

Time complexity of concurrent programs

Elena Giachino, Einar Broch Johnsen, Cosimo Laneve, Ka I Pun

► **To cite this version:**

Elena Giachino, Einar Broch Johnsen, Cosimo Laneve, Ka I Pun. Time complexity of concurrent programs: a technique based on behavioural types. FACS 2015, Oct 2015, Niterói, Rio de Janeiro, Brazil. <hal-01229068>

HAL Id: hal-01229068

<https://hal.inria.fr/hal-01229068>

Submitted on 16 Nov 2015

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

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

Time complexity of concurrent programs[★]

– a technique based on behavioural types –

Elena Giachino¹, Einar Broch Johnsen², Cosimo Laneve¹, and Ka I Pun²

¹ Dept. of Computer Science and Engineering, University of Bologna – INRIA FOCUS

² Dept. of Informatics, University of Oslo

Abstract. We study the problem of automatically computing the time complexity of concurrent object-oriented programs. To determine this complexity we use intermediate abstract descriptions that record relevant information for the time analysis (cost of statements, creations of objects, and concurrent operations), called *behavioural types*. Then, we define a translation function that takes behavioural types and makes the parallelism explicit into so-called *cost equations*, which are fed to an automatic off-the-shelf solver for obtaining the time complexity.

1 Introduction

Computing the cost of a sequential algorithm has always been a primary question for every programmer, who learns the basic techniques in the first years of their computer science or engineering curriculum. This cost is defined in terms of the input values to the algorithm and over-approximates the number of the executed instructions. In turn, given an appropriate abstraction of the CPU speed of a runtime system, one can obtain the expected computation time of the algorithm.

The computational cost of algorithms is particularly relevant in mainstream architectures, such as the cloud. In that context, a service is a concurrent program that must comply with a so-called *service-level agreement* (SLA) regulating the cost in time and assigning penalties for its infringement [3]. The service provider needs to make sure that the service is able to meet the SLA, for example in terms of the end-user response time, by deciding on a resource management policy and determining the appropriate number of virtual machine instances (or containers) and their parameter settings (e.g., their CPU speeds). To help service providers make correct decisions about the resource management before actually deploying the service, we need static analysis methods for resource-aware services [6]. In previous work by the authors, cloud deployments expressed in the formal modeling language ABS [8] have used a combination of cost analysis and simulations to analyse resource management [1], and a Hoare-style proof system to reason about end-user deadlines has been developed for sequential executions [7]. In contrast, we are here interested in statically estimating the computation time of concurrent services deployed on the cloud with a given dynamic resource management policy.

[★] Supported by the EU projects FP7-610582 ENVISAGE: Engineering Virtualized Services (<http://www.envisage-project.eu>).

Technically, this paper proposes a behavioural type system expressing the resource costs associated with computations and study how these types can be used to soundly calculate the time complexity of parallel programs deployed on the cloud. To succinctly formulate this problem, our work is developed for `tml`, a small formally defined concurrent object-oriented language which uses asynchronous communications to trigger parallel activities. The language defines virtual machine instances in terms of dynamically created concurrent object groups with bounds on the number of cycles they can perform per time interval. As we are interested in the concurrent aspects of these computations, we abstract from sequential analysis in terms of a statement `job(e)`, which defines the number of processing cycles required by the instruction – this is similar to the `sleep(n)` operation in Java.

The analysis of behavioural types is defined by translating them in a code that is adequate for an off-the-shelf solver – the CoFloCo solver [4]. As a consequence, we are able to determine the computational cost of algorithms in a parametric way with respect to their inputs.

Paper overview. The language is defined in Section 2 and we discuss restrictions that ease the development of our technique in Section 3. Section 4 presents the behavioural type system and Section 5 explains the analysis of computation time based on these behavioural types. In Section 6 we outline our correctness proof of the type system with respect to the cost equations. In Section 7 we discuss the relevant related work and in Section 8 we deliver concluding remarks.

2 The language `tml`

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

Syntax. A `tml` program is a sequence of method definitions $T \overline{m(T \overline{x})} \{ \overline{F y} ; s \}$, ranged over by M , plus a main body $\{ \overline{F z} ; s' \}$ **with** k . In `tml` we distinguish between *simple types* T which are either integers `Int` or classes `Class` (there is just one class in `tml`), and *types* F , which also include *future types* `Fut<T>`. These future types let asynchronous method invocations be typed (see below). The notation $\overline{T x}$ denotes any finite sequence of *variable declarations* $T x$. 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 possibly empty sequence otherwise.

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

$$\begin{aligned} s &::= x = z \mid \text{if } e \{ s \} \text{ else } \{ s \} \mid \text{job}(e) \mid \text{return } e \mid s ; s \\ z &::= e \mid e!m(\overline{e}) \mid e.m(\overline{x}) \mid e.\text{get} \mid \text{new Class with } e \mid \text{new local Class} \\ e &::= \text{this} \mid se \mid nse \end{aligned}$$

A statement s may be either one of the standard operations of an imperative language or the job statement `job(e)` that delays the continuation by e cycles of the machine executing it.

An expression z may change the state of the system. In particular, it may be an *asynchronous* method invocation of the form $e!m(\overline{e})$, which does not suspend the caller's

execution. When the value computed by the invocation is needed, the caller performs a *non-blocking get* operation: if the value needed by a process is not available, then an awaiting process is scheduled and executed, i.e., *await-get*. Expressions z also include standard synchronous invocations $e.m(\vec{e})$ and `new local Class`, which creates a new object. The intended meaning is to create the object in the same machine – called *cog* or *concurrent object group* – of the caller, thus sharing the processor of the caller: operations in the same virtual machine interleave their evaluation (even if in the following operational semantics the parallelism is not explicit). Alternatively, one can create an object on a different cog with `new Class with e` thus letting methods execute in parallel. In this case, e represents the capacity of the new cog, that is, the number of cycles the cog can perform per time interval. We assume the presence of a special identifier `this.capacity` that returns the capacity of the corresponding cog.

A *pure* expression e can be the reserved identifier `this` or an integer expression. Since the analysis in Section 5 cannot deal with generic integer expressions, we parse expressions in a careful way. In particular we split them into *size expressions* se , which are expressions in Presburger arithmetics (this is a decidable fragment of Peano arithmetics that only contains addition), and *non-size expressions* nse , which are the other type of expressions. The syntax of size and non-size expressions is the following:

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

In the paper, we assume that sequences of declarations $\overline{T} \overline{x}$ and method declarations \overline{M} do not contain duplicate names. We also assume that `return` statements have no continuation.

