

# Static analysis of cloud elasticity\*

Abel Garcia   Cosimo Laneve   Michael Lienhardt

Dept. of Computer Science and Engineering, University of Bologna – INRIA Focus

## 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) codes 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 records relevant informations for resource usage (creations, releases, and concurrent operations), (ii) a translation function that takes behavioural types and return cost equations, and (iii) an automatic solver for the the cost equations.

We have experimentally evaluated our technique using a cost analysis solver and we report some results. The technique in this paper may be also applied to estimate (heap) memory consumptions in object-oriented languages.

## 1 Introduction

The analysis of resource usage in a program is of great interest because an accurate assessment could reduce 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 facts, 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, (1) have pricing models that allow one to hire on demand virtual machine instances and paying them for the time they are in use, and (2) have APIs that include instructions for requesting and releasing virtual machine instances.

While it is relatively easy to estimate worst-case costs for sample codes, extrapolating this information for fully real-life complex programs could be

---

\*Partly funded by the EU project FP7-610582 ENVISAGE: Engineering Virtualized Services.

cumbersome and highly error-sensitive. The first attempts about the analysis of resource usages dates back to Wegbreit’s pioneering work in 1975 [20], 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 languages 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, 9]) and to make the connection between programs and closed-form expressions precise (see for instance [19, 13]). We postpone the discussion of the related work in the literature in Section 8.

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, it is used in a very constrained way [4]. On the other hand, cloud computing elasticity requests powerful acquire operations *as well as* release ones. Let us consider the following 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? Under the assumption that one can acquire a virtual machine that has been previously released. A solution to this problem 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 [5]. 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  $x$  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 the cost of methods like `double_release` in a precise way, in Section 4 we associate methods with abstract descriptions that also carry informations about parameter states and their identities. These descriptions are called *behavioural types* and are formally connected to the programs by means of a typing system.

In Section 5 we therefore analyse behavioural type descriptions by translating them in codes that are adequate for a powerful off-the-shelf solver – the `CoFloCo` solver [10]. 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 facts, 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 the two arguments equal then its cost cannot be computed by 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.

In Section 7 we have we report the results of some of our experimental evaluation. In particular, we compute the cost of `double_release` and two implementation of the factorial functions by means of `CoFloCo`.

Our technique target a simple concurrent language with explicit operations of creation and release of resources. The language is defined in Section 2 and we discuss restriction that ease the development of our technique in Section 3. We discuss how these restriction can be removed and outline our correctness proof in Section 6. We deliver concluding remarks in Section 9.

In this paper we use the metaphor of cloud computing and virtual machines. We observe that our technique actually addresses every type of resources that retain operations of acquire (or creation) and release, such as heap usages in concurrent object-oriented languages.

## 2 The language `vm1`

The syntax and the semantics of `vm1` are defined in the following two subsections; the third subsection discusses a number of examples.

**Syntax.** A `vm1` program is a sequence of method definitions  $T \mathfrak{m}(\overline{T x})\{\overline{F y} ; s\}$ , ranged over by  $M$ , plus a main body  $\{\overline{F z} ; s'\}$ . In `vm1` we distinguish between *simple types*  $T$  which are either integers `Int` or virtual machines `vm`, and *types*  $F$ , which also include *future types* `Fut<Int>`. These future types let asynchronous method invocations be typed (see below). The notation  $\overline{T x}$  denotes any finite sequence of *variable declaration*  $T x$ . The elements of the sequence are separated by commas. When we write  $\overline{T x}$ ; we mean a sequence  $T_1 x_1 ; \dots ; T_n x_n$ ; when the sequence is not empty; we mean the empty sequence otherwise.

The syntax of statements  $s$ , expressions with side-effects  $z$  and expressions  $e$  of `vm1` is defined by the following grammar:

$$\begin{aligned}
s & ::= x = z \mid \text{if } e \{ s \} \text{ else } \{ s \} \mid \text{return } e \mid s ; s \mid \text{release}(e) \\
z & ::= e \mid e!m(\bar{e}) \mid e.\text{get} \mid \text{new } vm \\
e & ::= \text{this} \mid se \mid nse
\end{aligned}$$

A (*pure*) expression  $e$  is either an integer constant  $p$ , or a variable  $x$ , or the reserved identifier **this**, or the standard arithmetic, relational and boolean operations. Since our analysis will be parametric with respect to the inputs, we will parse expressions in a careful way. In particular we split them into *size expressions*  $se$ , which are expressions in presburger arithmetics, and *non-size expressions*  $nse$ , which are the other type of expressions. The syntax of size and non-size expressions is the following:

$$\begin{aligned}
nse & ::= p \mid x \mid nse \leq nse \mid nse \text{ and } nse \mid nse \text{ or } nse \mid nse + nse \mid nse - nse \\
& \quad \mid nse \times nse \mid nse/nse \\
se & ::= ve \mid ve \leq ve \mid se \text{ and } se \mid se \text{ or } se \\
ve & ::= p \mid x \mid ve + ve \mid p \times ve \\
p & ::= \text{integer constants}
\end{aligned}$$

An expression  $z$  may change the state of the system. In particular, it may be an *asynchronous* method invocation that does not suspend caller's execution: when the value computed by the invocation is needed then the caller performs a **get** operation. Expressions  $z$  also include **new vm** that creates a new virtual machine. Operations taking place on different virtual machines may execute in parallel, while operations in the same virtual machine interleave their evaluation.

A statement  $s$  may be either one of the standard operations of an imperative language or the **release** operation. The operation **release**( $x$ ) marks the virtual machine  $x$  for disposal. Method invocations performed by a released machine, as well as, creations and releases of machines, always return erroneous values.

In the whole paper, we assume that sequences of declarations  $\overline{T x}$  and method declarations  $\overline{M}$  do not contain duplicate names.

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

$$\begin{aligned}
cn & ::= \epsilon \mid fut(f, v) \mid vm(o, a, p, q) \mid invoc(o, f, m, \bar{v}) \mid cn \ cn \\
p & ::= \{l \mid s\} \mid \text{idle} \\
q & ::= \epsilon \mid \{l \mid s\} \mid q \ q \\
v & ::= \text{integer constants} \mid o \mid f \mid \perp \mid \top \mid \text{err} \\
l & ::= [\dots, x \mapsto v, \dots]
\end{aligned}$$

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  $cn$ ; the empty configuration is denoted by  $\epsilon$ . The transition relation uses two infinite sets of names: *vm names*, ranged over by  $o, o', \dots$  and *future names*, ranged over by  $f, f', \dots$ . We assume there are infinitely many vm names and future names. The function `fresh()` returns either a fresh vm name or a fresh future name; the context will disambiguate between the twos. We also use  $l$  to range over maps from variables to values. The map  $l$  also binds the special name destiny to a future value.

*Runtime values*  $v$  are either integers or vm and future names, or two distinct special values denoting a machine alive ( $\top$ ) or dead ( $\perp$ ), or an erroneous value  $err$ .

The elements of configurations are

- *virtual machines*  $vm(o, a, p, q)$  where  $o$  is a vm name;  $a$  is either  $\top$  or  $\perp$  according to the machine is alive or dead,  $p$  is either  $\{l \mid \varepsilon\}$ , representing a terminated statement, or is the *active process*  $\{l \mid s\}$ , where  $l$  returns the values of local variables and  $s$  is the continuation;  $q$  is a set of processes to evaluate.
- *future binders*  $fut(f, v)$ . When the value  $v$  is  $\perp$  then the value of  $f$  has still to be computed.
- *method invocations*  $invoc(o, f, \mathbf{m}, \bar{v})$ .

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 variables that are stored either in  $l$ . Operations in `vm1` are also defined on the value  $err$ : when one of the arguments is  $err$ , every 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  $\llbracket f \rrbracket_l = f$ .
- $\text{bind}(o, f, \mathbf{m}, \bar{v}) = \{\{\text{destiny} \mapsto f, \bar{x} \mapsto \bar{v} \mid s\}^{o/\text{this}}\}$ , where  $T \mathbf{m}(\overline{T x})\{\overline{T' z}; s\}$  belongs to the program.

The transition relation rules are collected in Figure 1. They define transitions of virtual machines  $vm(o, a, p, q)$  according to the shape of the statement in  $p$ . We focus on rules concerning the method invocations and the management of virtual machines in `vm1`, since the other ones are standard.

(NEW-VM) creates a virtual machine and makes it active. If the virtual machine executing `new vm` has been already released, then the operation returns an error – rule (NEW-VM-ERR). A virtual machine is disposed by means of the operation `release(x)`: this amounts to update its state  $a$  to  $\perp$ .

Rule (ASYNC-CALL) defines asynchronous method invocation  $x = e! \mathbf{m}(\bar{e})$ . This rule creates a fresh future name that is assigned to the identifier  $x$ . The evaluation of the called method is then transferred to the callee virtual machine – rule (BIND-MTD) – and the caller progresses without waiting for callee’s termination. If the caller has been already disposed then the invocation returns  $err$  – rule (ASYNC-CALL-ERR) The invocation binds  $err$  to the future name when either the caller has been released – rule (ASYNC-CALL-ERR) – or the callee machine has been

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

Figure 1: Semantics of `vm1`.

disposed – rule (BIND-MTD-ERR). Rule (READ-FUT) allows the caller to retrieve the value returned by the callee.

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

$$ob(\text{start}, \top, \{[\text{destiny} \mapsto f_{\text{start}}] \mid s\}, \emptyset)$$

where *start* is a special virtual machine and  $f_{\text{start}}$  is a fresh future name. As usual, let  $\longrightarrow^*$  be the reflexive and transitive closure of  $\longrightarrow$ .

