

Appendix A. The proof of correctness of the type system

Appendix A.1. Type system correctness

As discussed in Section 5, the correctness of our analysis technique consists of two parts. The first part addresses the correctness of the type system in Section 4, which is usually expressed by a subject reduction theorem. This theorem states that if a configuration cn of the operational semantics is well typed and $cn \rightarrow cn'$ then cn' is well-typed as well. It is worth to observe that we cannot hope to demonstrate a statement guaranteeing type-preservation because our types are “behavioural” and change during the evolution of the systems. However, it is critical for the correctness of the cost analysis that there exists a relation between the type of cn and the type of cn' .

Therefore, a subject reduction for the type system of Section 4 requires

1. the extension of the typing to configurations;
2. the definition of an evaluation relation \rightsquigarrow between behavioural types.

The extension of the typing to configurations. When typing configurations there are two kinds of vm names: *static-time vm names* and *runtime vm names*. The former names are ranged over by α, β, \dots ; the latter names, which also include *err*, are ranged over by o, o', \dots . With an abuse of notation, in the following, α and \oplus will also range over runtime vm names o and err . This abuse will let us to reuse the typing rules in Figure 4 (with one exception, see below). The syntax of runtime types is:

$\mathbb{k} ::= 0 \mid f(\perp) \mid f(_) \mid f(o) \mid (f, \mathfrak{m} \ o(\overline{\mathfrak{f}})) \mid vm(o, a \mid \mathfrak{f}; \overline{\mathfrak{f}}) \mid \mathbb{k} \parallel \mathbb{k}$	runtime types
$\mathfrak{f} ::= (f, \mathfrak{c}, S)$	process types
$\mathfrak{a} ::= \dots \mid [\oplus]$	extended atom

The syntactic category \mathbb{k} models behavioural types only existing at runtime, in particular, virtual machines, futures, and processes. The term 0 correspond to an element of the running program that does not have any effect on its cost (e.g., ϵ); $f(\perp)$, $f(_)$, and $f(o)$ type a future binder: f is the name of the future, \perp or $_$ or o identify the value of the future, namely either \perp , if it is still not computed, or $_$ or o , if it has been already computed (see the rest of the paragraph for more details); $(f, \mathfrak{m} \ o(\overline{\mathfrak{f}}))$ corresponds to an invocation message. The type $vm(o, a \mid \mathfrak{f}; \overline{\mathfrak{f}})$ corresponds to a virtual machine (alive if $a = \top$, dead otherwise) executing the process \mathfrak{f} and with a set of processes $\overline{\mathfrak{f}}$ to execute. Processes are triples (f, \mathfrak{c}, S) where:

- f is the identity of the process (i.e. the name of its future);
- \mathfrak{c} is the type of the statement the process will execute;
- S is a set of runtime vm names.

The term $\mathbb{k} \parallel \mathbb{k}$ types the parallel composition of runtime configurations. Atoms are extended with the term $[\oplus]$ that represents the returned value of a method – see the discussion below.

The runtime type system consists of rules in Figure 4 plus the rules in Figure A.11. The reader may notice that the first rule (BT-RETURN) replaces the definition of the judgment $\Gamma \vdash_S s : \mathfrak{c} \triangleright \Gamma'$ for the `return` statement in Figure 4. This modification, together with the extensions of atoms, is used to refine the (static-time) typing rules in order to deal

$$\begin{array}{c}
\text{(BT-RETURN)} \\
\frac{\Gamma \vdash e : \odot \quad \Gamma \vdash \text{destiny} : \odot' \quad \odot \notin S}{\Gamma \vdash_S \text{return } e : [\odot] \triangleright \Gamma[\odot' \mapsto \Gamma(\odot)]} \\
\\
\text{BT-FUT-RUNNING} \\
\frac{\Gamma(f) = (\odot, o, \text{Ft}, R)}{\Gamma \vdash \text{fut}(f, \perp) : f(\perp)} \\
\\
\text{BT-FUT-COMPUTED} \\
\frac{\Gamma \vdash v : \odot \quad \Gamma \vdash f : \odot}{\Gamma \vdash \text{fut}(f, v) : f(\odot)} \\
\\
\text{BT-FUT-UNSYNC} \\
\frac{\Gamma(f) = (\odot, o, \text{Ft}, R) \quad v \neq \perp \quad \Gamma \vdash v : \odot' \quad \alpha \in R[\odot' / \odot] \Rightarrow \Gamma(\alpha) = \emptyset \perp}{\Gamma \vdash \text{fut}(f, v) : f(\odot')} \\
\\
\text{BT-EMPTY} \\
\Gamma \vdash \epsilon : 0 \\
\\
\text{BT-INVOKER} \\
\frac{\Gamma \vdash m o(\bar{v}) : \odot, R \quad \Gamma(f) = (\odot, o, \text{Ft}, R) \quad \Gamma \vdash \bar{v} : \bar{\Gamma}}{\Gamma \vdash \text{invoc}(o, f, m, \bar{v}) : (f, o m(\bar{\Gamma}))} \\
\\
\text{BT-VM} \\
\frac{\Gamma \vdash o : \text{Fa} \quad (\Gamma \vdash_{S_i} p_i : \mathbb{f}_i \quad \mathbb{f}_i = (f_i, c_i, S_i))^{i \in 1..n}}{\Gamma \vdash \text{vm}(o, a, p_1; \{p_2 \dots p_n\}) : \text{vm}(o, a \mid \mathbb{f}_1; \{\mathbb{f}_2, \dots, \mathbb{f}_n\})} \\
\\
\text{BT-PARALLEL} \\
\frac{\Gamma \vdash cn : \mathbb{k} \quad \Gamma \vdash cn' : \mathbb{k}'}{\Gamma \vdash cn \text{ } cn' : \mathbb{k} \parallel \mathbb{k}'} \\
\\
\text{BT-PROCESS} \\
\frac{\Gamma[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \odot] \vdash_S s : C[\mathbb{a}_1 \triangleright \Gamma_1] \cdots [\mathbb{a}_n \triangleright \Gamma_n] \quad \begin{array}{l} l = [(x_i \mapsto v_i)^{i \in I}, \text{destiny} \mapsto f] \quad (\Gamma \vdash v_i : \mathbb{x}_i)^{i \in I} \quad \Gamma(f) = (\odot, o, \text{Ft}, R) \\ \Gamma_i(\gamma) = \Gamma_j(\gamma) \end{array}^{i, j \in 1..n, \gamma \in S} \quad R = (S \cup \text{fv}(\odot)) \cap \{\gamma \mid \Gamma_1(\gamma) = \text{F}\perp\}}{\Gamma \vdash_S \{l \mid s\} : (f, C[\mathbb{a}_1 \triangleright \emptyset] \cdots [\mathbb{a}_n \triangleright \emptyset], S)}
\end{array}$$