Semantics. The semantics of `tml` is defined by a transition system whose states are *configurations* cn that are defined by the following syntax.

$$\begin{aligned}
cn &::= \varepsilon \mid fut(f, val) \mid ob(o, c, p, q) \mid invoc(o, f, m, \vec{v}) & act &::= o \mid \varepsilon \\
&\quad \mid cog(c, act, k) \mid cn \quad cn & val &::= v \mid \perp \\
p &::= \{l \mid s\} \mid \text{idle} & l &::= [\dots, x \mapsto v, \dots] \\
q &::= \emptyset \mid \{l \mid s\} \mid q \quad q & v &::= o \mid f \mid k
\end{aligned}$$

A *configuration* cn is a set of concurrent object groups (cogs), objects, invocation messages and futures, and the empty configuration is written as ε . The associative and commutative union operator on configurations is denoted by whitespace. A *cog* is given as a term $cog(c, act, k)$ where c and k are respectively the identifier and the capacity of the cog, and act specifies the currently active object in the cog. An object is written as $ob(o, c, p, q)$, where o is the identifier of the object, c the identifier of the cog the object belongs to, p an *active process*, and q a pool of *suspended processes*. A *process* is written as $\{l \mid s\}$, where l denotes local variable bindings and s a list of statements. An *invocation message* is a term $invoc(o, f, m, \vec{v})$ consisting of the callee o , the future f to which the result of the call is returned, the method name m , and the set of actual parameter values for the call. A *future* $fut(f, val)$ contains an identifier f and a reply value val , where \perp indicates the reply value of the future has not been received.

$$\begin{array}{c}
\text{(ASSIGN-LOCAL)} \\
\frac{x \in \text{dom}(l) \quad v = \llbracket e \rrbracket_l}{ob(o, c, \{l \mid x = e\}; s), q} \\
\rightarrow ob(o, c, \{l \mid x \mapsto v\}; s), q
\end{array}
\quad
\begin{array}{c}
\text{(COND-TRUE)} \\
\frac{\text{true} = \llbracket e \rrbracket_l}{ob(o, c, \{l \mid \text{if } e \{s_1\} \text{ else } \{s_2\}; s\}, q)} \\
\rightarrow ob(o, c, \{l \mid s_1\}; s), q
\end{array}
\quad
\begin{array}{c}
\text{(COND-FALSE)} \\
\frac{\text{false} = \llbracket e \rrbracket_l}{ob(o, c, \{l \mid \text{if } e \{s_1\} \text{ else } \{s_2\}; s\}, q)} \\
\rightarrow ob(o, c, \{l \mid s_2\}; s), q
\end{array}$$

$$\begin{array}{c}
\text{(NEW)} \\
\frac{c' = \text{fresh}() \quad o' = \text{fresh}() \quad k = \llbracket e \rrbracket_l}{ob(o, c, \{l \mid x = \text{new Class with } e\}; s), q)} \\
\rightarrow ob(o, c, \{l \mid x = o'\}; s), q \quad ob(o', c', \text{idle}, \emptyset) \quad cog(c', o', k)
\end{array}
\quad
\begin{array}{c}
\text{(NEW-LOCAL)} \\
\frac{o' = \text{fresh}()}{ob(o, c, \{l \mid x = \text{new local Class}\}; s), q)} \\
\rightarrow ob(o, c, \{l \mid x = o'\}; s), q \quad ob(o', c, \text{idle}, \emptyset)
\end{array}$$

$$\begin{array}{c}
\text{(GET-TRUE)} \\
\frac{f = \llbracket e \rrbracket_l \quad v \neq \perp}{ob(o, c, \{l \mid x = e.\text{get}\}; s), q} \quad fut(f, v) \\
\rightarrow ob(o, c, \{l \mid x = v\}; s), q \quad fut(f, v)
\end{array}
\quad
\begin{array}{c}
\text{(GET-FALSE)} \\
\frac{f = \llbracket e \rrbracket_l}{ob(o, c, \{l \mid x = e.\text{get}\}; s), q} \quad fut(f, \perp) \\
\rightarrow ob(o, c, \text{idle}, q \cup \{l \mid x = e.\text{get}\}; s) \quad fut(f, \perp)
\end{array}$$

$$\begin{array}{c}
\text{(SELF-SYNC-CALL)} \\
\frac{o = \llbracket e \rrbracket_l \quad \bar{v} = \llbracket \bar{e} \rrbracket_l \quad f' = l(\text{destiny}) \quad f = \text{fresh}() \quad \{l' \mid s'\} = \text{bind}(o, f, m, \bar{v})}{ob(o, c, \{l \mid x = e.m(\bar{e})\}; s), q)} \\
\rightarrow ob(o, c, \{l' \mid s'; \text{cont}(f')\}, q \cup \{l \mid x = f.\text{get}\}; s) \quad fut(f, \perp)
\end{array}
\quad
\begin{array}{c}
\text{(SELF-SYNC-RETURN-SCHED)} \\
\frac{f = l'(\text{destiny})}{ob(o, c, \{l \mid \text{cont}(f)\}, q \cup \{l' \mid s'\})} \\
\rightarrow ob(o, c, \{l' \mid s'\}, q)
\end{array}$$

$$\begin{array}{c}
\text{(COG-SYNC-CALL)} \\
\frac{o' = \llbracket e \rrbracket_l \quad \bar{v} = \llbracket \bar{e} \rrbracket_l \quad f' = l(\text{destiny}) \quad f = \text{fresh}() \quad \{l' \mid s'\} = \text{bind}(o', f, m, \bar{v})}{ob(o, c, \{l \mid x = e.m(\bar{e})\}; s), q)} \\
\rightarrow ob(o, c, \text{idle}, q \cup \{l \mid x = f.\text{get}\}; s) \quad fut(f, \perp) \\
ob(o', c, \{l' \mid s'; \text{cont}(f')\}, q') \quad cog(c, o', k)
\end{array}
\quad
\begin{array}{c}
\text{(COG-SYNC-RETURN-SCHED)} \\
\frac{f = l'(\text{destiny})}{ob(o, c, \{l \mid \text{cont}(f)\}, q) \quad cog(c, o, k)} \\
\rightarrow ob(o, c, \text{idle}, q' \cup \{l' \mid s'\}) \\
\rightarrow ob(o, c, \text{idle}, q) \quad cog(c, o', k) \\
ob(o', c, \{l' \mid s'\}, q')
\end{array}$$

$$\begin{array}{c}
\text{(ASYNC-CALL)} \\
\frac{o' = \llbracket e \rrbracket_l \quad \bar{v} = \llbracket \bar{e} \rrbracket_l \quad f = \text{fresh}()}{ob(o, c, \{l \mid x = e!m(\bar{e})\}; s), q)} \\
\rightarrow ob(o, c, \{l \mid x = f\}; s), q \quad \text{invoc}(o', f, m, \bar{v}) \quad fut(f, \perp)
\end{array}
\quad
\begin{array}{c}
\text{(BIND-MTD)} \\
\frac{\{l \mid s\} = \text{bind}(o, f, m, \bar{v})}{ob(o, c, p, q) \quad \text{invoc}(o, f, m, \bar{v})} \\
\rightarrow ob(o, c, p, q \cup \{l \mid s\})
\end{array}$$

$$\begin{array}{c}
\text{(RELEASE-COG)} \\
ob(o, c, \text{idle}, q) \quad cog(c, o, k) \\
\rightarrow ob(o, c, \text{idle}, q) \quad cog(c, \varepsilon, k)
\end{array}
\quad
\begin{array}{c}
\text{(ACTIVATE)} \\
ob(o, c, \text{idle}, q \cup \{l \mid s\}) \quad cog(c, \varepsilon, k) \\
\rightarrow ob(o, c, \{l \mid s\}, q) \quad cog(c, o, k)
\end{array}$$

$$\begin{array}{c}
\text{(RETURN)} \\
\frac{v = \llbracket e \rrbracket_l \quad f = l(\text{destiny})}{ob(o, c, \{l \mid \text{return } e\}, q) \quad fut(f, \perp)} \\
\rightarrow ob(o, c, \text{idle}, q) \quad fut(f, v)
\end{array}
\quad
\begin{array}{c}
\text{(JOB-0)} \\
\frac{\llbracket e \rrbracket_l = 0}{ob(o, c, \{l \mid \text{job}(e)\}; s), q)} \\
\rightarrow ob(o, c, \{l \mid s\}, q)
\end{array}
\quad
\begin{array}{c}
\text{(CONTEXT)} \\
\frac{cn \rightarrow cn'}{cn \quad cn'' \rightarrow cn' \quad cn''}
\end{array}$$

Fig. 1. The transition relation of tml – part 1.