**Examples.** In order to illustrate the features of `vm1` we discuss few examples. For every example we also examine the type of 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 ; VM z ;
    if (n==0) { return 1 ; }
    else { z = new VM ; x = z!fact(n-1) ; m = x.get ; release z ; return n*m ;
    }
}
Int cheap_fact(Int n, Int r){
    Fut<Int> x, y ; Int m ; VM z ;
    if (n==0) { return r ; }
    else { z = new VM ; x = z!prod(n,r) ; m = x.get ; release z ;
          y = this!cheap_fact(n-1,m) ; m = y.get ; return m ;
    }
}

```

( $\text{prod}(x,y)$  has been omitted: it just returns  $x*y$ ). The method `fact` is the standard definition of factorial with the recursive invocation `fact(n-1)` performed on a new virtual machine `z`. The caller waits for its result, let it be `m`, then it releases the machine `z` and delivers the value `n*m`. Notice that every vm creation occurs before any release operation. As a consequence, `fact` will create as many virtual machines as the argument `n`. Therefore, in order to be executed, `fact` needs `n` virtual machines (in addition to the one where the method is performed).

While `cheap_fact` also computes the factorial, its behaviour is different. In particular it implements the function

$$f(n,r) = \begin{cases} r & \text{if } n = 0 \\ f(n-1, n*r) & \text{otherwise} \end{cases}$$

which is initially invoked with  $f(n,1)$ . In `cheap_fact` the computation of  $n*r$  is performed on a new virtual machine, which is released *before* the recursive invocation. For this reason, one might always *reuse* the same virtual machine (from a pool). In facts, it turns out that the cost of `cheap_fact` is 1 (in addition to the virtual machine where the method is performed).

The analysis of the two implementations of factorial has been easy because the `release` operation always carries a locally created virtual machine. Yet, in `vm1`, `release` may also apply to method arguments and this is the major source of difficulties for the analysis. Consider for instance the following code:

```

Int first_method() {
    VM x ; Fut<Int> f ; Fut<Int> g ;
    x = new vm ;
    f = x!unknown_method(this) ; f.get ;
    g = x!second_method() ; g.get ;
    release x ;
    return 0 ; }

```

A rough analysis might indicate that `first_method` creates a virtual machine, invokes two methods on that machine, and then releases it. However, this analysis is wrong when `unknown_method` releases its argument(s). For instance, if `unknown_method` releases the argument `this`, the invocation of `second_method`

and the statement `release x` will not be executed. In order to let the cost analysis be compositional, we record the effects of methods on virtual machines in the arguments and we compute the cost analysis accordingly. One might argue that compositionality might be achieved, in this case, by admitting an over- or under-approximate output instead of recording method's effects on arguments. Actually these approximations are not reasonable if one does not consider releases on the arguments:

- an over-approximation might return very imprecise results. Consider, for instance the case when the invocation `x!unknown_method(this)` releases the two arguments `x` and `this`. The over-approximation will compute the cost of `second_method`, even if it will never be called. This means that the analysis might output a very large cost because `second_method` creates a large number of virtual machines.
- an under-approximation might return erroneous results, as it will consider that `release x` will be executed (corresponding to a `-1` cost) while we actually have a null cost.

### 3 Problems, solutions and restrictions

In this section, we present the two new important concepts linked to resource removal, their properties, and which restrictions we put on input programs to keep our analysis from being too complex (indeed, most of the restriction presented here can be relaxed by increasing the complexity of our analysis; this will be discussed in Section 9).

**Method's effects.** As discussed in the foregoing example `first_method`, in order to augment the precision of the cost analysis and to avoid erroneous outputs, our analysis records the effects that a method has on its interface. To keep the formalism as simple as possible, we restrict these effects to be *a set of virtual machines in method's interface*, which is noted `R` in Section 4. This simplicity – `R` is a set – has a price: the virtual machines in the interface of a method that will be released in *every execution path* is always the same. That is, a method as

```
Int ugly_release(VM x1, ... , VM xn) {
  VM x ; Fut<Int> f ;
  x = new vm ; f = x!release_all(x1, ... , xn) ; release x ; f.get ;
  return 0; }
```

cannot be handled by our analysis (`release_all(x1, ... , xn)` disposes the virtual machines `x1, ... , xn`). In facts, since `release x` is performed before the synchronisation with `release_all` – statement `f.get` –, the method `release_all` can be stopped at any point of its execution, thus making the set of released virtual machines non-determinate. In order to ban methods like `ugly_release`, we constrain definitions as follows:



1. every method invocation is synchronized within caller's body. In this way every effect of a method is computed before its termination.
2. it is not possible to release a machine that is executing a method.

**Virtual machines' identity.** Removals of virtual machines may have side effects on other machines. The following method `double_release` illustrates the point:

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

This method takes two machines and simply releases them. This means that, when the virtual machines in input are active *and different*, the cost of `double_release` is -2. In fact, this is the case of the following method `user1`:

```
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 ; }
```

An easy computation returns a (net) cost of 0 for `user1`. However, the cost of `double_release` is *not always* -2. For example, consider the `user2` above. This method creates one virtual machine (cost +1) and invokes `double_release` with a duplicated argument: in this case the cost of `double_release` is -1, not -2, (and the cost of `user2` is 0). Said more explicitly, the cost of a method depends on the identity of its arguments, not only on their states (alive or dead).

It is also worth to notice that the identity of arguments has impact on method's effects. Consider the following code snippet:

```
VM x = new vm; VM y = new vm;
x!double_release(x,y);
```

This code creates two new virtual machines and releases them by invoking `double_release`. The awkward point is that the callee machine coincides with the first argument. Therefore `double_release` will fail to release the second argument `y`.

In order to comply with these issue we annotate every operation that may have side effects with *identity sharing informations*, noted  $\Theta$  in Section 4 – and generate a different cost function for every  $\Theta$ . For instance, we have two (net) cost functions for the `double_release` method:

1. one where we consider `x` and `y` to be different: here its cost will be -2 when the two machines are alive,
2. one where we consider the two parameters to be equal: here its cost will be -1 when the two machines are alive.

In addition, in order to avoid the problematic issue of the code above, we forbid the callee machine to have the same identity of the (other) arguments of the invocation.

**Release of carrier’s machine.** To simplify our analysis, we admit releases of the `this` machine – the carrier – to be the last statement (before `return`). While this restriction may be easily dropped, it avoids duplications of rules in Figure 3 to deal with the fact that the virtual machine `this` could be dead. In any case, we notice that this constraint does not affect the expressive power of the language.

## 4 The behavioral type system of `vm1`

Our analysis uses abstract descriptions, called *behavioural types*, which are intermediate codes highlighting the features of `vm1` programs that are relevant for the analysis of resources in Section 5. These types support compositional reasonings and are associated to programs by means of a type system.

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  $\mathbf{S}, \mathbf{S}', \mathbf{R}, \dots$ , and sets of future names will be ranged over by  $\mathbf{F}, \mathbf{F}', \dots$ . The syntactic rules are the following

$\circ$	::=	$- \mid \alpha$	basic value
$\mathfrak{t}$	::=	$\alpha \mid \partial \mid \perp \mid \top$	vm value
$se$	::=	<i>integer constant</i> $\mid x \mid (\mathfrak{t} \leq \perp) \mid (\mathfrak{t} \leq \top) \mid se \text{ op}' se$	size expression
$\text{op}'$	::=	$+ \mid - \mid = \mid \leq \mid \geq \mid \wedge \mid \vee$	linear operation
$\mathfrak{r}, \mathfrak{s}$	::=	$\circ \mid se$	typing value
$\mathfrak{z}$	::=	$(\circ, \alpha, \Theta, \mathbf{R}) \mid \circ$	future value
$\mathfrak{x}$	::=	$- \mid \mathbf{F}\mathfrak{t} \mid f \mid \mathfrak{z}$	extended value
$\mathfrak{a}$	::=	$\mathbf{0} \mid \nu\alpha[\Theta] \mid \alpha^\vee[\Theta] \mid \nu f : \mathfrak{m} \alpha(\overline{\mathfrak{s}}) \mid f^\vee[\Theta, \mathbf{R}]$	atom
$\mathfrak{c}$	::=	$\mathfrak{a} \triangleright \Gamma \mid \mathfrak{a} \ddagger \mathfrak{c} \mid \mathfrak{c} + \mathfrak{c} \mid (se)\{\mathfrak{c}\}$	behavioural type

Behavioural types express creations of virtual machines ( $\nu\alpha$ ) and their removal ( $\alpha^\vee$ ), method invocations ( $\nu f : \mathfrak{m} \alpha(\overline{\mathfrak{s}})$ ) and corresponding retrieval of the value ( $f^\vee[\Theta, \mathbf{R}]$ ) and the conditionals (respectively  $(se)\{\mathfrak{c}\} + (\neg se)\{\mathfrak{c}'\}$  or  $\mathfrak{c} + \mathfrak{c}'$ , according to the boolean guard is a size expressions that depends on the arguments of a method or not). Behavioural types also carry *vm name environments*  $\Theta, \Theta', \dots$ . These environments map vm names to extended values  $\mathbf{F}\mathfrak{t}$ , which are called *vm states* in the following. This feature is new in behavioural types and, as discussed in Section 3, it is needed during the analysis to manage the identity of methods’ arguments. We will provide additional examples in the rest of the paper to explain vm name environments and their use.

In order to have a more precise type of continuations, the leaves of behavioural types are labelled with *environments*, ranged over by  $\Gamma, \Gamma', \dots$ . Environments are maps from method names  $\mathfrak{m}$  to terms  $\alpha(\overline{\mathfrak{r}}) : \circ, \mathbf{R}$ , from variables to extended values  $\mathfrak{x}$ , from future names to future values, and from vm names to

vm states. These environments are used in the typing proofs and are dropped in the final types (method types and the main statement type).

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  $\Gamma \vdash \mathfrak{m} \alpha(\bar{\mathfrak{r}}) : \mathbb{o}, \mathbb{R}$  for methods.
- $\Gamma \vdash_{\mathfrak{s}} z : \mathbb{x}, \mathfrak{c} \triangleright \Gamma'$  for expressions with side effects  $z$ , where  $\mathbb{x}$  is the value,  $\mathfrak{c}$  is the behavioural type for  $z$  and  $\Gamma'$  is the environment  $\Gamma$  *with updates* of variables and future names.
- $\Gamma \vdash_{\mathfrak{s}} s : \mathfrak{c}$ , in this case the updated environments are inside the behavioural type  $\Gamma'$ , in correspondence of every branch of its.

Since  $\Gamma$  is a function, we use the standard predicates  $x \in \text{dom}(\Gamma)$  or  $x \notin \text{dom}(\Gamma)$ . Moreover we define

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

The following operation and notations are going to be used:

- vm values  $\mathfrak{t}$  are partially ordered by the relation  $\leq$  defined by  $\partial \leq \top$  and  $\partial \leq \perp$ . 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  $\partial \sqcap \alpha$  is not defined.
- the *update of a vm name environment*  $\Theta$  with respect to  $f$  and  $\Gamma$  (we remind that  $\Theta$  is defined on vm names only), written  $\Theta \searrow_f^{\Gamma} \Gamma$ , returns a vm name environment defined as follows:

$$\Theta \searrow_f^{\Gamma} \Gamma \stackrel{\text{def}}{=} [\alpha \mapsto (\mathbf{F}' \setminus \{f\})(\mathfrak{t} \sqcap \mathfrak{t}')]^{\alpha \in \text{dom}(\Theta), \Theta(\alpha) = \mathbf{F}\mathfrak{t}, \Gamma(\alpha) = \mathbf{F}'\mathfrak{t}'}$$

- the *multihole contexts*  $\mathcal{C}[\ ]$  defined by the following syntax:

$$\mathcal{C}[\ ] ::= [\ ] \mid \mathfrak{a} \mathfrak{s} \mathcal{C}[\ ] \mid \mathcal{C}[\ ] + \mathcal{C}[\ ] \mid (\text{se})\{\mathcal{C}[\ ]\}$$

and, whenever  $\mathfrak{c} = \mathcal{C}[\mathfrak{a}_1 \triangleright \Gamma_1] \cdots [\mathfrak{a}_n \triangleright \Gamma_n]$ , then  $\mathfrak{c}[x \mapsto \mathbb{x}]$  is defined as  $\mathcal{C}[\mathfrak{a}_1 \triangleright \Gamma_1[x \mapsto \mathbb{x}]] \cdots [\mathfrak{a}_n \triangleright \Gamma_n[x \mapsto \mathbb{x}]]$ .

The type systems for expressions, expressions with side effects and statements are reported in Figures 2 and 3. It is worth to notice that the type system for expressions 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. As regards the rules for statements, we discuss (T-INVOKE), (T-GET), (T-RELEASE), and (T-NEW) because the other ones are straightforward.

Rule (T-INVOKE) types method invocations  $e!\mathfrak{m}(\bar{e})$  by means of a new future name  $f$  that is associated to the method name, the vm name of the callee and the arguments. The relevant point is the value of  $f$  in the updated environment.

$$\begin{array}{c}
\text{(T-VAR)} \quad \frac{x \in \text{dom}(\Gamma)}{\Gamma \vdash x : \Gamma(x)} \quad \text{(T-PRIMITIVE)} \quad \Gamma \vdash \mathbf{p} : \mathbf{p} \quad \text{(T-OP)} \quad \frac{\Gamma \vdash e_1 : se_1 \quad \Gamma \vdash e_2 : se_2}{\Gamma \vdash e_1 \text{ op}' e_2 : se_1 \text{ op}' se_2} \quad \text{T-UNIT} \quad \frac{\Gamma \vdash e : se}{\Gamma \vdash e : -} \\
\\
\text{(T-OP-UNIT)} \quad \frac{\Gamma \vdash e_1 : - \quad \text{or} \quad \Gamma \vdash e_2 : - \quad \text{or} \quad \text{op} \in \{*, /\}}{\Gamma \vdash e_1 \text{ op} e_2 : -} \quad \text{(T-PURE)} \quad \frac{\Gamma \vdash e : \mathbf{x}}{\Gamma \vdash e : \mathbf{x}, \mathbf{0} \triangleright \Gamma} \\
\\
\text{(T-METHOD-SIG)} \quad \frac{\Gamma(\mathbf{m}) = \alpha(\bar{\mathbf{r}}) : \mathfrak{o}, \mathbf{R} \quad \bar{\beta} \subseteq fv(\alpha, \bar{\mathbf{r}}, \mathfrak{o}) \quad \sigma \text{ is a vm renaming such that } \mathfrak{o} \notin fv(\alpha, \bar{\mathbf{r}}) \text{ implies } \sigma(\mathfrak{o}) \text{ fresh}}{\Gamma \vdash \mathbf{m} \sigma(\alpha)(\sigma(\bar{\mathbf{r}})) : \sigma(\mathfrak{o}), \sigma(\mathbf{R})}
\end{array}$$

Figure 2: Typing rules for expressions

This value contains the returned value, the vm name of the callee, a vm name environment and the set of vm names that the method is going to remove. The vm name environment records the state of vm names when the method is invoked and it will be used when the method is synchronized to update the environment of the synchronisation (see rule (T-GET)) with the changes performed by methods in parallel. It is important to observe that the environment returned by the judgment is updated with informations about vm names released by the method: every such name will contain  $f$  in its state. Next we discuss the constraints in the second line and third line of the premise of (T-INVOKE). Assuming that the callee has not been already released ( $\Gamma(\alpha) \neq \mathbf{F}\perp$ ), there are two cases:

- (i) either  $\Gamma(\alpha) = \emptyset\top$  or  $\alpha$  is the caller object, 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 subclasses, 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 because the nondeterminism of the execution makes the analysis too much imprecise (in our view). In facts, as discussed in Section 3, in case (ii.a), we cannot determine what subset of  $\mathbf{R}$  will be actually released (because of the parallel method releasing  $\alpha$ ). In case (ii.b), being  $\alpha$  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  $\mathbf{R}$  that will be released. It is worth to notice that the constraint  $\mathbf{R} = \emptyset$  of case (ii) enforces a programming style that reduces uncertainty and provides a more precise cost analysis. The constraint in the third line of the premise of (T-INVOKE) applies the same constraint of the callee to the other invocations in parallel and to the object executing  $e!m(\bar{e})$ : it is not possible to remove a vm name that is a callee of a parallel method.

Rule (T-GET) defines the synchronisation with a method invocation that corresponds to a future  $f$ . Let  $(\mathfrak{o}, \alpha, \Theta, \mathbf{R})$  be the value of  $f$  in the environment.

$$\begin{array}{c}
\text{(T-INVOKE)} \\
\frac{\Gamma \vdash e : \alpha \quad \Gamma \vdash \bar{e} : \bar{\mathbb{S}} \quad \Gamma \vdash \mathbf{m} \alpha(\bar{\mathbb{S}}) : \circ, \mathbf{R} \quad \Gamma \vdash \mathbf{this} : \alpha' \\
\Gamma(\alpha) \neq \mathbf{F}\perp \quad \text{and } ((\Gamma(\alpha) \neq \circ\top \text{ and } \alpha \neq \alpha') \text{ implies } \mathbf{R} = \emptyset) \\
\mathbf{R} \cap (\{\alpha'\} \cup \{\beta \mid f' \in \text{dom}(\Gamma) \text{ and } \Gamma(f') = (\circ', \beta, \Theta, \mathbf{R}')\}) = \emptyset \\
f \text{ fresh} \quad \Gamma' = \Gamma[\beta \mapsto (\{f\} \cup \mathbf{F}')\mathbf{t}]^{\beta \in \mathbf{R}, \Gamma(\beta) = \mathbf{F}'\mathbf{t}}}{\Gamma \vdash_{\mathbf{S}} \mathbf{elm}(\bar{e}) : f, \nu f : \mathbf{m} \alpha(\bar{\mathbb{S}}) \triangleright \Gamma'[f \mapsto (\circ, \alpha, \Gamma|_{\mathbf{S} \cup \text{fv}(\alpha, \bar{\mathbb{S}})}, \mathbf{R})]} \\
\\
\text{(T-INVOKE-BOT)} \quad \frac{\Gamma \vdash e : \alpha \quad \Gamma(\alpha) = \mathbf{F}\perp \quad f \text{ fresh}}{\Gamma \vdash_{\mathbf{S}} \mathbf{elm}(\bar{e}) : f, 0 \triangleright \Gamma'[f \mapsto \_]} \quad \text{(T-NEW)} \quad \frac{\beta \text{ fresh}}{\Gamma \vdash_{\mathbf{S}} \mathbf{new} \mathbf{vm} : \beta, \nu \beta[\Gamma|_{\mathbf{S}}] \triangleright \Gamma[\beta \mapsto \circ\top]} \\
\\
\text{(T-GET)} \quad \frac{\Gamma \vdash x : f \quad \Gamma \vdash f : (\circ, \alpha, \Theta, \mathbf{R}) \\
\Theta' = \Theta \not\vee \Gamma \quad \mathbf{R}' = \text{fv}(\circ) \setminus (\mathbf{R} \cup \text{dom}(\Theta)) \\
\Gamma' = \Gamma[\beta \mapsto \emptyset\perp]^{\beta \in \mathbf{R}}[\beta' \mapsto \Theta'(\alpha)]^{\beta' \in \mathbf{R}'}[f \mapsto \circ]}{\Gamma \vdash_{\mathbf{S}} x.\mathbf{get} : \circ, f'[\Theta', \mathbf{R}] \triangleright \Gamma'} \quad \text{(T-GET-DONE)} \quad \frac{\Gamma \vdash x : f \quad \Gamma \vdash f : \circ}{\Gamma \vdash_{\mathbf{S}} x.\mathbf{get} : \circ, 0 \triangleright \Gamma} \\
\\
\text{(T-RELEASE)} \quad \frac{\Gamma \vdash x : \alpha}{\Gamma \vdash_{\mathbf{S}} \mathbf{release}(x) : \alpha'[\Gamma|_{\mathbf{S} \cup \{\alpha\}}] \triangleright \Gamma[\alpha \mapsto \emptyset\perp]} \quad \text{(T-ASSIGN-VAR)} \quad \frac{\Gamma(x) = \mathbf{x} \quad \Gamma \vdash_{\mathbf{S}} z : \mathbf{x}', \mathbf{c}}{\Gamma \vdash_{\mathbf{S}} x = z : \mathbf{c}[x \mapsto \mathbf{x}']} \\
\\
\text{(T-IF)} \quad \frac{\Gamma \vdash e : se \quad \Gamma \vdash_{\mathbf{S}} s_1 : \mathbf{c}_1 \quad \Gamma \vdash_{\mathbf{S}} s_2 : \mathbf{c}_2}{\Gamma \vdash_{\mathbf{S}} \mathbf{if} e \{s_1\} \mathbf{else} \{s_2\} : (se)\{\mathbf{c}_1\} + (\neg se)\{\mathbf{c}_2\}} \quad \text{(T-IF-ND)} \quad \frac{\Gamma \vdash e : \_ \quad \Gamma \vdash_{\mathbf{S}} s_1 : \mathbf{c}_1 \quad \Gamma \vdash_{\mathbf{S}} s_2 : \mathbf{c}_2}{\Gamma \vdash_{\mathbf{S}} \mathbf{if} e \{s_1\} \mathbf{else} \{s_2\} : \mathbf{c}_1 + \mathbf{c}_2} \\
\\
\text{(T-SEQ)} \quad \frac{\Gamma \vdash_{\mathbf{S}} s_1 : \mathcal{C}[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n] \quad \Gamma_i \vdash_{\mathbf{S}} s_2 : \mathbf{c}'_i}{\Gamma \vdash_{\mathbf{S}} s_1; s_2 : \mathcal{C}[\mathbf{a}_1 \mathbin{\&} \mathbf{c}'_1] \cdots [\mathbf{a}_n \mathbin{\&} \mathbf{c}'_n]} \quad \text{(T-RETURN)} \quad \frac{\Gamma \vdash e : \circ \quad \Gamma \vdash \mathbf{destiny} : \circ' \\
\circ \in \mathbf{S} \cup \{\_\} \quad \text{if and only if } \circ = \circ'}{\Gamma \vdash_{\mathbf{S}} \mathbf{return} e : 0 \triangleright \Gamma[\mathbf{destiny} \mapsto \circ']}
\end{array}$$

Figure 3: Type rules for expressions with side effects and statements.

Since  $\mathbf{R}$  defines the resources of the caller that are released, we record in the returned environment  $\Gamma'$  that these resources are no more available.  $\Gamma'$  also records the state of the returned vm name, if it is a virtual machine that has been created by the method of  $f$ . This state is the same of the callee vm name (which may have been updated since the invocation), namely the value of  $(\Theta \not\vee \Gamma)(\alpha)$ . As regards the behavioural type  $f'[\Theta \not\vee \Gamma, \mathbf{R}]$  of  $x.\mathbf{get}$ , it carries two arguments (*i*) the vm name environment of the invocation updated with the invocation of parallel methods and (*ii*) the vm names to be released. These two arguments will let us disable the removal of names in the cost analysis when these names are removed by methods in parallel. In this case, the removal is computed in the caller method (and therefore it counts  $-1$ , instead of  $-1$  in every method in parallel) – see definition of  $\mathbf{translate}$  in case of  $f'[\Theta \not\vee \Gamma, \mathbf{R}]$  in Section 5.

Rule (T-RELEASE) models the removal of a vm name  $\alpha$ . Notice that the be-

havioural type is not just  $\alpha^\vee$ , which should correspond to a  $-1$  in the cost analysis. Rather, it is a “conditional” removal because the name may be removed by some method in parallel, or even may have been already removed when the method has been invoked. This is the reason for the presence of the vm name environment  $[\Gamma]_{\mathcal{S} \cup \{\alpha\}}$ . We also observe that this rule applies only when the caller has not been already released.

For reasons similar to the rules discussed above, (T-NEW) adds an environment  $\Gamma|_{\mathcal{S}}$  to the type  $\nu\beta$  expressing a new resource. In facts, this environment will make the cost increase by  $1$  provided the caller method has not been already released – see definition of `translate` in case of  $\nu\beta[\Gamma|_{\mathcal{S}}]$  in Section 5.

The type system of `vml` is completed with the rules for method declarations and programs:

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

$$\begin{array}{c}
\text{(T-PROGRAM)} \\
\frac{\Gamma \vdash \overline{M} : \overline{\mathcal{C}} \quad \Gamma \vdash_{\text{start}} s : \mathcal{C}[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n]}{\Gamma \vdash \overline{M} \{ \overline{F x} ; s \} : \overline{\mathcal{C}}, \mathcal{C}[\mathbf{a}_1 \triangleright \emptyset] \cdots [\mathbf{a}_n \triangleright \emptyset]}
\end{array}$$

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  $\circ$  may be either  $\_$  or a vm name in  $\mathbf{S}$  or a fresh vm name. In this last case, the premise  $(\Gamma_i(\gamma) = \Gamma_j(\gamma))^{i,j \in 1..n, \gamma \in \mathcal{S} \cup \text{fv}(\circ)}$  guarantees that every branch in the returned behavioural type creates a new vm name and, by rule (T-RETURN), the chosen vm name must be always the same.

We display behavioural types at work by typing codes of Section 2 and 3. Actually, the following types do not abstract a lot from codes’ details because the programs of the previous sections have been designed for highlighting the issues of our technique.

The behavioural types of `fact` and `cheap_fact` are the following ones

<pre> <b>fact</b> <math>\alpha(n)</math> {   (<math>n==0</math>){ 0 }   + (<math>n&gt;0</math>){ <math>\nu\beta[\alpha \mapsto \emptyset\alpha]</math>;               <math>\nu x : \mathbf{fact} \beta(n-1)</math>;               <math>x^\vee[\alpha \mapsto \emptyset\alpha, \beta \mapsto \emptyset\top]</math>;               <math>\beta^\vee[\alpha \mapsto \emptyset\alpha, \beta \mapsto \emptyset\top]</math>;             } - , { } </pre>	<pre> <b>cheap_fact</b> <math>\alpha(n, -)</math> {   (<math>n==0</math>){ 0 }   + (<math>n&gt;0</math>){ <math>\nu\beta[\alpha \mapsto \emptyset\alpha]</math>;               <math>\nu x : \mathbf{prod} \beta(n, -)</math>;               <math>x^\vee[\alpha \mapsto \emptyset\alpha, \beta \mapsto \emptyset\top]</math>;               <math>\beta^\vee[\alpha \mapsto \emptyset\alpha, \beta \mapsto \emptyset\top]</math>;               <math>\nu y : \mathbf{cheap\_fact} \beta(n-1, -)</math>;               <math>y^\vee[\alpha \mapsto \emptyset\alpha, \beta \mapsto \emptyset\perp]</math>;             } - , { } </pre>
---	---

We notice that the types record the order between recursive invocations and

releases of machines, which is different in the two codes for factorial.

The behavioural types of `double_release` and `user1` are the following ones

<pre>double_release α(β, γ) {   β<sup>✓</sup>[α ↦ ∅α, β ↦ ∅β, γ ↦ ∅γ];   γ<sup>✓</sup>[α ↦ ∅α, β ↦ ∅⊥, γ ↦ ∅γ]; } -, {β, γ}</pre>	<pre>user1 α( ) {   νβ[α ↦ ∅α] ; νγ[α ↦ ∅α];   ν f : double_free α(β, γ);   f<sup>✓</sup>[α ↦ ∅α, β ↦ ∅⊥, γ ↦ ∅⊥]; } -, { }</pre>
---	---

It is worth to notice that the releases  $\beta^\checkmark$  and  $\gamma^\checkmark$  in `double_release` are conditioned by the values of  $\beta$  and  $\gamma$  when the method is invoked. In facts, in case of `user1` these values are  $\emptyset\top$  and this will mean that the cost of  $f^\checkmark[\alpha \mapsto \emptyset\alpha, \beta \mapsto \emptyset\top, \gamma \mapsto \emptyset\top]$  will be -2.

## 5 The analysis of behavioural types

The behavioural types returned by system in Section 4 are used to compute the resource elasticity of a `vml` program. This computation is performed by an off-the-shelf solver – the `CoFloCo` solver [10] – and, in this section, we discuss the translation of a behavioural type program into a set of *recurrence relations* that fed the solver.

Basically, our translation maps method types into functions from parameters to cost, where

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

There are two difficulties that entangle our translation that pertain to method invocations: the management of arguments' identities and of argument's values.

**Argument's identities.** Consider the code