Figure A.11: Runtime Typing rules

with a naming issue in the subject reduction theorem. In particular, when a method returns a freshly created virtual machine, we associate to it a fresh name in (T-INVOKER) that we can use later in (T-GET) and in the rest of the typing of the (static) program. The problem is that, when we type the method invocation (i.e. when we collect the information about the method and its return value in order to type future usage of this value), the identity of the returned virtual machine is still unknown because the virtual machine has not yet been created. In order to solve this problem, we extend atoms with the term $[\odot]$, which retains the returned value and we use the new rule (BT-RETURN) to type the return statements.

The other rule of Figure A.11 that we comment is (BT-VM). This rule allows one to type a vm term containing a number of processes p . To this aim, one has to type every p by means of the judgment $\Gamma[\text{this} \mapsto o] \vdash_S p : (f, c, S)$; namely the index S and the third argument of the process type must be the same. In fact, this argument records the actual vm names of the corresponding method invocation – see rule (R-BIND-MTD) in Figure A.12. In turn, this set is used in (BT-PROCESS) to type the statements of the processes.

Lemma Appendix A.1 (Substitution Lemma). *Let $\Gamma \vdash cn : \mathbb{k}$ and let ι be an injective renaming of names, vm names, and future names. Let $\iota(\Gamma) \stackrel{\text{def}}{=} [\iota(x) \mapsto \iota(\Gamma(x))]^{x \in \text{dom}(\Gamma)}$.*

Then $\iota(\Gamma) \vdash \iota(cn) : \iota(\mathbb{k})$. Similarly for $\Gamma \vdash_S \{l \mid s\} : (f, c)$ and $\Gamma \vdash_S s : c \triangleright \Gamma'$ (in these two cases also S is mapped by ι) and for $\Gamma \vdash e : \mathbb{x}$.

Proof. The proof proceeds straightforward by induction on the structure of s . □

Runtime Types Evolution. As usual with behavioural types, our types change while the programs execute. These changes can be defined and, more importantly, can be related to the cost analysis. The modifications of types, noted \rightsquigarrow , is defined as a reduction relation in Figure A.12. The rules are mostly an adaptation of the operational semantics of vm1 at the type level. Therefore, to ease the understanding of the relation \rightsquigarrow , we use similar names for similar rules. Nevertheless, there are few differences from the operational semantics of vm1: *i*) rule (ASSIGN) has no corresponding rule in the definition of \rightsquigarrow (assignments are managed at the type system level); *ii*) rule (ACTIVATE) is replaced by the rule (R-GC) (there are no scheduling policies in our runtime types, so the only effect of the rule (ACTIVATE) at the type level is the deletion of a finished process); *iii*) similarly, rule (ACTIVATE-GET) has no corresponding rule in the definition of \rightsquigarrow ; and *iv*) because at runtime all objects identities are known, the rules (READ-FUT), (ASYNC-CALL), (BIND-MTD) and (NEW-VM) have been adjusted in order to replace the names generated a static time with the objects' actual identities.