The following auxiliary function is used in the semantic rules for invocations. Let $T' \ m(\overline{T \ x})\{ \overline{F \ x'}; s \}$ be a method declaration. Then

$$\text{bind}(o, f, m, \bar{v}) = \{ [\text{destiny} \mapsto f, \bar{x} \mapsto \bar{v}, \bar{x}' \mapsto \perp] \mid s\{^o/\text{this}\} \}$$

The *transition rules* of tml are given in Figures 1 and 2. We discuss the most relevant ones: object creation, method invocation, and the **job**(*e*) operator. The creation of objects is handled by rules **NEW** and **NEW-LOCAL**: the former creates a new object inside a new cog with a given capacity *e*, the latter creates an object in the local cog. Method invocations can be either synchronous or asynchronous. Rules **SELF-SYNC-CALL** and **COG-SYNC-CALL** specify synchronous invocations on objects belonging to the same cog of the caller. Asynchronous invocations can be performed on every object.

$$\begin{array}{c}
\text{(Tick)} \\
\frac{\text{strongstable}_t(cn)}{cn \rightarrow \Phi(cn, t)} \\
\text{where} \\
\Phi(cn, t) = \begin{cases}
ob(o, c, \{l \mid \mathbf{job}(k') ; s\}, q) \Phi(cn', t) & \text{if } cn = ob(o, c, \{l \mid \mathbf{job}(e) ; s\}, q) cn' \\
& \text{and } cog(c, o, k) \in cn' \\
& \text{and } k' = \llbracket e \rrbracket_l - k * t \\
ob(o, c, \mathbf{idle}, q) \Phi(cn', t) & \text{if } cn = ob(o, c, \mathbf{idle}, q) cn' \\
cn & \text{otherwise.}
\end{cases}
\end{array}$$

Fig. 2. The transition relation of tml – part 2: the strongly stable case

In our model, the unique operation that consumes time is $\mathbf{job}(e)$. We notice that the reduction rules of Figure 1 are not defined for the $\mathbf{job}(e)$ statement, except the trivial case when the value of e is 0. This means that time does not advance while non-job statements are evaluated. When the configuration cn reaches a *stable* state, *i.e.*, no other transition is possible apart from those evaluating the $\mathbf{job}(e)$ statements, then the time is advanced by the minimum value that is necessary to let at least *one* process start. In order to formalize this semantics, we define the notion of stability and the *update operation* of a configuration cn (with respect to a time value t). Let $\llbracket e \rrbracket_l$ return the value of e when variables are bound to values stored in l .

Definition 1. Let $t > 0$. A configuration cn is t -stable, written $\text{stable}_t(cn)$, if any object in cn is in one of the following forms:

1. $ob(o, c, \{l \mid \mathbf{job}(e); s\}, q)$ with $cog(c, o, k) \in cn$ and $\llbracket e \rrbracket_l / k \geq t$,
2. $ob(o, c, \mathbf{idle}, q)$ and
 - i. either $q = \emptyset$,
 - ii. or, for every $p \in q$, $p = \{l \mid x = e.\mathbf{get}; s\}$ with $\llbracket e \rrbracket_l = f$ and $\text{fut}(f, \perp)$,
 - iii. or, $cog(c, o', k) \in cn$ where $o \neq o'$, and o' satisfies Definition 1.1.

A configuration cn is *strongly* t -stable, written $\text{strongstable}_t(cn)$, if it is t -stable and there is an object $ob(o, c, \{l \mid \mathbf{job}(e); s\}, q)$ with $cog(c, o, k) \in cn$ and $\llbracket e \rrbracket_l / k = t$.

Notice that t -stable (and, consequently, strongly t -stable) configurations cannot progress anymore because every object is stuck either on a job or on unresolved get statements. The update of cn with respect to a time value t , noted $\Phi(cn, t)$ is defined in Figure 2. Given these two notions, rule Tick defines the time progress.

The initial configuration of a program with main method $\{\overline{F} x; s\}$ **with** k is

$$\begin{array}{c}
ob(\text{start}, \text{start}, \{[\text{destiny} \mapsto f_{\text{start}}, \bar{x} \mapsto \perp] \mid s\}, \emptyset) \\
cog(\text{start}, \text{start}, k)
\end{array}$$

where start and start are special cog and object names, respectively, and f_{start} is a fresh future name. As usual, \rightarrow^* is the reflexive and transitive closure of \rightarrow .

Examples. To begin with, we discuss the Fibonacci method. It is well known that the computational cost of its sequential recursive implementation is exponential. However, this is not the case for the parallel implementation. Consider

```

Int fib(Int n) {
    if (n<=1) { return 1; }
    else { Fut<Int> f; Class z; Int m1; Int m2;
          job(1);
          z = new Class with this.capacity ;
          f = this!fib(n-1); g = z!fib(n-2);
          m1 = f.get; m2 = g.get;
          return m1 + m2;
        }
    }

```

Here, the recursive invocation `fib(n-1)` is performed on the `this` object while the invocation `fib(n-2)` is performed on a new cog with the same capacity (i.e., the object referenced by `z` is created in a new cog set up with `this.capacity`), which means that it can be performed in parallel with the former one. It turns out that the cost of the following invocation is n .

```

Class z; Int m; Int x;
x = 1;
z = new Class with x;
m = z.fib(n);

```

Observe that, by changing the line `x = 1;` into `x = 2;` we obtain a cost of $n/2$.

Our semantics does not exclude paradoxical behaviours of programs that perform infinite actions without consuming time (preventing rule Tick to apply), such as this one

```

Int foo() { Int m; m = this.foo(); return m; }

```

This kind of behaviours are well-known in the literature, (*cf.* Zeno behaviours) and they may be easily excluded from our analysis by constraining recursive invocations to be prefixed by a **job**(e)-statement, with a positive e . It is worth to observe that this condition is not sufficient to eliminate paradoxical behaviours. For instance the method below does not terminate and, when invoked with `this.fake(2)`, where `this` is in a cog of capacity 2, has cost 1.

```

Int fake(Int n) {
    Int m; Class x;
    x = new Class with 2*n; job(1); m = x.fake(2*n); return m;
}

```

Imagine a parallel invocation of the following method

```

Int one() { job(1); }

```

on an object residing in a cog of capacity 1. At each stability point the **job**(1) of the latter method will compete with the **job**(1) of the former one, which will win every time, since having a greater (and growing) capacity it will require always less time. So at

the first stability point we get $\mathbf{job}(1 - 1/2)$ (for the method one), then $\mathbf{job}(1 - 1/2 - 1/4)$ and so on, thus this sum will never reach 0.

In the examples above, the statement $\mathbf{job}(e)$ is a cost annotation that specifies how many processing cycles are needed by the subsequent statement in the code. We notice that this operation can also be used to program a timer which suspends the current execution for e units of time. For instance, let

```
Int wait(Int n) { job(n); return 0; }
```

Then, invoking `wait` on an object with capacity 1

```
Class timer; Fut<Class> f; Class x;
timer = new Class with 1;
f = timer!wait(5); x = f.get;
```

one gets the suspension of the current thread for 5 units of time.