```
Int free() { release(this) ; return(0) ; }

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

The behavioural types of these methods are

$$\begin{aligned} & \mathbf{free} \alpha() \{ \alpha^\vee [\alpha \mapsto \emptyset \top] \} -, \{ \alpha \} \\ & \mathbf{m} \alpha(\beta, \gamma) \{ \nu f : \mathbf{free} \beta() ; \gamma^\vee [\Theta] ; f^\vee [\Theta'] \} -, \{ \beta, \gamma \} \end{aligned}$$

where

$$\begin{aligned} \Theta &= \alpha \mapsto \emptyset \alpha, \beta \mapsto \{ \mathbf{free} \} \beta, \gamma \mapsto \emptyset \gamma \\ \Theta' &= \alpha \mapsto \emptyset \alpha, \beta \mapsto \{ \mathbf{free} \} \beta, \gamma \mapsto \emptyset \perp \end{aligned}$$

We notice that, in the type of  $\mathbf{m}$ , there is not enough information to determine whether  $\gamma^\vee$  will have a cost equal to  $-1$  or  $0$ . In facts, while in typing rules of methods the arguments are assumed to be pairwise different – see rule (T-METHOD) –, it is not the case for invocations. For instance, if  $\mathbf{m}$  is invoked with two arguments that are equal –  $\beta = \gamma$  – then  $\gamma$  is going to be released by the invocation  $\mathbf{free}(\beta)$  and therefore it counts  $0$ . To solve this problem of arguments' identity, we refine even more the translation of method, which now depends on an equivalence relation telling which of the vm names in parameter are actually equal or not. Hence, in general, our method  $\mathbf{m}$  will be translated in four cost functions:  $\mathbf{m}_{\text{peak}}^{\{1\},\{2\}}(x, y)$  and  $\mathbf{m}_{\text{net}}^{\{1\},\{2\}}(x, y)$ , which correspond to the invocations where  $x \neq y$ , and  $\mathbf{m}_{\text{peak}}^{\{1,2\}}(x)$  and  $\mathbf{m}_{\text{net}}^{\{1,2\}}(x)$ , which correspond to the invocations where  $x = y$ . (The equivalence relation in the superscript never mention **this**, which is also an argument, because we assume that **this** is always different from the other arguments, see below.)