R-SKIP $\frac{vm(o, \top \mid (f, 0 \circledast c, S); \bar{f})}{\rightsquigarrow vm(o, \top \mid (f, c, S); \bar{f})}$	R-READ-FUT $\frac{vm(o, \top \mid (f, f'_o \checkmark \circledast c, S); \bar{f}) \parallel f'(\circ')}{\rightsquigarrow vm(o, \top \mid (f, c[\circ'/\circ], S); \bar{f}) \parallel f'(\circ')}$	R-ASYNC-CALL $\frac{f'' \text{ fresh}}{vm(o, \top \mid (f, \nu f' : m \circ'(\bar{v}, \bar{\beta}) \circledast c, S); \bar{f}) \rightsquigarrow vm(o, \top \mid (f, c[f''/f'], S); \bar{f}) \parallel (f'', m \circ'(\bar{v}, \bar{\beta}))}$
R-BIND-MTD $\frac{m \alpha(\text{Int } \bar{x}, \bar{M} \bar{z}) = c, \circ, R \quad S = \{o, \bar{o}\}}{vm(o, \top \mid \bar{f}; \bar{f}) \parallel (f, o \ m(\bar{v}, \bar{o})) \rightsquigarrow vm(o, \top \mid \bar{f}; \bar{f} \cup \{(f, c[\bar{v}, \bar{o}/\bar{x}, \bar{z}][\circ/\alpha], S)\}) \parallel f(\perp)}$	R-LCHOICE $\frac{vm(o, \top \mid (f, c + c', S); \bar{f})}{\rightsquigarrow vm(o, \top \mid (f, c, S); \bar{f})}$	R-RCHOICE $\frac{vm(o, \top \mid (f, c + c', S); \bar{f})}{\rightsquigarrow vm(o, \top \mid (f, c', S); \bar{f})}$
R-ITE-T $\frac{se = \text{true}}{vm(o, \top \mid (f, (se) \{c\}, S); \bar{f}) \rightsquigarrow vm(o, \top \mid (f, c, S); \bar{f})}$	R-ITE-F $\frac{se = \text{false}}{vm(o, \top \mid (f, (se) \{c\}, S); \bar{f}) \rightsquigarrow vm(o, \top \mid (f, 0, S); \bar{f})}$	R-NEW-VM $\frac{o' \text{ fresh}}{vm(o, \top \mid (f, \nu \beta \circledast c, S); \bar{f}) \rightsquigarrow vm(o, \top \mid (f, c[\circ'/\beta], S); \bar{f}) \parallel vm(o', \top \mid \emptyset)}$
R-RELEASE-VM $\frac{vm(o, \top \mid (f, o' \checkmark \circledast c, S); \bar{f}) \parallel vm(o', a' \mid \bar{f}'; \bar{f}')}{\rightsquigarrow vm(o, \top \mid (f, c, S); \bar{f}) \parallel vm(o', \perp \mid \bar{f}'; \bar{f}')}$	R-GC $\frac{vm(o, a \mid (f, 0 \triangleright \emptyset, S); \bar{f} \cup \{\bar{f}\})}{\rightsquigarrow vm(o, a \mid \bar{f}; \bar{f})}$	R-RETURN $\frac{vm(o, \top \mid (f, [\circ], S); \bar{f}) \parallel f(\perp)}{\rightsquigarrow vm(o, \top \mid (f, 0 \triangleright \emptyset, S); \bar{f}) \parallel f(\circ)}$
R-RELEASE-BOT $\frac{vm(o, \perp \mid (f, c, S); \bar{f}) \parallel f(\perp)}{\rightsquigarrow vm(o, \perp \mid (f, 0 \triangleright \emptyset, S); \bar{f}) \parallel f(err)}$	R-BIND-MTD-ERR $\frac{vm(o', \perp \mid \bar{f}; \bar{f}) \parallel (f, o' \ m(\bar{v}))}{\rightsquigarrow vm(o', \perp \mid \bar{f}; \bar{f}) \parallel f(err)}$	R-BIND-MTD-SERR $(f, err \ m(\bar{v})) \rightsquigarrow f(err)$
R-RELEASE-GC $\frac{vm(o, \top \mid (f, \dots o' \checkmark \dots, S); \bar{f}) \parallel vm(o', \perp \mid \dots)}{\rightsquigarrow vm(o, \top \mid (f, \dots 0 \dots, S); \bar{f}) \parallel vm(o', \perp \mid \dots)}$	R-CONTEXT $\frac{k \rightsquigarrow k'}{k \parallel k'' \rightsquigarrow k' \parallel k''}$	

Figure A.12: Operational Semantics of Runtime Types

The subject reduction theorem. Based on the previous construction, we present the subject reduction theorem. Lemma Appendix A.2 relates the static time type system to the runtime type system: this corresponds to the first step of the subject reduction property.

Lemma Appendix A.2. *Let $P = \overline{M} \{ \overline{F} \ x \ ; \ s \}$ be a vml program, and let $\Gamma \vdash P : \overline{\mathbb{C}}, c$. Then*

$$\Gamma[start \mapsto \top][f_{start} \mapsto (_, start, \emptyset \top, \emptyset)] \vdash vm(start, \top, \{[destiny \mapsto f_{start}] \mid s\}, \emptyset) : vm(start, \top \mid (f_{start}, c, \emptyset); \emptyset) .$$

Proof. The proof follows by the application of rules (BT-VM) and (BT-PROCESS), and using the hypothesis. \square

The main subject reduction theorem states that any well-typed runtime program reduces into another well-typed runtime program.

Theorem Appendix A.3 (Subject Reduction). *Let $\Gamma \vdash cn : \mathbb{k}$. If $cn \rightarrow cn'$ then there exists Γ' such that $\Gamma' \vdash cn' : \mathbb{k}'$ and $\mathbb{k} \rightsquigarrow^* \mathbb{k}'$.*

Proof. As usual for Subject Reduction theorems we proceed by case on the reduction rule used in $cn \rightarrow cn'$. Since our types are behavioral we cannot assume that the types of a runtime programs state and its reduction are equivalent, instead we prove that this types $(\mathbb{k}, \mathbb{k}')$ are related according to the operational semantics for runtime types, see Figure A.12. We present the proof of some cases, the rest are either straightforward or similar to those below. Whenever the construction of ι is omitted it is considered as the identity function. Finally, if not stated otherwise, we implicitly use the notation used in the considered reduction rule:

- Case (ASSIGN).

$$\frac{v = \llbracket e \rrbracket_{\iota}}{vm(o, \top, \{l \mid x = e; s\}, q) \rightarrow vm(o, \top, \{l[x \mapsto v] \mid s\}, q)}$$