3 Issues in computing the cost of tml programs

The computation time analysis of tml programs is demanding. To highlight the difficulties, we discuss a number of methods.

```
Int wrapper(Class x) {
    Fut<Int> f; Int z;
    job(1) ; f = x!server(); z = f.get;
    return z;
}
```

Method `wrapper` performs an invocation on its argument x . In order to determine the cost of `wrapper`, we notice that, if x is in the same cog of the carrier, then its cost is (assume that the capacity of the carrier is 1): $1 + \mathit{cost}(\mathit{server})$ because the two invocations are sequentialized. However, if the cogs of x and of the carrier are different, then we are not able to compute the cost because we have no clue about the state of the cog of x . Next consider the following definition of `wrapper`

```
Int wrapper_with_log(Class x) {
    Fut<Int> f; Fut<Int> g; Int z;
    job(1) ; f = x!server(); g = x!print_log(); z = f.get;
    return z;
}
```

In this case the wrapper also asks the server to print its log and this invocation is not synchronized. We notice that the cost of `wrapper_with_log` is not anymore $1 + \mathit{cost}(\mathit{server})$ (assuming that x is in the same cog of the carrier) because `print_log` might be executed *before* `server`. Therefore the cost of `wrapper_with_log` is $1 + \mathit{cost}(\mathit{server}) + \mathit{cost}(\mathit{print_log})$.

Finally, consider the following wrapper that also logs the information received from the server on a new cog without synchronising with it:


```

Int wrapper_with_external_log(Class x) {
  Fut<Int> f; Fut<Int> g; Int z; Class y;
  job(1) ; f = x!server(); g = x!print_log(); z = f.get;
  y = new Class with 1;
  f = y!external_log(z) ;
  return z;
}

```

What is the cost of `wrapper_with_external_log`? Well, the answer here is debatable: one might discard the cost of `y!external_log(z)` because it is useless for the value returned by `wrapper_with_external_log`, or one might count it because one wants to count every computation that has been triggered by a method in its cost. In this paper we adhere to the second alternative; however, we think that a better solution should be to return different cost for a method: a *strict cost*, which spots the cost that is necessary for computing the returned value, and an *overall cost*, which is the one computed in this paper.

Anyway, by the foregoing discussion, as an initial step towards the time analysis of `tml` programs, we simplify our analysis by imposing the following constraint:

- *it is possible to invoke methods on objects either in the same cog of the caller or on newly created cogs.*

The above constraint means that, if the callee of an invocation is one of the arguments of a method then it must be in the same cog of the caller. It also means that, if an invocation is performed on a returned object then this object must be in the same cog of the carrier. We will enforce these constraints in the typing system of the following section – see rule T-INVOKE.

4 A behavioural type system for `tml`

In order to analyse the computation time of `tml` programs we use abstract descriptions, called *behavioural types*, which are intermediate codes highlighting the features of `tml` programs that are relevant for the analysis in Section 5. These abstract descriptions support compositional reasoning and are associated to programs by means of a type system. The syntax of behavioural types is defined as follows:

$t ::=$	$- \mid se \mid c[se]$	basic value
$x ::=$	$f \mid t$	extended value
$a ::=$	$e \mid vc[se] \mid m(\bar{t}) \rightarrow t \mid \nu f: m(\bar{t}) \rightarrow t \mid f'$	atom
$b ::=$	$a \triangleright I \mid a \S b \mid (se)\{b\} \mid b + b$	behavioural type

where c, c', \dots range over cog names and f, f', \dots range over future names. Basic values t are either generic (non-size) expressions $-$ or size expressions se or the type $c[se]$ of an object of cog c with capacity se . The extended values add future names to basic values.

Atoms a define creation of cogs ($vc[se]$), synchronous and asynchronous method invocations ($m(\bar{t}) \rightarrow t$ and $\nu f: m(\bar{t}) \rightarrow t$, respectively), and synchronizations on asynchronous invocations (f'). We observe that cog creations always carry a capacity, which has to be a size expression because our analysis in the next section cannot deal

with generic expressions. Behavioural types b are sequences of atoms $a \sharp b'$ or conditionals, typically $(se)\{b\} + (\neg se)\{b'\}$ or $b + b'$, according to whether the boolean guard is a size expression that depends on the arguments of a method or not. In order to type sequential composition in a precise way (see rule T-SEQ), the leaves of behavioural types are labelled with *environments*, ranged over by Γ, Γ', \dots . Environments are maps from method names m to terms $(\bar{t}) \rightarrow t$, from variables to extended values x , and from future names to values that are either t or t^\checkmark .

The abstract behaviour of methods is defined by *method behavioural types* of the form: $m(t_r, \bar{t})\{b\} : t_r$, where t_r is the type value of the receiver of the method, \bar{t} are the type value of the arguments, b is the abstract behaviour of the body, and t_r is the type value of the returned object. The subterm t_r, \bar{t} of the method contract is called *header*; t_r is called *returned type value*. We assume that names in the header occur linearly. Names in the header *bind* the names in b and in t_r . The header and the returned type value, written $(t_r, \bar{t}) \rightarrow t_r$, are called *behavioural type signature*. Names occurring in b or t_r may be *not bound* by header. These *free names* correspond to new cog creations and will be replaced by fresh cog names during the analysis. We use C to range over method behavioural types.

The type system uses judgments of the following form:

- $\Gamma \vdash e : x$ for pure expressions e , $\Gamma \vdash f : t$ or $\Gamma \vdash f : t^\checkmark$ for future names f , and $\Gamma \vdash m(\bar{t}) : t$ for methods.
- $\Gamma \vdash z : x$, $[a \triangleright \Gamma']$ for expressions with side effects z , where x is the value, $a \triangleright \Gamma'$ is the corresponding behavioural type, where Γ' is the environment Γ with possible updates of variables and future names.
- $\Gamma \vdash s : b$, in this case the updated environments Γ' are inside the behavioural type, in correspondence of every branch of its.

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

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

The *multi-hole contexts* $C[\]$ are defined by the following syntax:

$$C[\] ::= [\] \mid a \sharp C[\] \mid C[\] + C[\] \mid (se)\{C[\]\}$$

and, whenever $b = C[a_1 \triangleright \Gamma_1] \cdots [a_n \triangleright \Gamma_n]$, then $b[x \mapsto x]$ is defined as $C[a_1 \triangleright \Gamma_1[x \mapsto x]] \cdots [a_n \triangleright \Gamma_n[x \mapsto x]]$.