The following function  $\mathbf{EqRel}$  computes the equivalence relation corresponding to a specific method call;  $\mathbf{EqRel}$  takes a tuple of vm names and returns an equivalence relation on indices of the tuple:

$$\mathbf{EqRel}(\alpha_1, \dots, \alpha_n) = \bigcup_{i \in 1..n} \{ \{j \mid \alpha_j = \alpha_i\} \}$$

Let  $\mathbf{EqRel}(\alpha_1, \dots, \alpha_n)(\beta_1, \dots, \beta_n)$  be the tuple  $(\beta_{i_1}, \dots, \beta_{i_k})$ , where  $i_1, \dots, i_k$  are *canonical representatives* of the sets in  $\mathbf{EqRel}(\alpha_1, \dots, \alpha_n)$  (we take the vm name with the least index in every set). We observe that, by definition,  $\mathbf{EqRel}(\alpha_1, \dots, \alpha_n)(\alpha_1, \dots, \alpha_n)$  is a tuple of pairwise different vm names (in  $\alpha_1, \dots, \alpha_n$ ).

We will always assume that  $\{1\} \in \mathbf{EqRel}(\alpha_1, \dots, \alpha_n)$ . Since  $\alpha_1$  always represents the **this** object, this constraint enforces that the other arguments of invocations are always different from **this** (see the discussion in the paragraph about virtual machines identity in Section 3).

**Argument's values.** Consider the code

```
Int m(VM x) {
  Fut<Int> f ; Fut<Int> g ; VM y ; VM z ; Int u ;
  y = new vm ; z = new vm ;
  f = x!free() ; g = y!foo(x, z) ;
  u = g.get ; u = f.get ; return 0 ;
}
```



where **free** is as above and **foo** is unspecified (we assume that **foo** does not release any machine). The behavioural types of **m** is

$$\mathbf{m} \alpha(\beta) \{ \nu \gamma [\Theta_\gamma] ; \nu \gamma' [\Theta_{\gamma'}]; \nu f : \mathbf{free} \beta(\ ) ; \nu g : \mathbf{foo} \gamma(\beta, \gamma') ; g^\vee [\Theta_2]; f^\vee [\Theta_1] \} \\ \_ , \{ \beta \}$$

where

$$\begin{aligned} \Theta_\gamma &= [\alpha \mapsto \emptyset \alpha, \beta \mapsto \emptyset \beta] \\ \Theta_{\gamma'} &= [\alpha \mapsto \emptyset \alpha, \beta \mapsto \emptyset \beta, \gamma \mapsto \emptyset \top] \\ \Theta_2 &= [\alpha \mapsto \emptyset \alpha, \beta \mapsto \{f\} \beta, \gamma \mapsto \emptyset \top, \gamma' \mapsto \emptyset \top] \\ \Theta_1 &= [\alpha \mapsto \emptyset \alpha, \beta \mapsto \{f\} \beta, \gamma \mapsto \emptyset \top, \gamma' \mapsto \emptyset \top] \end{aligned}$$

The arguments of the invocation of **foo** are  $\beta$  and  $\gamma'$  and, in order to compute its cost, it is necessary to instantiate these arguments with actual vm values. This instantiation is straightforward for  $\gamma'$ : there is no concurrent future releasing it, so its state is  $\top$ . However, this is not the case for  $\beta'$ : as specified in  $\Theta_2$ , its vm state is tagged with a nonempty set of futures because **free** is concurrently releasing it. Hence, we need to translate this state into a simple one ( $\perp$ ,  $\top$ , or  $\partial$ ). In facts, the value of the first argument of **foo** may be  $\partial$  if  $\beta$  was originally either  $\top$  or  $\partial$ , or it may be  $\perp$  if  $\beta$  was already released. To deal with this case, we introduce an operator on vm values  $\beta \downarrow$  whose meaning is the one just described:

$$\beta \downarrow = \begin{cases} \partial & \text{if } \beta = \top \text{ or } \beta = \partial \\ \perp & \text{else} \end{cases}$$

Since it is not possible to know the value of  $\beta$  during the translation, we use  $\beta \downarrow$  in the syntax of our cost expressions. This term will be evaluated during the solving process.

To conclude, the translation of vm states into vm values (extended with  $\alpha \downarrow$ ), written  $\mathbf{Ft} \downarrow$ , is defined as follows:

$$\mathbf{Ft} \downarrow \stackrel{def}{=} \begin{cases} \mathfrak{t} & \text{if } \mathbf{F} = \emptyset \\ \partial & \text{if } \mathbf{F} \neq \emptyset \text{ and } \mathfrak{t} = \top \\ \perp & \text{if } \mathbf{F} \neq \emptyset \text{ and } \mathfrak{t} = \perp \\ \alpha \downarrow & \text{if } \mathbf{F} \neq \emptyset \text{ and } \mathfrak{t} = \alpha \end{cases}$$

and we write  $(\mathbf{F}_1 \mathfrak{t}_1, \dots, \mathbf{F}_n \mathfrak{t}_n) \downarrow$  for  $(\mathbf{F}_1 \mathfrak{t}_1 \downarrow, \dots, \mathbf{F}_n \mathfrak{t}_n \downarrow)$ .

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

1.  $\Delta$  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  $x$ . Therefore we can use  $\Delta$  also as a substitution operation.

2.  $\Psi$  is the *translation environment* which stores temporary information about futures that are active (unsynchronized);
3.  $\alpha$  is the vm name of the virtual machine of the current behaviour type;
4.  $(se)\{\bar{e}\}$  is sequence of costs of the current execution branch: the size expression  $se$  stores all the conditions corresponding to the current execution branch, while  $\bar{e}$  is the sequence of (over-approximated) costs that branch takes during its execution.
5. the behavioural type that is translated; it may be either  $a$ ,  $c$  or  $\overline{C}$ .

The auxiliary functions below let us reduce the number of cases of the definition of **translate**:

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

This 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.

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

This function is used when a virtual machine is released (in correspondence of atoms  $\beta^\vee$ ). The release is effectively computed – value  $-1$  – only when the virtual machine that is executing the code is alive ( $\alpha = \top$ ).

The **translate** function also uses the *merge operation*, noted  $\Theta[\Delta]$ , that takes a vm name environment  $\Theta$  and an equivalence relation on vm names  $\Delta$  and returns a *substitution*. We remind that vm name environments have been introduced in Section 4 to manage identities of arguments in method calls. Take, for instance, the atom  $f^\vee[\Theta, R]$  within a behavioural type that binds  $f$  to  $\text{foo } \alpha(\beta, \gamma)$ . Assume to evaluate this behavioural type for  $m_{\text{peak}}^\Delta$  where  $\Delta = \{\beta, \gamma\}$ . That is, the two arguments are actually identical. What are the values of  $\beta$  and  $\gamma$  for evaluating  $\text{foo}_{\text{peak}}^\Delta$  and  $\text{foo}_{\text{net}}^\Delta$ ? Well, we have

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

For example, let  $\Theta = [\alpha \mapsto \emptyset\alpha, \beta \mapsto \emptyset\beta, \gamma \mapsto \emptyset\gamma]$  and  $\Delta(\beta) = \Delta(\gamma) = \beta$ . We expect that a value for the item 2 above is  $\emptyset\beta$  and the substitution of the item 3 is  $\{\emptyset\beta, \emptyset\beta / \beta, \gamma\}$ . Formally, the operation returning the value for 2 is noted  $\otimes^\beta$

below (it applies to vm values and vm states) and the operation returning the substitution of item 3 is the merge operation.

$$\mathfrak{t} \otimes^\alpha \mathfrak{t}' = \begin{cases} \perp & \text{if } \mathfrak{t} = \perp \text{ or } \mathfrak{t}' = \perp \\ \alpha \downarrow & \text{if } (\mathfrak{t} = \gamma \downarrow \text{ and } \mathfrak{t}' \neq \perp) \text{ or } (\mathfrak{t} \neq \perp \text{ and } \mathfrak{t}' = \gamma \downarrow) \\ \alpha & \text{if } \mathfrak{t} = \gamma \text{ and } \mathfrak{t}' = \gamma' \end{cases}$$

$$\mathbf{F}_1 \mathfrak{t}_1 \otimes^\alpha \mathbf{F}_2 \mathfrak{t}_2 = (\mathbf{F}_1 \cup \mathbf{F}_2)(\mathfrak{t}_1 \otimes^\alpha \mathfrak{t}_2)$$

$$\Theta[\Delta] : \left[ \alpha \mapsto \bigotimes^{\Delta(\alpha)} \{ \Theta(\beta) \mid \beta \in \text{dom}(\Theta) \text{ and } \Delta(\beta) = \Delta(\alpha) \} \right]^{\alpha \in \text{dom}(\Theta)}$$

We notice that the definition of  $\otimes^\alpha$  is not necessary for vm values as  $\partial$  or  $\top$  because we merge vm names whose image by  $\Theta$  can only be either  $\mathbf{F}\beta$  or  $\mathbf{F}\beta \downarrow$  or  $\mathbf{F}\perp$ . As a notational remark, we observe that the substitution  $\Theta[\Delta]$  is noted as a map  $[\alpha_1 \cdots \mathbf{F}_1 \mathfrak{t}_1, \cdots, \alpha_n \cdots \mathbf{F}_n \mathfrak{t}_n]$  instead of the standard notation  $\{\mathbf{F}_1 \mathfrak{t}_1, \cdots, \mathbf{F}_n \mathfrak{t}_n / \alpha_1, \dots, \alpha_n\}$ . These two notations are clearly equivalent: we prefer the former one because it will let us to write  $\Theta[\Delta](\alpha)$  or even  $\Theta[\Delta](\alpha_1, \dots, \alpha_n)$  with the obvious meanings.