By hypothesis we have that $\Gamma \vdash vm(o, \top, \{l \mid x = e; s\}, q) : \mathbb{k}$. By application of the corresponding typing rules we have that $\mathbb{k} = vm(o, \top \mid (f, 0 \circledast \mathbb{C}, S), \bar{\mathbb{F}})$. Let $\Gamma' = \Gamma$, then by application of rules (BT-VM) and (BT-PROCESS), and by hypothesis we have that $\Gamma' \vdash vm(o, \top, \{l[x \mapsto v] \mid s\}, q) : \mathbb{k}'$, with $\mathbb{k}' = vm(o, \top \mid (f, \mathbb{C}, S), \bar{\mathbb{F}})$. We notice that \mathbb{k} and \mathbb{k}' are related by rule (R-SKIP)

- Case (READ-FUT).

$$\frac{(READ-FUT) \quad f = \llbracket e \rrbracket_l \quad v \neq \perp}{vm(o, \top, \{l \mid x = e.get; s\}, q) \quad fut(f, v) \rightarrow vm(o, \top, \{l \mid x = v; s\}, q) \quad fut(f, v)}$$

By hypothesis we have that $\Gamma \vdash vm(o, \top, \{l \mid x = e.get; s\}, q) \quad fut(f, v) : \mathbb{k}$. Let consider the case in which $\Gamma(f) = (\mathbb{O}, \alpha, F\mathbb{t}, R)$ (the other case, where $\Gamma(f) = \mathbb{O}$, is trivial). By reconstruction of the proof tree of the hypothesis, with $l = \{x_i \mapsto \mathbb{x}_i \mid i \in I\}$, we derive that $\mathbb{k} = vm(o, \top \mid (f', f_o' \circledast \mathbb{C}, S), \bar{\mathbb{F}}) \parallel f(\mathbb{O}'')$ and that $\Gamma[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}'][\beta \mapsto \emptyset \perp]^{\beta \in R}[\beta' \mapsto \emptyset \mathbb{t}']^{\beta' \in R'}[f \mapsto \mathbb{O}] \vdash s : \mathbb{C}$ with $R' = fv(\mathbb{O}) \setminus R$ and $\mathbb{t}' = \mathbb{t} \sqcap (\Gamma(\alpha) \Downarrow)$.

Let $\iota = [\mathbb{O}''/\mathbb{O}]$: we have that $\iota(\Gamma) = \Gamma$. Let $\Gamma' = \Gamma[f \mapsto \mathbb{O}'']$, by lemma Appendix A.1 we have that $\Gamma'[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}'][\beta \mapsto \emptyset \perp]^{\beta \in \iota(R)}[\beta' \mapsto \emptyset \mathbb{t}']^{\beta' \in R''} \vdash s : \iota(\mathbb{C})$ with $R'' = fv(\mathbb{O}'') \setminus \iota(R)$. To prove the theorem in this case, we need to show that s is typable with the environment $\Gamma'[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}']$. First, because of the side condition in rule (BT-FUT-UNSYNCR), we have that $\Gamma'(\beta) = \emptyset \perp$ for all $\beta \in \iota(R)$. We thus have

$$\Gamma'[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}'][\beta \mapsto \emptyset \perp]^{\beta \in \iota(R)}[\beta' \mapsto \emptyset \mathbb{t}']^{\beta' \in R''} = \Gamma'[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}'][\beta' \mapsto \emptyset \mathbb{t}']^{\beta' \in R''}$$

Now, to remove the last substitution $[\beta' \mapsto \emptyset \mathbb{t}']^{\beta' \in R''}$, we have four cases:

- either R'' is empty, in which case that substitution is the identity;
- either \mathbb{t}' is \top , in which case, by construction of the set R , we have $\Gamma'(\beta') = \Gamma(\beta) = \top$: the substitution does not change Γ' ;
- either \mathbb{t}' is ∂ and $\Gamma'(\beta')$ is \top , in which case $\Gamma'[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}'] \vdash s : \iota(\mathbb{C})$ holds;
- either \mathbb{t}' is ∂ and $\Gamma'(\beta')$ is \perp , in which case $\Gamma'[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}'] \vdash s : \mathbb{C}'$ holds, with $\iota(\mathbb{C})$ and \mathbb{C}' being related with few (possibly none) application of the rule (R-RELEASE-GC): statements `release` β are now typed with (T-RELEASE-BOT) instead of (T-RELEASE).

We can then conclude that $\Gamma' \vdash vm(o, \top, \{l \mid x = v; s\}, q) \quad fut(f, v) : \mathbb{k}'$ with $\mathbb{k}' = vm(o, \top \mid (f', \mathbb{C}', S), \bar{\mathbb{F}}) \parallel f(\mathbb{O}'')$, and behavioral types \mathbb{k} and \mathbb{k}' are related by (R-READ-FUT), and possibly few application of (R-RELEASE-GC).

- Case (ASYNC-CALL).

$$\frac{(ASYNC-CALL) \quad o' = \llbracket e \rrbracket_l \quad \bar{v} = \llbracket \bar{e} \rrbracket_l \quad f = \text{fresh}(\)}{vm(o, \top, \{l \mid x = e!m(\bar{e}); s\}, q) \rightarrow vm(o, \top, \{l \mid x = f; s\}, q) \quad \text{invoc}(o', f, m, \bar{v})}$$

Lets consider the case when $\Gamma(o') = \top$ (the opposite case $\Gamma(o') = \perp$ is almost identical). By hypothesis we have that $\Gamma \vdash vm(o, \top, \{l \mid x = e!m(\bar{e}); s\}, q) : \mathbb{k}$. By reconstruction we have that $\mathbb{k} = vm(o, \top \mid (f', \nu f'' : m o'(\bar{\mathbb{S}}) \circledast \mathbb{C}, S'), \bar{\mathbb{F}})$ and that there is Γ'' such that $\Gamma[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}'] \vdash !m(\bar{e}) : \nu f'' : m o'(\bar{\mathbb{S}}) \triangleright \Gamma''$ and $\Gamma'' \vdash s : \mathbb{C}$.