The typing rules for expressions are defined in Figure 3. These rules are not standard because (size) expressions containing method's arguments are typed with the expressions themselves. This is crucial to the cost analysis in Section 5. In particular, *cog creation* is typed by rule T-NEW, with value $c[se]$, where c is the fresh name associated with the new cog and se is the value associated with the declared capacity. The behavioural type for the cog creation is $vc[se] \triangleright \Gamma[c \mapsto se]$, where the newly created cog is added to Γ . In this way, it is possible to verify whether the receiver of a method invocation is within a locally created cog or not by testing whether the receiver belongs to $\text{dom}(\Gamma)$ or not, respectively (cf. rule T-VOKE). *Object creation* (cf. rule T-NEW-LOCAL) is typed as the cog creation, with the exception that the cog name and the capacity value are taken

$$\begin{array}{c}
\text{(T-VAR)} \\
\frac{x \in \text{dom}(\Gamma)}{\Gamma \vdash x : \Gamma(x)} \\
\\
\text{(T-SE)} \quad \Gamma \vdash se : se \\
\text{(T-NSE)} \quad \Gamma \vdash nse : - \\
\\
\text{(T-METHOD)} \\
\frac{\Gamma(\mathbf{m}) = (\bar{t}) \rightarrow t' \quad fv(t') \setminus fv(\bar{t}) \neq \emptyset \quad \text{implies} \quad \sigma(t') \text{ fresh}}{\Gamma \vdash \mathbf{m}(\sigma(\bar{t})) : \sigma(t')} \\
\\
\text{(T-NEW)} \\
\frac{\Gamma \vdash e : se \quad c \text{ fresh}}{\Gamma \vdash \text{new Class } \mathbf{with} \ e : c[se], [\nu c[se] \triangleright \Gamma[c \mapsto se]]} \\
\text{(T-NEW-LOCAL)} \\
\frac{\Gamma \vdash \mathbf{this} : c[se]}{\Gamma \vdash \text{new local Class} : c[se], [\emptyset \triangleright \Gamma]} \\
\\
\text{(T-INVOKE-SYNC)} \\
\frac{\Gamma \vdash e : c[se] \quad \Gamma(\mathbf{this}) = c[se] \quad \Gamma \vdash \bar{e} : \bar{t} \quad \Gamma \vdash \mathbf{m}(c[se], \bar{t}) : t'}{\Gamma \vdash e.\mathbf{m}(\bar{e}) : t', [\mathbf{m}(c[se], \bar{t}) \rightarrow t' \triangleright \Gamma]} \\
\text{(T-INVOKE)} \\
\frac{\Gamma \vdash e : c[se] \quad (c \in \text{dom}(\Gamma) \text{ or } \Gamma(\mathbf{this}) = c[se]) \quad \Gamma \vdash \bar{e} : \bar{t} \quad \Gamma \vdash \mathbf{m}(c[se], \bar{t}) : t' \quad f \text{ fresh}}{\Gamma \vdash e!\mathbf{m}(\bar{e}) : f, [\nu f : \mathbf{m}(c[se], \bar{t}) \rightarrow t' \triangleright \Gamma[f \mapsto t']]} \\
\\
\text{(T-GET)} \\
\frac{\Gamma \vdash e : f \quad \Gamma(f) = t}{\Gamma \vdash e.\mathbf{get} : t, [f' \triangleright \Gamma[f \mapsto t']]} \\
\text{(T-GET-TOP)} \\
\frac{\Gamma \vdash e : f \quad \Gamma(f) = t'}{\Gamma \vdash e.\mathbf{get} : t, [\emptyset \triangleright \Gamma]}
\end{array}$$

Fig. 3. Typing rules for expressions

from the local cog and the behavioural type is empty. Rule T-INVOKE types *method invocations* $e!\mathbf{m}(\bar{e})$ by using a fresh future name f that is associated to the method name, the cog name of the callee and the arguments. In the updated environment, f is associated with the returned value. Next we discuss the constraints in the premise of the rule. As we discussed in Section 2, asynchronous invocations are allowed on callees located in the current cog, $\Gamma(\mathbf{this}) = c[se]$, or on a newly created object which resides in a fresh cog, $c \in \text{dom}(\Gamma)$. Rule T-GET defines the *synchronization* with a method invocation that corresponds to a future f . The expression is typed with the value t of f in the environment and behavioural type f' . Γ is then updated for recording that the synchronization has been already performed, thus any subsequent synchronization on the same value would not imply any waiting time (see that in rule T-GET-TOP the behavioural type is \emptyset). The *synchronous method invocation* in rule T-INVOKE-SYNC is directly typed with the return value t' of the method and with the corresponding behavioural type. The rule enforces that the cog of the callee coincides with the local one.

The typing rules for statements are presented in Figure 4. The behavioural type in rule T-JOB expresses the time consumption for an object with capacity se' to perform se processing cycles: this time is given by se/se' , which we observe is in general a rational number. We will return to this point in Section 5.

The typing rules for method and class declarations are shown in Figure 5.

Examples The behavioural type of the `fib` method discussed in Section 2 is

$$\begin{array}{l}
\text{fib}(c[x], n) \{ \\
\quad (n \leq 1) \{ \emptyset \triangleright \emptyset \} \\
\quad + \\
\}
\end{array}$$

$$\begin{array}{c}
\text{(T-ASSIGN)} \\
\frac{\Gamma \vdash rhs : x, [a \triangleright \Gamma']}{\Gamma \vdash x = rhs : a \triangleright \Gamma'[x \mapsto x]} \\
\\
\text{(T-RETURN)} \\
\frac{\Gamma \vdash e : t \quad \Gamma \vdash \mathbf{destiny} : t}{\Gamma \vdash \mathbf{return} e : \mathbf{0} \triangleright \Gamma} \\
\\
\text{(T-IF-NSE)} \\
\frac{\Gamma \vdash e : - \quad \Gamma \vdash s : b \quad \Gamma \vdash s' : b'}{\Gamma \vdash \mathbf{if} e \{ s \} \mathbf{else} \{ s' \} : b + b'} \\
\\
\text{(T-JOB)} \\
\frac{\Gamma \vdash e : se \quad \Gamma \vdash \mathbf{this} : c[se']}{\Gamma \vdash \mathbf{job}(e) : se/se' \triangleright \Gamma} \\
\\
\text{(T-SEQ)} \\
\frac{\Gamma \vdash s : C[a_1 \triangleright \Gamma_1] \cdots [a_n \triangleright \Gamma_n] \quad \Gamma_i \vdash s' : b'_i}{\Gamma \vdash s ; s' : C[a_1 \mathbin{\&}; b'_1] \cdots [a_n \mathbin{\&}; b'_n]} \\
\\
\text{(T-IF-SE)} \\
\frac{\Gamma \vdash e : se \quad \Gamma \vdash s : b \quad \Gamma \vdash s' : b'}{\Gamma \vdash \mathbf{if} e \{ s \} \mathbf{else} \{ s' \} : (se)\{b\} + (\neg se)\{b'\}}
\end{array}$$

Fig. 4. Typing rules for statements

$$\begin{array}{c}
\text{(T-METHOD)} \\
\frac{\Gamma(m) = (t_i, \bar{t}) \rightarrow t_r \quad \Gamma[\mathbf{this} \mapsto t_i][\mathbf{destiny} \mapsto t_r][\bar{x} \mapsto \bar{t}] \vdash s : C[a_1 \triangleright \Gamma_1] \cdots [a_n \triangleright \Gamma_n]}{\Gamma \vdash T m(\bar{T} \bar{x}) \{ s \} : m(t_i, \bar{t})\{C[a_1 \triangleright \emptyset] \cdots [a_n \triangleright \emptyset]\} : t_r} \\
\\
\text{(T-CLASS)} \\
\frac{\Gamma \vdash \bar{M} : \bar{C} \quad \Gamma[\mathbf{this} \mapsto \mathbf{start}[k]][\bar{x} \mapsto \bar{t}] \vdash s : C[a_1 \triangleright \Gamma_1] \cdots [a_n \triangleright \Gamma_n]}{\Gamma \vdash \bar{M} \{ \bar{T} \bar{x} ; s \} \mathbf{with} k : \bar{C}, C[a_1 \triangleright \emptyset] \cdots [a_n \triangleright \emptyset]}
\end{array}$$