Every preliminary notion is in place for defining **translate**. We begin with the translation of atoms.

$$\text{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}; e\}, \mathfrak{a}) = \begin{cases} (\Psi, (se)\{\bar{e}; e\}) & \text{when } \mathfrak{a} = 0 \\ (\Psi, (se)\{\bar{e}; e; e + \mathbf{CNEW}(\mathfrak{t})\}) & \text{when } \mathfrak{a} = \nu\beta[\Theta] \\ & \text{and } \Theta[\Delta](\alpha) = \mathbf{F}\mathfrak{t} \\ (\Psi, (se)\{\bar{e}; e; e + \mathbf{CREL}(\mathfrak{t})\}) & \text{when } \mathfrak{a} = \beta^\vee[\Theta] \\ & \text{and } \Theta[\Delta](\beta) = \mathbf{F}\mathfrak{t} \\ (\Psi[f \mapsto \mathfrak{m} \Delta(\beta)(\overline{s\bar{e}}, \Delta(\overline{\beta}))], (se)\{\bar{e}; e; e + f\}) & \text{when } \mathfrak{a} = (\nu f = \mathfrak{m} \beta(\overline{s\bar{e}}, \overline{\beta})) \\ (\Psi \setminus f, (se)\{\bar{e}; e\}\sigma; (\bar{e}; e)\sigma' + \sum_{\gamma \in \Delta(\mathbf{R}), \Theta[\Delta](\gamma) = \mathbf{F}\mathfrak{t}, \mathbf{F} \neq \emptyset} \mathbf{CREL}(\mathfrak{t})\}) & \text{when } \mathfrak{a} = f^\vee[\Theta, \mathbf{R}] \\ & \text{where} \\ & \Psi(f) = \mathfrak{m} \beta(\overline{s\bar{e}}, \overline{\beta}) \\ & \mathbf{EqRel}(\beta, \overline{\beta}) = \Xi \\ & \sigma = \{\mathfrak{m}_{\text{peak}}^\Xi(\overline{s\bar{e}}, \Theta[\Delta](\Xi(\beta, \overline{\beta})) \Downarrow) / f\} \\ & \sigma' = \{\mathfrak{m}_{\text{net}}^\Xi(\overline{s\bar{e}}, \Theta[\Delta](\Xi(\beta, \overline{\beta})) \Downarrow) / f\} \end{cases}$$

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  $f^\vee[\Theta, \mathbf{R}]$ . In this last case, let  $(se)\{\bar{e}; e\}$  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  $\sigma$  – and the resources used upon the termination of (the method corresponding to)  $f$  – the *net cost* used in the substitution  $\sigma'$ . In particular, this last value has to be decreased by the number of the resources released by

the method. This the purpose of the addend  $\sum_{\gamma \in \Delta(\mathbf{R}), \Theta[\Delta](\gamma) = \mathbf{Ft}, \mathbf{F} \neq \emptyset} \mathbf{CREL}(\dagger)$  that remove 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  $\mathbf{m}_{\text{peak}}$  and  $\mathbf{m}_{\text{net}}$  that are invoked are those corresponding to the equivalence relation of the tuple  $(\beta, \bar{\beta})$ .

The definition of **translate** for behavioural type is given below. It follows straightforwardly by composing the definitions of the atoms. We also observe that, in this case, the sequences of costs are sets.

$$\mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathfrak{c}) = \begin{cases} \{(se')\{\bar{e}'\}\} & \text{when } \mathfrak{c} = \mathfrak{a} \triangleright \emptyset \text{ and } \mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathfrak{a}) = (\Psi', (se')\{\bar{e}'\}) \\ C'' & \text{when } \mathfrak{c} = \mathfrak{a} \ddagger \mathfrak{c}' \text{ and } \mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathfrak{a}) = (\Psi', \{(se')\{\bar{e}'\}\}) \\ & \text{and } \mathbf{translate}(\Delta, \Psi', \alpha, (se')\{\bar{e}'\}, \mathfrak{c}') = (\Psi'', C'') \\ C' \cup C'' & \text{when } \mathfrak{c} = \mathfrak{c}_1 + \mathfrak{c}_2 \text{ and } \mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathfrak{c}_1) = (\Psi', C') \\ & \text{and } \mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathfrak{c}_2) = (\Psi'', C'') \\ C' & \text{when } \mathfrak{c} = (se')\{\mathfrak{c}'\} \text{ and } \mathbf{translate}(\Delta, \Psi, \alpha, (se \wedge se')\{\bar{e}\}, \mathfrak{c}') = (\Psi', C') \end{cases}$$

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

$$\mathbf{translate}(\mathbf{m} \alpha_1(\bar{x}, \alpha_2, \dots, \alpha_n) \{ \mathfrak{c} \} : \mathbb{O}, \mathbf{R}) = \bigcup_{\Xi \in \mathcal{P}} \mathbf{translate}(\Xi, \mathbf{m} \alpha_1(\bar{x}, \alpha_2, \dots, \alpha_n) \{ \mathfrak{c} \} : \mathbb{O}, \mathbf{R})$$

where  $\mathbf{translate}(\Xi, \mathbf{m} \alpha_1(\bar{x}, \alpha_2, \dots, \alpha_n) \{ \mathfrak{c} \} : \mathbb{O}, \mathbf{R})$  is defined as follows. Let

$$\Delta = \{ \{ \alpha_{i_1}, \dots, \alpha_{i_m} \} \mid \{ i_1, \dots, i_m \} \in \Xi \}$$

and

$$\mathbf{translate}(\Delta, \emptyset, \alpha_1, (\mathbf{true})\{0\}, \mathfrak{c}) = \bigcup_{i=1}^n (se_i)\{e_{1,i}; \dots; e_{h_i,i}\}$$

Then

$$\mathbf{translate}(\Xi, \mathbf{m} \alpha_1(\bar{x}, \alpha_2, \dots, \alpha_k) \{ \mathfrak{c} \} : \mathbb{O}, \mathbf{R}) = \begin{cases} \mathbf{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = 0 & [\alpha_1 = \perp] \\ \mathbf{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{1,1} & [se_1 \wedge \alpha_1 \neq \perp] \\ \vdots & \vdots \\ \mathbf{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_1,1} & [se_1 \wedge \alpha_1 \neq \perp] \\ \mathbf{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{1,2} & [se_2 \wedge \alpha_1 \neq \perp] \\ \vdots & \vdots \\ \mathbf{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_n,n} & [se_n \wedge \alpha_1 \neq \perp] \\ \mathbf{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = 0 & [\alpha_1 = \perp] \\ \mathbf{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = \mathbf{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \dots, \alpha_n]) & [\alpha_1 = \emptyset] \\ \mathbf{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{n_1,1} & [se_1 \wedge \alpha_1 = \top] \\ \vdots & \vdots \\ \mathbf{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{n_n,n} & [se_n \wedge \alpha_1 = \top] \end{cases}$$

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

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

We show the output of `translate` when applied to the behavioural types of `double_release`, `user1`, and `user2` that we have described in Section 4 (well, the behavioural type of `user2` has not been shown: it is left as an exercise). Since `double_release` has two arguments, we generate two sets of equations, as discussed above. On the contrary, methods `user1` and `user2` carry one argument and therefore there is one set of equations only. In order to ease the reading, we omit the equivalence classes of arguments that label function names: the reader may grasp them from the number of arguments. For the same reason, we represent a partition  $\{\{1\}, \{2\}, \{3\}\}$  corresponding to vm names  $\alpha_1$ ,  $\alpha_2$  and  $\alpha_3$  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}_{\text{double\_release}} \} : -, \{\alpha_2, \alpha_3\}) = \\ &\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. \\ &\text{translate}([\alpha_1, \alpha_2], \text{double\_release } \alpha_1(\alpha_2) \{ \mathbb{c}_{\text{double\_release}} \} : -, \{\alpha_2\}) = \\ &\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{CNEW}(\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}$$

$$\begin{aligned}
& \text{translate}([\alpha_1], \text{user1 } \alpha_1() \{ \mathfrak{c}_{\text{user1}} \} : -, \{ \}) = \\
& \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) & [\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) & [\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) & [\alpha_1 = \top]
\end{array} \right. \\
& \text{translate}([\alpha_1], \text{user2 } \alpha_1() \{ \mathfrak{c}_{\text{user2}} \} : -, \{ \}) = \\
& \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}$$

## 6 Outline of the proof of correctness

The proof of correctness of our technique is long even if almost standard (see [11] for a similar proof). In this section we overview it by highlighting the main difficulties.

The first part of the proof 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. It is worth to observe that we cannot hope to demonstrate a statement guaranteeing type-preservation (the types of  $cn$  and  $cn'$  are equal) because our types are “behavioural”. However, it is critical for the correctness of the cost analysis that there is a relation between the type of  $cn$ , let it be  $\mathfrak{c}$ , and the type of  $cn'$ , let it be  $\mathfrak{c}'$ .

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.

Once 1 and 2 above have been defined, it is possible to demonstrate (let  $\rightsquigarrow^*$  be the reflexive and transitive closure of  $\rightsquigarrow$ ):

**Theorem 6.1** (Subject Reduction). *Let  $cn$  be a configuration of a `vm1` program and let  $\mathfrak{c}$  be its behavioural type. If  $cn \rightarrow cn'$  then there is  $\mathfrak{c}'$  typing  $cn'$  such that  $\mathfrak{c} \rightsquigarrow^* \mathfrak{c}'$ .*

The proof of this theorem is by case on the reduction rule applied and it is usually not complex because the relation  $\rightsquigarrow$  mimics the `vml` transitions in Section 2.