Let $\iota = [f/f']$ and lets choose Γ' such that $\Gamma'' = \Gamma'[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}']$, by lemma Appendix A.1 we have that $\iota(\Gamma')[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \mathbb{O}'] \vdash vm(o, \top, \{l \mid x = e!m(\bar{e}); s\}, q) \rightarrow vm(o, \top, \{l \mid x = f; s\}, q) \quad \text{invoc}(o', f, m, \bar{v}) : \mathbb{k}'$ with $\mathbb{k}' = vm(o, \top \mid (f', 0 \circledast \mathbb{C}', S'), \bar{\mathbb{F}}) \parallel (f'', m o'(\bar{\mathbb{T}}))$, where $\mathbb{C}' = \iota(\mathbb{C})$. Behavioral types \mathbb{k} and \mathbb{k}' are related by (R-ASYNC-CALL).

- Case (RELEASE-VM).

$$\frac{\text{(RELEASE-VM)} \quad o' = \llbracket e \rrbracket_l \quad o \neq o'}{vm(o, \top, \{l \mid \text{release } e; s\}, q) \quad vm(o', a', p', q') \rightarrow vm(o, \top, \{l \mid s\}, q) \quad vm(o', \perp, p', q')}$$

By hypothesis we have that $\Gamma \vdash vm(o, \top, \{l \mid \text{release } e; s\}, q) \quad vm(o', a', p', q') : \mathbb{k}$ and we can derive that $\mathbb{k} = vm(o, a \mid (f, o' \checkmark \circ \mathbb{C}, S), \bar{\mathbb{f}}) \parallel vm(o', a' \mid \bar{\mathbb{f}}')$. Let $\Gamma' = \Gamma[o' \mapsto \emptyset \perp]$, we have that $\Gamma' \vdash vm(o, \top, \{l \mid s\}, q) \quad vm(o', \perp, p', q') : \mathbb{k}'$ with $\mathbb{k}' = vm(o, a \mid (f, \mathbb{C}, S), \bar{\mathbb{f}}) \parallel vm(o', \perp \mid \bar{\mathbb{f}}')$. Behavioral types \mathbb{k} and \mathbb{k}' are related by (R-RELEASE-VM)

- Case (RETURN).

$$\frac{\text{(RETURN)} \quad v = \llbracket e \rrbracket_l \quad f = l(\text{destiny})}{vm(o, \top, \{l \mid \text{return } e\}, q) \quad fut(f, \perp) \rightarrow vm(o, \top, \{l \mid \varepsilon\}, q) \quad fut(f, v)}$$

By reconstruction of the proof tree of the hypothesis after the application of rules (BT-PARALLEL), (BT-VM), (BT-PROCESS) and (BT-RETURN), we can derive that $\mathbb{k} = vm(o, \top \mid (f, [\odot], S), \bar{\mathbb{f}}) \parallel f(\perp)$. Let $\Gamma' = \Gamma$, by hypothesis and by rule (BT-FUT-UNSYNC) we have that $\Gamma' \vdash vm(o, \top, \{l \mid \varepsilon\}, p_2 \dots p_n) \quad fut(f, v) : \mathbb{k}'$ with $\mathbb{k}' = vm(o, \top \mid (f, 0 \triangleright \emptyset, S), \bar{\mathbb{f}}) \parallel f(\odot)$. Note that the side condition $\alpha \in R[\odot' / \odot] \Rightarrow \Gamma(\alpha) = \emptyset \perp$ of the rule (BT-FUT-UNSYNC) is a consequence of the constraints $R = (S \cup fv(\odot)) \cap \{\gamma \mid \Gamma_1(\gamma) = F \perp\}$ of the (BT-PROCESS) rule. Finally, the behavioral types \mathbb{k} and \mathbb{k}' are related by (R-RETURN).

- Case (VM-ERR-RETURN).

$$\frac{\text{(VM-ERR-RETURN)} \quad f = l(\text{destiny})}{vm(o, \perp, \{l \mid s\}, q) \quad fut(f, \perp) \rightarrow vm(o, \perp, \{l; \varepsilon\}, q) \quad fut(f, err)}$$

Let $\Gamma \vdash vm(o, \perp, \{l \mid s\}, q) \quad fut(f, \perp) : \mathbb{k}$ where $\mathbb{k} = vm(o, \perp \mid (f, 0 \circ \mathbb{C}, S), \bar{\mathbb{f}}) \parallel f(\perp)$. Let $\Gamma' = \Gamma$, the typing of left and right configurations produces an almost identical derivation tree, hence $\Gamma' \vdash vm(o, \perp, \{l \mid \varepsilon\}, q) \quad fut(f, err) : \mathbb{k}'$, and $\mathbb{k}' = vm(o, \perp \mid (f, 0 \triangleright \emptyset, S), \bar{\mathbb{f}}) \parallel fut(f, err)$, which are related by rule (R-RELEASE-BOT).

□

Appendix A.2. Cost Computation

The second part of the correctness intends to demonstrate that, given a computation, the costs of configurations therein do not increase. This means that the cost of the first configuration is the greatest one and, therefore, the cost of the program gives an upper bound of the actual cost of every computation of its.

In order to prove this property, we need to extend the notion of cost to runtime types, and show two properties of these cost: i) they are an upper bound for the number of alive vm in the typed configuration; and ii) they decrease when the configuration executes, thus showing that the cost computed for the initial program is an upper bound for the number of alive vm at every step of its execution.