Fig. 5. Typing rules for declarations

$$\begin{array}{l}
(n \geq 2) \{ \\
\quad 1/x \mathbin{\&} d[x] \mathbin{\&} vf : \mathbf{fib}(c[x], n-1) \rightarrow - \mathbin{\&} vg : \mathbf{fib}(d[x], n-2) \rightarrow - \mathbin{\&} \\
\quad \quad f' \mathbin{\&} g' \mathbin{\&} \mathbf{0} \triangleright \emptyset \\
\} \\
\} : -
\end{array}$$

5 The time analysis

The behavioural types returned by the system defined in Section 4 are used to compute upper bounds of time complexity of a `tml` program. This computation is performed by an off-the-shelf solver – the `CoFloCo` solver [4] – and, in this section, we discuss the translation of a behavioural type program into a set of *cost equations* that are fed to the solver. These cost equations are terms

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

where m is a (cost) function symbol, exp is an expression that may contain (cost) function symbol applications (we do not define the syntax of exp , which may be derived

by the following equations; the reader may refer to [4]), and se is a size expression whose variables are contained in \bar{x} . Basically, our translation maps method types into cost equations, where (i) method invocations are translated into function applications, and (ii) cost expressions se occurring in the types are left unmodified. The difficulties of the translation is that the cost equations must account for the parallelism of processes in different cogs and for sequentiality of processes in the same cog. For example, in the following code:

```
x = new Class with c; y = new Class with d;
f = x!m(); g = y!n(); u = g.get ; u = f.get ;
```

the invocations of m and n will run in parallel, therefore their cost will be $\max(t, t')$, where t is the time of executing m on x and t' is the time executing n on y . On the contrary, in the code

```
x = new local Class; y = new local Class;
f = x!m(); g = y!n(); u = g.get; u = f.get;
```

the two invocations are queued for being executed on the same cog. Therefore the time needed for executing them will be $t+t'$, where t is time needed for executing m on x , and t' is the time needed for executing n on y . To abstract away the execution order of the invocations, the execution time of *all unsynchronized* methods from the same cog are taken into account when one of these methods is synchronized with a `get`-statement. To avoid calculating the execution time of the rest of the unsynchronized methods in the same cog more than necessary, their estimated cost are ignored when they are later synchronized.

In this example, when the method invocation `y!n()` is synchronized with `g.get`, the estimated time taken is $t + t'$, which is the sum of the execution time of the two unsynchronized invocations, including the time taken for executing m on x because both x and y are residing in the same cog. Later when synchronizing the method invocation `x!m()`, the cost is considered to be *zero* because this invocation has been taken into account earlier.

The translate function. The translation of behavioural types into cost equations is carried out by the function `translate`, defined below. This function parses atoms, behavioural types or declarations of methods and classes. We will use the following auxiliary function that removes cog names from (tuples of) t terms:

$$[-] = - \quad [e] = e \quad [c[e]] = e \quad [t_1, \dots, t_n] = [t_1], \dots, [t_n]$$

We will also use *translation environments*, ranged over by Ψ, Ψ', \dots , which map future names to pairs $(e, m(\bar{t}))$ that records the (over-approximation of the) time when the method has been invoked and the invocation.

In the case of atoms, `translate` takes four inputs: a *translation environment* Ψ , the cog name of the carrier, an over-approximated cost e of an execution branch, and the atom a . In this case, `translate` returns an updated translation environment and the

cost. It is defined as follows.

$$\text{translate}(\Psi, c, e, a) = \begin{cases} (\Psi, e + e') & \text{when } a = e' \\ (\Psi, e) & \text{when } a = \nu c[e'] \\ (\Psi, e + m(\overline{[t]})) & \text{when } a = m(\overline{t}) \rightarrow t' \\ (\Psi[f \mapsto (e, m(\overline{t}))], e) & \text{when } a = (\nu f: m(\overline{t}) \rightarrow t') \\ (\Psi \setminus F, e + e_1)) & \text{when } a = f' \text{ and } \Psi(f) = (e_f, m_f(c[e'], \overline{t_f})) \\ & \text{let } F = \{g \mid \Psi(g) = (e_g, m_g(c[e'], \overline{t_g}))\} \text{ then} \\ & \text{and } e_1 = \sum \{m_g(\overline{[t'_g]}) \mid (e_g, m_g(\overline{t'_g})) \in \Psi(F)\} \\ (\Psi \setminus F, \max(e, e_1 + e_2)) & \text{when } a = f' \text{ and } \Psi(f) = (e_f, m_f(c[e'], \overline{t_f})) \text{ and } c \neq c' \\ & \text{let } F = \{g \mid \Psi(g) = (e_g, m_g(c'[e'], \overline{t_g}))\} \text{ then} \\ & e_1 = \sum \{m_g(\overline{[t'_g]}) \mid (e_g, m_g(\overline{t'_g})) \in \Psi(F)\} \\ & \text{and } e_2 = \max\{e_g \mid (e_g, m_g(\overline{t'_g})) \in \Psi(F)\} \\ (\Psi, e) & \text{when } a = f' \text{ and } f \notin \text{dom}(\Psi) \end{cases}$$

The interesting case of `translate` is when the atom is f' . There are three cases:

1. The synchronization is with a method whose callee is an object of the same cog. In this case its cost must be *added*. However, it is not possible to know when the method will be actually scheduled. Therefore, we sum the costs of all the methods running on the same cog (worst case) – the set F in the formula – and we remove them from the translation environment.
2. The synchronization is with a method whose callee is an object on a different cog c' . In this case we use the cost that we stored in $\Psi(f)$. Let $\Psi(f) = (e_f, m_f(c[e'], \overline{t_f}))$, then e_f represents the time of the invocation. The cost of the invocation is therefore $e_f + m_f(e', \overline{[t_f]})$. Since the invocation is *in parallel* with the thread of the cog c , the overall cost will be $\max(e, e_f + m_f(e', \overline{[t_f]}))$. As in case 1, we consider the worst scheduler choice on c' . Therefore, instead of taking $e_f + m_f(e', \overline{[t_f]})$, we compute the cost of all the methods running on c' – the set F in the formula – and we remove them from the translation environment.
3. The future does not belong to Ψ . That is the cost of the invocation which has been already computed. In this case, the value e does not change.