The second part of the proofs relies on the definition of the notion of *direct cost of a behavioural type* (of a configuration), which is the number of virtual machines occurring in the type. We also observe that the number of alive virtual machines in a configuration is identical to the direct cost of the corresponding a behavioural type. Then it is also necessary to define

3. the extension of the function `translate` to compute the recurrence relations for behavioural types of configurations. These recurrence relations let us compute the *peak cost of a behavioural type* (of a configuration).

The proofs of the following two properties are preliminary to the correctness of our technique:

**Lemma 6.2** (Basic Cost Inclusion). *The direct cost of a behavioural type of a configuration is less or equal to its peak cost.*

**Lemma 6.3** (Reduction Cost Inclusion). *If  $\mathfrak{c} \rightsquigarrow \mathfrak{c}'$  then the peak cost of  $\mathfrak{c}'$  is less or equal to the peak cost of  $\mathfrak{c}$ .*

It is important to observe that the proofs of Lemmas 6.2 and 6.3 are given using the (theoretical) solution of recurrence relations in [2]. This lets us to circumvent possible errors in implementations of the theory, such as `CoFloCo` [10] or `PUBS` [2]. Given the basic cost and reduction cost inclusions, we can demonstrate the correctness theorem for our technique.

**Theorem 6.4** (Correctness). *Let  $\overline{M} \{ \overline{Fz} ; s' \}$  be a well-typed program and let  $\overline{\mathbb{C}}, \mathfrak{c}$  be its behavioural type. Let also  $n$  be a solution of the function `translate`( $\overline{\mathbb{C}}, \mathfrak{c}$ ). Then  $n$  is an upper bound of the number of virtual machines used during the execution of  $cn$ .*

The proof outline is as follows. Since the cost of the initial configuration  $cn$  is the direct cost of  $\mathfrak{c}$  then, by Lemma 6.2, this value is less or equal to the peak cost of  $\mathfrak{c}$ . Let  $n$  be a solution of this cost. The argument proceeds by induction on the number of reduction steps:

- for the base case, when the program doesn't reduce, it turns out that  $n \geq 1$ ;
- for the inductive case, let  $cn \rightarrow cn'$ . By applying Theorem 6.1 and Lemma 6.3, one derives that  $n$  is bigger than the peak cost of the behavioural type of  $cn'$ . Thus, by Lemma 6.2, we have that  $n$  is larger than the number of alive virtual machines in  $cn'$ .

## 7 Integration with a cost analysis tool and experiments

In this section we discuss technical details about the translation of the recurrence relations in Section 5 into the cost analysis tool we use – the CoFloCo analyser [10] – and we examine the outputs obtained for the running examples of this paper. It is worth to notice that instead of targeting CoFloCo, we might also target the PUBS analyser [2], which also has similar recurrence relations for input.

In order to comply with usual input formats of tools, we need to define encodings for vm values and for the functions CNEW and CREL. We therefore define

- $\top$  is modelled by 1,  $\partial$  is modelled by 2,  $\perp$  is modelled by 3, and  $\alpha$  by  $\alpha$ . As regards  $\alpha \downarrow$ , it is modelled by the conditional value  $[\alpha = 3]3 + [1 \leq \alpha \leq 2]2$ ;
- the auxiliary functions CNEW and CREL are translated in recurrence relations as follows:

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

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

We begin our experiments with the translation of `double_release` when used by `user1`. In this case, the arguments of `double_release` are all different, therefore we use  $double\_release_{peak}(\alpha_1, \alpha_2, \alpha_3)$  and  $double\_release_{net}(\alpha_1, \alpha_2, \alpha_3)$ . We write these functions in CoFloCo as `doubleRelease123_peak(A,B,C)` and `doubleRelease123_net(A,B,C)`. The input for the cost analyzer is shown in Figure 4 and, in order to evaluate it, we need to specify a so-called *entry point*. This entry point has the following format:

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

where the first argument always represents the state of the carrier virtual machine. The following table report the output of CoFloCo with respect to the entry point. We observe that the computed cost is exactly what we anticipated

Entry Point	Cost
<code>entry(user1_net(1): [])</code> .	Maximum cost of <code>user1_net(1)</code> : 0
<code>entry(user1_peak(1): [])</code> .	Maximum cost of <code>user1_peak(1)</code> : 2

Table 1: Costs of programs `double_release` and `user1`

in Section 4.

Figure 5 describes the program of `double_release` when used by `user2`. In this case, the arguments of `double_release` are equal, therefore we use



```

eq(doubleRelease123_peak(A,B,C), 0, [], [A = 3]).
eq(doubleRelease123_peak(A,B,C), 0, [], [A < 3]).
eq(doubleRelease123_peak(A,B,C), 0, [crel(B)], [A < 3]).
eq(doubleRelease123_peak(A,B,C), 0, [crel(B), crel(C)], [A < 3]).

eq(doubleRelease123_net(A,B,C), 0, [], [A = 3]).
eq(doubleRelease123_net(A,B,C), 0, [doubleRelease123_peak(A,B,C)], [A = 2]).
eq(doubleRelease123_net(A,B,C), 0, [crel(B), crel(C)], [A = 1]).

eq(user1_peak(A), 0, [], [A = 3]).
eq(user1_peak(A), 0, [], [A < 3]).
eq(user1_peak(A), 0, [cnew(A)], [A < 3]).
eq(user1_peak(A), 0, [cnew(A),cnew(A)], [A < 3]).
eq(user1_peak(A), 0, [cnew(A),cnew(A),doubleRelease123_peak(A, 1, 1)], [A < 3]).
eq(user1_peak(A), 0, [cnew(A),cnew(A),doubleRelease123_net(A, 1, 1)], [A < 3]).

eq(user1_net(A), 0, [], [A = 3]).
eq(user1_net(A), 0, [user1_peak(A)], [A = 2]).
eq(user1_net(A), 0, [cnew(A),cnew(A),doubleRelease123_net(A, 1, 1)], [A = 1]).

```

Figure 4: Cost equations of `double_release` and `user1` in CoFloCo format

$double\_release_{peak}(\alpha_1, \alpha_2)$  and  $double\_release_{net}(\alpha_1, \alpha_2)$ . We write these functions in CoFloCo as `doubleRelease12_peak(A,B)` and `doubleRelease12_net(A,B)`.

```

eq(doubleRelease12_peak(A,B), 0, [], [A = 3]).
eq(doubleRelease12_peak(A,B), 0, [], [A < 3]).
eq(doubleRelease12_peak(A,B), 0, [crel(B)], [A < 3]).
eq(doubleRelease12_peak(A,B), 0, [crel(B), crel(3)], [A < 3]).

eq(doubleRelease12_net(A,B), 0, [], [A = 3]).
eq(doubleRelease12_net(A,B), 0, [doubleRelease12_peak(A,B)], [A = 2]).
eq(doubleRelease12_net(A,B), 0, [crel(B), crel(3)], [A = 1]).

eq(user2_peak(A), 0, [], [A = 3]).
eq(user2_peak(A), 0, [], [A < 3]).
eq(user2_peak(A), 0, [cnew(A)], [A < 3]).
eq(user2_peak(A), 0, [cnew(A),doubleRelease12_peak(A, 1)], [A < 3]).
eq(user2_peak(A), 0, [cnew(A),doubleRelease12_net(A, 1)], [A < 3]).

eq(user2_net(A), 0, [], [A = 3]).
eq(user2_net(A), 0, [user2_peak(A)], [A = 2]).
eq(user2_net(A), 0, [cnew(A),doubleRelease12_net(A, 1)], [A = 1]).

```

Figure 5: Cost equations of `double_release` and `user2` in CoFloCo format

The table below shows the output of the cost analyzer for the given equations, where, again, we consider only the case when the first argument is alive, that is, it is equal to 1. As before, the cost is exactly what we informally computed in Section 4.

Entry Point	Cost
<code>entry(user2_net(1):[]).</code>	0
<code>entry(user2_peak(1):[]).</code>	1

Table 2: Costs of programs `double_release` and `user2`

We conclude this section by discussing the cost of the two factorial programs `fact` and `cheap_fact` discussed in Section 2. The list of cost equations generated by `translate` is given in Figure 6 below.

```

eq(fact_peak(A,B), 0, [], [A = 3]).
eq(fact_peak(A,B), 0, [], [B = 0]).
eq(fact_peak(A,B), 0, [], [B > 0]).
eq(fact_peak(A,B), 0, [cnew(A)], [B > 0]).
eq(fact_peak(A,B), 0, [cnew(A), fact_peak(1, B-1)], [B > 0]).
eq(fact_peak(A,B), 0, [cnew(A), fact_net(1, B-1)], [B > 0]).
eq(fact_peak(A,B), 0, [cnew(A), fact_net(1, B-1), crel(1)], [B > 0]).

eq(fact_net(A,B), 0, [], [A = 3]).
// eq(fact_net(A,B), 0, [fact_peak(A,B)], [A = 2]).
eq(fact_net(A,B), 0, [cnew(A), fact_net(1, B-1), crel(1)], [A = 1, B > 0]).
eq(fact_net(A,B), 0, [], [A = 1, B = 0]).

eq(cheap_fact_peak(A,B,C), 0, [], [A = 3]).
eq(cheap_fact_peak(A,B,C), 0, [], [B = 0]).
eq(cheap_fact_peak(A,B,C), 0, [], [B > 0]).
eq(cheap_fact_peak(A,B,C), 0, [cnew(A)], [B > 0]).
eq(cheap_fact_peak(A,B,C), 0, [cnew(A), prod_peak(1,B,C)], [B > 0]).
eq(cheap_fact_peak(A,B,C), 0, [cnew(A), prod_net(1,B,C)], [B > 0]).
eq(cheap_fact_peak(A,B,C), 0, [cnew(A), prod_net(1,B,C), crel(1)], [B > 0]).
eq(cheap_fact_peak(A,B,C), 0, [cnew(A), prod_net(1,B,C), crel(1),
    cheap_fact_peak(1, B-1, D)], [B > 0]).
eq(cheap_fact_peak(A,B,C), 0, [cnew(A), prod_net(1,B,C), crel(1),
    cheap_fact_net(1, B-1, D)], [B > 0]).

eq(cheap_fact_net(A,B,C), 0, [], [A = 3]).
//eq(cheap_fact_net(A,B,C), 0, [cheap_fact_peak(A,B,C)], [A = 2]).
eq(cheap_fact_net(A,B,C), 0, [cnew(A), prod_net(1,B,C), crel(1),
    cheap_fact_net(1, B-1, D)], [A = 1, B > 0]).
eq(cheap_fact_net(A,B,C), 0, [], [A = 1, B = 0]).

eq(prod_peak(A,B,C), 0, [], []).
eq(prod_net(A,B,C), 0, [], []).