Cost of Runtime Types. In this paragraph, we extend the definition of the `translate` function to runtime type \mathbb{k} . we first introduce our extension with some useful getters on *well-formed* runtime types (i.e., types corresponding to running configurations)

Definition Appendix A.4. A runtime type \mathbb{k} is well-formed iff all the names (object, future, etc) it uses are declared and if it contains the object: $vm(\text{start}, a \mid (f_{\text{start}}, \mathbb{C}, S); \bar{\mathbb{f}})$. Given a well-formed runtime type \mathbb{k} , we write: **alive**(\mathbb{k}) the number of alive vm in \mathbb{k} (i.e., the vm whose state is \top in \mathbb{k}); **mapping**(\mathbb{k}) the mapping giving the state of vm in \mathbb{k} ; and $\Delta_{\mathbb{k}}$ the set $\{\{o\} \mid o \text{ is an object of } \mathbb{k}\}$ (i.e., $\Delta_{\mathbb{k}}$ is the identity equivalence relation generated from \mathbb{k}).

We extend the function `translate` as presented in Figure A.13 and define:

Extension of translate on extended atoms:

$$\text{translate}[\Delta, \Psi, \alpha](\bar{e}; e)(\mathfrak{a}) = \begin{cases} (\Psi, \bar{e}; e) & \text{when } \mathfrak{a} = [o] \\ (\Psi, (\bar{e}; e)\sigma; (e)\sigma') & \text{when } \mathfrak{a} = f_{\circ}^{\vee} \text{ and } f \notin \text{dom}(\Psi), \text{ where } \sigma = [f_{\text{peak}}/f] \text{ and } \sigma' = [f_{\text{net}}/f] \\ (\Psi', (\bar{e}; e)\sigma; (e)\sigma') & \text{when } \Psi(f) = (\Psi', R, \mathfrak{m}\beta(\bar{\tau}, \bar{\beta}) \rightarrow \circ) \text{ and } \mathfrak{a} = f_{\circ}^{\vee} \\ & \text{and } \text{EqRel}(\beta, \bar{\beta}) = \Xi \text{ and } \Theta = \Psi' \searrow \Psi \\ & \text{and } \sigma = [\mathfrak{m}_{\text{peak}}^{\Xi}(\bar{\tau}, \Theta(\Xi(\beta, \bar{\beta})))/f] \text{ and } \sigma' = [\mathfrak{m}_{\text{net}}^{\Xi}(\bar{\tau}, \Theta(\Xi(\beta, \bar{\beta})))/f] \end{cases}$$

Extension of translate on process types:

$$\text{translate}(\Delta, \Psi, o, (f, \mathbb{C}, S)) = \begin{cases} f_{\text{peak}} = e_{1,1} & [se_1] \\ \vdots & \\ f_{\text{peak}} = e_{h_n, n} & [se_n] \\ f_{\text{net}} = f_{\text{peak}} & [\alpha = \partial] \\ f_{\text{net}} = e_{h_1, 1} & [se_1 \wedge \alpha = \top] \\ \vdots & \\ f_{\text{net}} = e_{h_n, n} & [se_n \wedge \alpha = \top] \end{cases} \quad \text{with } \text{translate}(\Delta, \Psi[o \mapsto \alpha], (\text{true})0, \mathbb{C}) = \bigcup_{i=1}^n (se_i)\{e_{1,i}; \dots; e_{h_i,i}\}$$

Extension of translate on runtime types:

$$\begin{aligned} \text{translate}(\Delta, \Psi, f(\circ)) &= \{f_{\text{peak}} = 0, f_{\text{net}} = 0\} & \text{translate}(\Delta, \Psi, \text{vm}(o, \top \mid \mathfrak{f}_1; \mathfrak{f}_2 \dots, \mathfrak{f}_n)) &= \bigcup_{i \in 1..n} \text{translate}(\Delta, \Psi, o, \mathfrak{f}_i) \\ \text{translate}(\Delta, \Psi, \text{vm}(o, \perp \mid \mathfrak{f}; \bar{\mathfrak{f}})) &= \emptyset & \frac{\Xi = \text{EqRel}(\bar{v}) \quad \bar{v}' = \Psi(\Xi(\bar{v}))}{\text{translate}(\Delta, \Psi, (f, o \mathfrak{m}(\bar{v})))} &= \begin{cases} f_{\text{peak}} = \mathfrak{m}_{\text{peak}}^{\Xi}(\bar{v}') \\ f_{\text{net}} = \mathfrak{m}_{\text{net}}^{\Xi}(\bar{v}') \end{cases} \\ \text{translate}(\Delta, \Psi, \mathbb{k} \parallel \mathbb{k}') &= \text{translate}(\Delta, \Psi, \mathbb{k}) \cup \text{translate}(\Delta, \Psi, \mathbb{k}') \end{aligned}$$

Figure A.13: Extension of translate to runtime types

Definition Appendix A.5. The cost equations corresponding to a well-formed runtime type \mathbb{k} , written $\text{translate}(\mathbb{k})$ is defined as follows:

$$\text{translate}(\mathbb{k}) \stackrel{\text{def}}{=} \text{translate}(\Delta_{\mathbb{k}}, \text{mapping}(\mathbb{k}), \mathbb{k})$$

Note in Figure A.13 that we also need to extend the definition of `translate` on atoms: we need to re-define the cost of a `get` operation as we now have not only the actual futures in the runtime type but also the ones coming from the static typing of the program (the ones that belong to Ψ). Moreover the initial definition of `translate` for the `get` included an extra adjustment derived from the fact that the state of some virtual machine were unknown, which is not the case at runtime.