In the case of behavioural types, `translate` takes as input a translation environment, the cog name of the carrier, an over-approximated cost of the current execution branch $(e_1)e_2$, where e_1 indicates the conditions corresponding to the branch, and the behavioural type a .

$$\text{translate}(\Psi, c, (e_1)e_2, b) = \begin{cases} \{(\Psi', (e_1)e'_2)\} & \text{when } b = a \triangleright \Gamma \text{ and } \text{translate}(\Psi, c, e_2, a) = (\Psi', e'_2) \\ C & \text{when } b = a \S b' \text{ and } \text{translate}(\Psi, c, e_2, a) = (\Psi', e'_2) \\ & \text{and } \text{translate}(\Psi', c, (e_1)e'_2, b') = C \\ C \cup C' & \text{when } b = b_1 + b_2 \text{ and } \text{translate}(\Psi, c, (e_1)e_2, b_1) = C \\ & \text{and } \text{translate}(\Psi, c, (e_1)e_2, b_2) = C' \\ C & \text{when } b = (e)\{b'\} \text{ and } \text{translate}(\Psi, c, (e_1 \wedge e)e_2, b') = C \end{cases}$$

The translation of the behavioural types of a method is given below. Let $\text{dom}(\Psi) = \{f_1, \dots, f_n\}$. Then we define $\Psi^\checkmark \stackrel{\text{def}}{=} f_1^\checkmark \wp \dots \wp f_n^\checkmark$.

$$\text{translate}(m(c[e], \bar{t})\{b\} : t) = \begin{cases} m(e, \bar{e}) = e'_1 + e''_1 & [e_1] \\ \vdots \\ m(e, \bar{e}) = e'_n + e''_n & [e_n] \end{cases}$$

where $\text{translate}(\emptyset, c, \mathbf{0}, b) = \{\Psi_i, (e_i)e'_i \mid 1 \leq i \leq n\}$, and $\bar{e} = [\bar{t}]$, and $e''_i = \text{translate}(\Psi_i, c, \mathbf{0}, \Psi_i^\checkmark \triangleright \emptyset)$.

In addition, $[e_i]$ are the conditions for branching the possible execution paths of method $m(e, \bar{e})$, and $e'_i + e''_i$ is the over-approximation of the cost for each path. In particular, e'_i corresponds to the cost of the synchronized operations in each path (e.g., **jobs** and **gets**), while e''_i corresponds to the cost of the asynchronous method invocations triggered by the method, but not synchronized within the method body.

Examples We show the translation of the behavioural type of fibonacci presented in Section 4. Let $b = (se)\{\mathbf{0} \triangleright \emptyset\} + (\neg se)\{b'\}$, where $se = (n \leq 1)$ and $b' = 1/e \wp vf: \text{fib}(c[e], n-1) \rightarrow - \wp vg: \text{fib}(c'[e], n-2) \rightarrow - \wp f^\checkmark \wp g^\checkmark \wp \mathbf{0} \triangleright \emptyset$. Let also $\Psi = \Psi_1 \cup \Psi_2$, where $\Psi_1 = [f \mapsto (1/e, \text{fib}(e, n-1))]$ and $\Psi_2 = [g \mapsto (1/e, \text{fib}(e, n-2))]$.

The following equations summarize the translation of the behavioural type of the fibonacci method.

$$\begin{aligned} \text{translate}(\emptyset, c, \mathbf{0}, b) &= \text{translate}(\emptyset, c, \mathbf{0}, (se)\{\mathbf{0} \triangleright \emptyset\}) \cup \text{translate}(\emptyset, c, \mathbf{0}, (\neg se)\{b'\}) \\ &= \text{translate}(\emptyset, c, (se)\mathbf{0}, \{\mathbf{0} \triangleright \emptyset\}) \cup \text{translate}(\emptyset, c, (\neg se)\mathbf{0}, \{1/e \wp \dots\}) \\ &= \{(se)\mathbf{0}\} \cup \text{translate}(\emptyset, c, (\neg se)(1/e), \{vf: \text{fib}(c[e], n-1) \rightarrow - \wp \dots\}) \\ &= \{(se)\mathbf{0}\} \cup \text{translate}(\Psi_1, c, (\neg se)(1/e), \{vg: \text{fib}(c'[e], n-2) \rightarrow - \wp \dots\}) \\ &= \{(se)\mathbf{0}\} \cup \text{translate}(\Psi, c, (\neg se)(1/e), \{f^\checkmark \wp g^\checkmark \wp \dots\}) \\ &= \{(se)\mathbf{0}\} \cup \text{translate}(\Psi_2, c, (\neg se)(1/e + \text{fib}(e, n-1)), \{g^\checkmark \wp \dots\}) \\ &= \{(se)\mathbf{0}\} \cup \text{translate}(\emptyset, c, (\neg se)(1/e + \max(\text{fib}(e, n-1), \text{fib}(e, n-2))), \{\mathbf{0} \triangleright \emptyset\}) \\ &= \{(se)\mathbf{0}\} \cup \{(\neg se)(1/e + \max(\text{fib}(e, n-1), \text{fib}(e, n-2)))\} \end{aligned}$$

$$\begin{aligned} \text{translate}(\emptyset, c, \mathbf{0}, \mathbf{0}) &= (\emptyset, \mathbf{0}) \\ \text{translate}(\emptyset, c, \mathbf{0}, 1/e) &= (\emptyset, 1/e) \\ \text{translate}(\emptyset, c, 1/e, vf: \text{fib}(c[e], n-1) \rightarrow -) &= (\Psi_1, 1/e) \\ \text{translate}(\Psi_1, c, 1/e, vg: \text{fib}(c'[e], n-2) \rightarrow -) &= (\Psi, 1/e) \\ \text{translate}(\Psi, c, 1/e, f^\checkmark) &= (\Psi_2, 1/e + \text{fib}(e, n-1)) \\ \text{translate}(\Psi_2, c, 1/e + \text{fib}(e, n-1), g^\checkmark) &= (\emptyset, 1/e + \max(\text{fib}(e, n-1), \text{fib}(e, n-2))) \end{aligned}$$

$$\text{translate}(\text{fib}(c[e], n)\{b\} : -) = \begin{cases} \text{fib}(e, n) = 0 & [n \leq 1] \\ \text{fib}(e, n) = 1/e + \max(\text{fib}(e, n-1), \text{fib}(e, n-2)) & [n \geq 2] \end{cases}$$

Remark 1. Rational numbers are produced by the rule T-Job of our type system. In particular behavioural types may manifest terms se/se' where se gives the processing cycles defined by the job operation and se' specifies the number of processing cycles per unit of time the corresponding cog is able to handle. Unfortunately, our backend solver – CoFloCo – cannot handle rationals se/se' where se' is a variable. This is the case, for instance, of our fibonacci example, where the cost of each iteration is $1/x$, where x is a parameter. In order to analyse this example, we need to determine *a priori* the capacity to be a constant – say 2 –, obtaining the following input for the solver:

```
eq(f(E,N),0,[],[-N>=1,2*E=1]).
eq(f(E,N),nat(E),[f(E,N-1)],[N>=2,2*E=1]).
eq(f(E,N),nat(E),[f(E,N-2)],[N>=2,2*E=1]).
```

Then the solver gives the following upper bound:

```
nat(N-1)* (1/2).
```

It is worth to notice that fixing the fibonacci method is easy because the capacity does not change during the evaluation of the method. This is not always the case, as in the following alternative definition of fibonacci:

```
Int fib_alt(Int n) {
  if (n<=1) { return 1; }
  else { Fut<Int> f; Class z; Int m1; Int m2;
        job(1);
        z = new Class with (this.capacity*2) ;
        f = this!fib_alt(n-1); g = z!fib_alt(n-2);
        m1 = f.get; m2 = g.get;
        return m1+m2; }
}
```

In this case, the recursive invocation $z!fib_alt(n-2)$ is performed on a cog with twice the capacity of the current one and CoFloCo is not able to handle it. It is worth to observe that this is a problem of the solver, which is otherwise very powerful for most of the examples. Our behavioural types carry enough information for dealing with more complex examples, so we will consider alternative solvers or combination of them for dealing with examples like `fib_alt`.

6 Properties

In order to prove the correctness of our system, we need to show that (i) the behavioural type system is correct, and (ii) the computation time returned by the solver is an upper bound of the actual cost of the computation.