```

Figure 6: Cost equations of `fact` and `cheap_fact` in CoFloCo format

We notice that Figure 6 has a couple of equations commented out. These equations are obtained by the `translate` function in Section 5. However it is not allowed as input in CoFloCo because those equations lead to mutually recursive chains, which is banned by this tool. Yet, in this case, the exclusion of this

equation does not affect the result because we assume that the virtual machine executing either `fact` or `cheap_fact` is always alive. Table 3 shows the output of CoFloCo.

Entry Point	Cost
<code>entry(fact_net(1,B) : [B&gt;=0]) .</code>	0
<code>entry(fact_peak(1,B) : [B&gt;=0]) .</code>	$\max([B, 1])$
<code>entry(cheap_fact_net(1,B) : [B&gt;=0]) .</code>	0
<code>entry(cheap_fact_peak(1,B) : [B&gt;=0]) .</code>	1

Table 3: Costs of programmms `fact` and `cheap_fact`

As in the previous examples the *net cost* is equal to 0 in both cases, because every created virtual machine is released before the end of the program. In the case of the *peak cost*, for method `fact` the number of virtual machines will depend on the depth of the recursion, in this case, given by parameter `B`. On the other hand for method `cheap_fact` at each step the created VM is released before the recursive call producing a *peak cost* equal to 1.

## 8 Related Work

After the pioneering work by Wegbreit in 1975 [20] that developed a technique for deriving upper-bounds costs of functional programs, a number of techniques for the cost analysis have been developed. These techniques may be divided into two categories.

The first category, that we call *classical techniques*, addresses cost analysis in three steps: (i) extracting relevant informations out of the original programs, e.g. abstracting data structures to their size and assigning a cost to every program expression, (ii) converting the abstract program into recurrence relations, and (iii) solving the cost equations with an automatic tool. Very powerful classical technique tools are [2, 12, 3, 9, 10] that produce very accurate upper bounds expressions for various kinds of programs of different complexity. We refer to [10] for a comparison among some of these techniques. The main drawback of these techniques is the lack of compositionality: it is difficult to scale the analysis to large programs and the translation from the original program to the recurrence relations is always unclear.

The second category, that we call *amortized techniques*, uses the technique based on type systems and amortized analysis developed for functional programs by Hofmann [15]. This approach is highly compositional because of the use of types, and more suitable for formal verification because the connection between the original program and the cost equations can be demonstrated by a standard subject-reduction theorem [15, 19, 16, 17, 13]. We follow this technique in the present work.

A common feature of classical and amortized techniques is that they analyze *cumulative resources*, that is resources that *do not decrease* during the

execution of the programs. This is the case, for example, for execution time, number of operations, memory (without an explicit `free` operation), etc. As we discussed, this assumption eases the analysis because it permits to compute over-approximated cost. On the contrary, the presence of *explicit or implicit* release operation entangles the analysis, as already discussed in [6] where a memory cost analysis is proposed for languages with garbage collection. It is worth to say that the scenario of [6] is not difficult because, by definition of garbage collection, released memories are 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*. As discussed in Section 5, the first refers to the worst case cost for an operation to complete, while the second refers to the cost of an operation after completion. For cumulative analysis the two notions coincide; however, 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 Albert *et al.* in [4] have analysed the cost of a language with explicit releases. However, the release operation they consider is used in a very restrictive way: releases can only be performed over locally acquired resources. This constraint ensures having no partially negative costs when analyzing block sequences thus maintaining the restriction of non-negative cost models.

Most cost analysis techniques usually address sequential program. Only few works also address concurrent programs [1, 5, 14]. In this cases, to reduce the imprecision of the analysis caused by the nondeterminism, the authors of [1, 5] use a clever technique for determining parallel codes, called *may-happen-in-parallel* [8]. However no one of these contributions consider a concurrent language with a powerful `release` operation that allows one disable the resources taken in input. In facts, without this operation, one can model the cost by simply aggregating the sets of operations that can occur in parallel, as in [5], and all the theoretical development is much easier.

## 9 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 informations about resource usages in programs, called behavioural types; an automatic translation that transforms these types into cost expressions; the application of solvers, like CoFloCo [10], 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.

For the future, we consider three lines of work. First, we will complete the technical development of this paper by delivering full proofs of correctness of our technique. Second, 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  $\mathbb{R}$  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  $\mathbb{R}$  with a condition specifying when such effect is yielded. Third, we intend to implement our analysis targeting a programming language with a formal model as ABS [18].

## References

- [1] Elvira Albert, Puri Arenas, Samir Genaim, Miguel Gómez-Zamalloa, and German Puebla. Cost analysis of concurrent OO programs. In *Proceedings of Programming Languages and Systems - APLAS 2011*, volume 7078 of *Lecture Notes in Computer Science*, pages 238–254. Springer, 2011.
- [2] Elvira Albert, Puri Arenas, Samir Genaim, and Germán Puebla. Automatic inference of upper bounds for recurrence relations in cost analysis. In *Proceedings SAS 2008*, volume 5079 of *Lecture Notes in Computer Science*, pages 221–237. Springer, 2008.
- [3] Elvira Albert, Puri Arenas, Samir Genaim, German Puebla, and Damiano Zanardini. Cost analysis of object-oriented bytecode programs. *Theoretical Computer Science*, 413(1):142–159, 2012.
- [4] Elvira Albert, Jesús Correas, and Guillermo Román-Díez. Non-Cumulative Resource Analysis. In *Proceedings of (TACAS 2015)*, Lecture Notes in Computer Science. Springer, 2015. To appear.
- [5] Elvira Albert, Jess Correas, and Guillermo Romn-Dez. Peak cost analysis of distributed systems. In *Proceedings of SAS 2014*, volume 8723 of *Lecture Notes in Computer Science*, pages 18–33. Springer, 2014.
- [6] Elvira Albert, Samir Genaim, and Miguel Gómez-Zamalloa. Parametric inference of memory requirements for garbage collected languages. *SIGPLAN Not.*, 45(8):121–130, 2010.
- [7] Diego Esteban Alonso-Blas and Samir Genaim. On the limits of the classical approach to cost analysis. In *Static Analysis*, pages 405–421. Springer, 2012.
- [8] Rajkishore Barik. Efficient computation of may-happen-in-parallel information for concurrent java programs. In *Proceedings of LCPC 2005*, volume 4339 of *Lecture Notes in Computer Science*, pages 152–169. Springer, 2006.
- [9] Marc Brockschmidt, Fabian Emmes, Stephan Falke, Carsten Fuhs, and Jürgen Giesl. Alternating runtime and size complexity analysis of integer

- programs. In *Proceedings of TACAS 2014*, volume 8413 of *Lecture Notes in Computer Science*, pages 140–155. Springer, 2014.
- [10] Antonio Flores Montoya and Reiner Hähnle. Resource analysis of complex programs with cost equations. In Jacques Garrigue, editor, *12th Asian Symposium on Programming Languages and Systems, Singapore*, volume 8858 of *Lecture Notes in Computer Science*, pages 275–295. Springer, November 2014.
  - [11] Elena Giachino, Cosimo Laneve, and Michael Lienhardt. A Framework for Deadlock Detection in ABS. *Software and Systems Modeling*, 2015. To Appear.
  - [12] Sumit Gulwani, Krishna K Mehra, and Trishul Chilimbi. Speed: precise and efficient static estimation of program computational complexity. In *ACM SIGPLAN Notices*, volume 44, pages 127–139. ACM, 2009.
  - [13] Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. Multivariate amortized resource analysis. In *ACM SIGPLAN Notices*, volume 46, pages 357–370. ACM, 2011.
  - [14] Jan Hoffmann and Zhong Shao. Automatic static cost analysis for parallel programs, 2015. [Online; accessed 11-February-2015].
  - [15] Martin Hofmann and Steffen Jost. Static prediction of heap space usage for first-order functional programs. In *ACM SIGPLAN Notices*, volume 38, pages 185–197. ACM, 2003.
  - [16] Martin Hofmann and Steffen Jost. Type-based amortised heap-space analysis. In *Programming Languages and Systems*, pages 22–37. Springer, 2006.
  - [17] Martin Hofmann and Dulma Rodriguez. Efficient type-checking for amortised heap-space analysis. In *Computer Science Logic*, pages 317–331. Springer, 2009.
  - [18] Einar Broch Johnsen, Reiner Hähnle, Jan Schäfer, Rudolf Schlatte, and Martin Steffen. ABS: A core language for abstract behavioral specification. In Bernhard Aichernig, Frank S. de Boer, and Marcello M. Bonsangue, editors, *Proc. 9th International Symposium on Formal Methods for Components and Objects (FMCO 2010)*, volume 6957 of *Lecture Notes in Computer Science*, pages 142–164. Springer, 2011.
  - [19] Wei ngan Chin, Huu Hai Nguyen, Shengchao Qin, and Martin Rinard. Memory usage verification for oo programs. In *In SAS 05*, pages 70–86. Springer, 2005.
  - [20] Ben Wegbreit. Mechanical program analysis. *Communications of the ACM*, 18(9):528–539, 1975.