Cost Solution. A cost solution Σ is a mapping from cost functions (like $\mathfrak{m}_{\text{net}}^{\Xi}(\bar{v})$ or f_{peak} , where f is a future name) to cost expressions. Basically, this cost solution gives cost expressions of the peak and net cost of all methods in the program, as well as, the peak and net cost of all processes in the runtime type. Following [13, Definition 6], we relate cost solutions and runtime types by assuming the existence of a *cost validation predicate* \models between cost solutions and cost equations as generated by the `translate` function in Section 5:

Definition Appendix A.6. Let C be a set of cost equations $\bigcup_{i=1}^n f_i(\bar{x}_i) = e_i[se_i]$; let Σ be a function that maps all $f_i(\bar{x}_i)$ to cost expressions. Σ is a cost solution of C , noted $\Sigma \models C$, if for all $1 \leq i \leq n$ and all substitutions $[\bar{v}/\bar{x}_i]$ such that $se_i[\bar{v}/\bar{x}_i]$ is true, we have that

$$\Sigma(f_i)(\bar{v}) \geq \Sigma(e_i[\bar{v}/\bar{x}_i])$$

where $\Sigma(e)$ replaces the function calls in e by their value in Σ .

Correctness of the Cost Analysis. The proof of the correction of our cost analysis is done in four steps. First, we relate the cost equation computed at stating time with the runtime version of the types:

Lemma Appendix A.7. *Let $\overline{M} \{ \overline{Fz} ; s' \}$ be a well-typed program, let $\overline{\mathbb{C}}, \mathbb{c}$ be its behavioural type, cn be its initial configuration and \mathbb{k} the runtime type of cn . The for all cost solution Σ with $\Sigma \models \text{translate}(\overline{\mathbb{C}}, \mathbb{c})$, there exists a cost solution Σ' such that $\Sigma' \models \text{translate}(\mathbb{k})$ and $\Sigma(\text{main}()) = \Sigma'(f_{\text{startpeak}}) + \text{alive}(\mathbb{k})$.*

Proof. Let define $\Sigma' = \Sigma[f_{\text{startpeak}} \mapsto \Sigma(\text{main}()) - 1]$. By construction (see the construction of the equation page 15 and Figure A.13), we have that

$$\text{translate}(\mathbb{C}_1 \dots \mathbb{C}_n, \mathbb{c}) = \left\{ \begin{array}{l} \text{translate}(\mathbb{C}_1) \dots \text{translate}(\mathbb{C}_n) \\ \text{main}() = 1 + e_{1,1} \quad [se_1] \\ \vdots \\ \text{main}() = 1 + e_{h_1,1} \quad [se_1] \\ \text{main}() = 1 + e_{1,2} \quad [se_2] \\ \vdots \\ \text{main}() = 1 + e_{h_m,m} \quad [se_m] \end{array} \right. \quad \text{translate}(\mathbb{k}) = \left\{ \begin{array}{l} f_{\text{startpeak}} = e_{1,1} \quad [se_1] \\ \vdots \\ f_{\text{startpeak}} = e_{h_1,1} \quad [se_1] \\ f_{\text{startpeak}} = e_{1,2} \quad [se_2] \\ \vdots \\ f_{\text{startpeak}} = e_{h_m,m} \quad [se_m] \end{array} \right.$$

$$\text{with } \text{translate}(\emptyset, \emptyset, \alpha, (\text{true})\{0\}, \mathbb{c}) = \bigcup_{j=1}^m (se_j)\{e_{1,j}; \dots; e_{h_j,j}\}$$

Hence, by construction, we have that $\Sigma' \models \text{translate}(\mathbb{k})$. Moreover, as $\text{alive}(\mathbb{k}) = 1$, we also have that $\Sigma(\text{main}()) = \Sigma'(f_{\text{startpeak}}) + \text{alive}(\mathbb{k})$. \square

Second, we show that any solution of the cost equation generated from a runtime type is an upper bound for the number of live vm in that runtime type:

Lemma Appendix A.8. *Let $\Gamma \vdash cn : \mathbb{k}$ and assume there is a cost solution Σ such that $\Sigma \models \text{translate}(\mathbb{k})$. Then $\text{alive}(cn) \leq \text{alive}(\mathbb{k}) + \Sigma(f_{\text{startpeak}})$ (recall that $f_{\text{startpeak}}$ is the peak cost of the main process type of \mathbb{k}).*

Proof. Observe that $\text{alive}(cn) = \text{alive}(\mathbb{k})$ by construction of \mathbb{k} . We moreover notice that $\Sigma(f_{\text{startpeak}})$ is positive, since peak costs are by construction always positive integers. We thus have the result. \square

Third, we show that an upper bound computed by our technique is stable w.r.t. runtime type evolution:

Lemma Appendix A.9. *Suppose given two well-formed runtime types \mathbb{k} and \mathbb{k}' such that $\mathbb{k} \rightsquigarrow \mathbb{k}'$. Then, for all cost solution Σ with $\Sigma \models \text{translate}(\mathbb{k})$, there exists Σ' such that: i) $\Sigma' \models \text{translate}(\mathbb{k}')$; and ii) $\text{alive}(\mathbb{k}') + \Sigma'(f_{\text{startpeak}}) \leq \text{alive}(\mathbb{k}) + \Sigma(f_{\text{startpeak}})$.*

Proof. By Case on the rule used in $\mathbb{k} \rightsquigarrow \mathbb{k}'$. If not stated otherwise, we implicitly use the notation used in the considered reduction rule):

- Cases (R-SKIP), (R-READ-FUT-KNOWN), (R-READ-FUT-UNKNOWN), (R-CHOICE), (R-GC), (R-RETURN), (R-ASYNC-CALL-ERR). Because $\text{alive}(\mathbb{k}) = \text{alive}(\mathbb{k}')$, it is clear that taking $\Sigma' = \Sigma$, we have and $\Sigma(f_{\text{startpeak}}) = \Sigma'(f_{\text{startpeak}})$. Hence, we have the result.
- Case (R-ASYNC-CALL). It is easy to remark that $\text{translate}(\mathbb{k}')$ differs from $\text{translate}(\mathbb{k})$ in the following points: i) the entries $[f'_{\text{peak}} = m_{\text{peak}}^{\Xi}(\overline{v}'); f'_{\text{net}} = m_{\text{net}}^{\Xi}(\overline{v}')] (with $\overline{v}' = \Psi(\Xi(\overline{v}))$ and $\Xi = \text{EqRel}(\overline{v}'))$ have been added to $\text{translate}(\mathbb{k}')$; and ii) some references to $m_{\text{peak}}^{\Xi}(\overline{v})$ and $m_{\text{net}}^{\Xi}(\overline{v})$ in $\text{translate}(\mathbb{k})$ have been replaced in $\text{translate}(\mathbb{k}')$ by references to f'_{peak} and f'_{net} . So if we define Σ' as follow, we have the result (as $\text{alive}(\mathbb{k}) = \text{alive}(\mathbb{k}')$):$

$$\Sigma' = \Sigma[f'_{\text{peak}} \mapsto \Sigma(m_{\text{peak}}^{\Xi}(\overline{v}')); f'_{\text{net}} \mapsto \Sigma(m_{\text{net}}^{\Xi}(\overline{v}))]$$

- Case (R-BIND-MTD). Given that \mathbb{c} is the behavior of $m \ o(\overline{v})$ and that $\text{translate}(\mathbb{k}_1 \parallel \mathbb{k}_2) = \text{translate}(\mathbb{k}_1) \cup \text{translate}(\mathbb{k}_2)$, $\text{translate}(\mathbb{k})$ and $\text{translate}(\mathbb{k}')$ produce the same set of equations up to the substitution of formal parameters. By definition of Σ we have the result for $\Sigma' = \Sigma$.

- Case (R-NEW-VM). Let f_i be all possible branches of f . We define

$$\Sigma' = \Sigma[f_{i_peak} \mapsto \Sigma(f_{i_peak}) - 1][f_{i_net} \mapsto \Sigma(f_{i_net}) - 1]$$

The definition of Σ' corresponds to the transfer of the cost +1 to the set of alive virtual machines. Hence we $\text{alive}(\mathbb{k}) = \text{alive}(\mathbb{k}') - 1$ and $\Sigma'(f_{start_peak}) = \Sigma(f_{start_peak}) - 1$, which gives us the result.

- Case (R-RELEASE-VM). If $a' = \top$ we have that $\text{alive}(\mathbb{k}) = \text{alive}(\mathbb{k}') + 1$, if $a' = \perp$ we have that $\text{alive}(\mathbb{k}) = \text{alive}(\mathbb{k}')$, thus taking $\Sigma' = \Sigma$, we have the result.

□

Finally, we can combine all the previous lemma to demonstrate the correctness of our technique.

Theorem 5.1 (Correctness)

Let $\overline{M} \{Fz ; s'\}$ be a well-typed program, let $\overline{\mathbb{C}}, \mathbb{C}$ be its behavioural type and cn be its initial configuration. Let also Σ be a solution of the cost functions $\text{translate}(\overline{\mathbb{C}}, \mathbb{C})$. Then, for every $cn \longrightarrow^* cn'$, $\text{alive}(cn') \leq \Sigma(\text{main}())$.

Proof. Let $n = \Sigma(\text{main}())$.

The argument proceeds by induction on the number of reduction steps:

- At step 0, we translate the program into its runtime equivalent cn . By Lemma Appendix A.7, there exists Σ' such that $\Sigma' \models \text{translate}(\mathbb{k})$ and that we have that $\Sigma'(f_{start_peak}) + \text{alive}(\mathbb{k}) = n$. Moreover, by Lemma Appendix A.8 we have that $\text{alive}(cn) \leq \text{alive}(\mathbb{k}) + \Sigma'(f_{start_peak}) = n$.
- for the inductive case, let Γ , a runtime configuration cn' , a runtime type \mathbb{k} and a cost solution Σ' such that $\Gamma \vdash cn : \mathbb{k}, \Sigma' \models \mathbb{k}$ and $\Sigma'(f_{start_peak}) + \text{alive}(\mathbb{k}) \leq n$ hold. Then by Theorem Appendix A.3, if $cn \rightarrow cn'$, there exists Γ' and a runtime type \mathbb{k}' such that $\mathbb{k} \rightsquigarrow^* \mathbb{k}'$ and $\Gamma' \vdash cn' : \mathbb{k}'$. By Lemma Appendix A.9, there exist a cost solution Σ'' such that $\Sigma'' \models \mathbb{k}'$ and $\text{alive}(\mathbb{k}') + \Sigma''(f_{start_peak}) \leq \text{alive}(\mathbb{k}) + \Sigma'(f_{start_peak})$. By hypothesis and Lemma Appendix A.8, we that $\text{alive}(cn') \leq \text{alive}(\mathbb{k}') + \Sigma''(f_{start_peak}) \leq n$.

□