The correctness of the type system in Section 4 is demonstrated by means of a subject reduction theorem expressing that if a runtime configuration cn is well typed and $cn \rightarrow cn'$ then cn' is well-typed as well, and the computation time of cn is larger or equal to that of cn' . In order to formalize this theorem we extend the typing to configurations and we also use extended behavioural types k with the following syntax

$$k ::= b \mid [b]_f^c \mid k \parallel k \quad \text{runtime behavioural type}$$

The type $[b]_f^c$ expresses the behaviour of an asynchronous method bound to the future f and running in the cog c ; the type $k \parallel k'$ expresses the parallel execution of methods in k and in k' .

We then define a relation \succeq_t between runtime behavioural types that relates types. The definition is algebraic, and $k \succeq_t k'$ is intended to mean that the computational time of k is at least that of $k' + t$ (or conversely the computational time of k' is at most that of $k - t$). This is actually the purpose of our theorems.

Theorem 1 (Subject Reduction). *Let cn be a configuration of a `tml` program and let k be its behavioural type. If cn is not strongly t -stable and $cn \rightarrow cn'$ then there exists k' typing cn' such that $k \succeq_0 k'$. If cn is strongly t -stable and $cn \rightarrow cn'$ then there exists k' typing cn' such that $k \succeq_t k'$.*

The proof of is a standard case analysis on the last reduction rule applied.

The second part of the proof requires an extension of the `translate` function to runtime behavioural types. We therefore define a cost of the equations \mathcal{E}_k returned by `translate(k)` – noted $\text{cost}(\mathcal{E}_k)$ – by unfolding the equational definitions.

Theorem 2 (Correctness). *If $k \succeq_t k'$, then $\text{cost}(\mathcal{E}_k) \geq \text{cost}(\mathcal{E}_{k'}) + t$.*

As a byproduct of Theorems 1 and 2, we obtain the correctness of our technique, modulo the correctness of the solver.

7 Related work

In contrast to the static time analysis for sequential executions proposed in [7], the paper proposes an approach to analyse time complexity for concurrent programs. Instead of using a Hoare-style proof system to reason about end-user deadlines, we estimate the execution time of a concurrent program by deriving the time-consuming behaviour with a type-and-effect system.

Static time analysis approaches for concurrent programs can be divided into two main categories: those based on type-and-effect systems and those based on abstract interpretation – see references in [9]. Type-and-effect systems (i) collect constraints on type and resource variables and (ii) solve these constraints. The difference with respect to our approach is that we do not perform the analysis during the type inference. We use the type system for deriving behavioural types of methods and, in a second phase, we use them to run a (non compositional) analysis that returns cost upper bounds. This dichotomy allows us to be more precise, avoiding unification of variables that are performed during the type derivation. In addition, we notice that the techniques in the literature are devised for programs where parallel modules of sequential code are running. The concurrency is not part of the language, but used for parallelising the execution.

Abstract interpretation techniques have been proposed addressing domains carrying quantitative information, such as resource consumption. One of the main advantages of abstract interpretation is the fact that many practically useful optimization techniques have been developed for it. Consequently, several well-developed automatic solvers for cost analysis already exist. These techniques either use finite domains or use expedients

(widening or narrowing functions) to guarantee the termination of the fix-point generation. For this reason, solvers often return inaccurate answers when fed with systems that are finite but not statically bounded. For instance, an abstract interpretation technique that is very close to our contribution is [2]. The analysis of this paper targets a language with the same concurrency model as ours, and the backend solver for our analysis, CoFloCo, is a slightly modified version of the solver used by [2]. However the two techniques differ profoundly in the resulting cost equations and in the way they are produced. Our technique computes the cost by means of a type system, therefore every method has an associated type, which is parametric with respect to the arguments. Then these types are translated into a bunch of cost equations that may be *composed* with those of other methods. So our approach supports a technique similar to *separate compilation*, and is able to deal with systems that create statically an unbounded but finite number of nodes. On the contrary, the technique in [2] is not compositional because it takes the whole program and computes the parts that may run in parallel. Then the cost equations are generated accordingly. This has the advantage that their technique does not have any restriction on invocations on arguments of methods that are (currently) present in our one.

We finally observe that our behavioural types may play a relevant role in a cloud computing setting because they may be considered as abstract descriptions of a method suited for SLA compliance.

8 Conclusions

This article presents a technique for computing the time of concurrent object-oriented programs by using behavioural types. The programming language we have studied features an explicit cost annotation operation that define the number of machine cycles required before executing the continuation. The actual computation activities of the program are abstracted by **job**-statements, which are the unique operations that consume time. The computational cost is then measured by introducing the notion of (strong) *t-stability* (cf. Definition 1), which represents the ticking of time and expresses that up to t time steps no control activities are possible. A Subject Reduction theorem (Theorem 1), then, relates this stability property to the derived types by stating that the consumption of t time steps by **job** statements is properly reflected in the type system. Finally, Theorem 2 states that the solution of the cost equations obtained by translation of the types provides an upper bound of the execution times provided by the type system and thus, by Theorem 1, of the actual computational cost.

Our behavioural types are translated into so-called cost equations that are fed to a solver that is already available in the literature – the CoFloCo solver [4]. As discussed in Remark 1, CoFloCo cannot handle rational numbers with variables at the denominator. In our system, this happens very often. In fact, the number pc of processing cycles needed for the computation of a **job**(pc) is divided by the speed s of the machine running it. This gives the cost in terms of time of the **job**(pc) statement. When the capacity is not a constant, but depends on the value of some parameter and changes over time, then we get the untreatable rational expression. It is worth to observe that this is a problem of the solver (otherwise very powerful for most of the examples), while

our behavioural types carry enough information for computing the cost also in these cases. We plan to consider alternative solvers or a combination of them for dealing with complex examples.

Our current technique does not address the full language. In particular we are still not able to compute costs of methods that contain invocations to arguments which do not live in the same machine (which is formalized by the notion of cog in our language). In fact, in this case it is not possible to estimate the cost without any indication of the state of the remote machine. A possible solution to this issue is to deliver costs of methods that are parametric with respect to the state of remote machines passed as argument. We will investigate this solution in future work.

In this paper, the cost of a method also includes the cost of the asynchronous invocations in its body that have not been synchronized. A more refined analysis, combined with the resource analysis of [5], might consider the cost of each machine, instead of the overall cost. That is, one should count the cost of a method *per* machine rather than in a cumulative way. While these values are identical when the invocations are always synchronized, this is not the case for unsynchronized invocation and a disaggregated analysis might return better estimations of virtual machine usage.

References

1. E. Albert, F. S. de Boer, R. Hähnle, E. B. Johnsen, R. Schlatte, S. L. Tapia Tarifa, and P. Y. H. Wong. Formal modeling of resource management for cloud architectures: An industrial case study using Real-Time ABS. *Journal of Service-Oriented Computing and Applications*, 8(4):323–339, 2014.
2. E. Albert, J. C. Fernández, E. B. Johnsen, and G. Román-Díez. Parallel cost analysis of distributed systems. In *Proceedings of SAS 2015*, volume 9291 of *Lecture Notes in Computer Science*. Springer, 2015. To appear.
3. R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic. Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. *Future Generation Comp. Sys.*, 25(6):599–616, 2009.
4. A. Flores Montoya and R. Hähnle. Resource analysis of complex programs with cost equations. In *Proceedings of 12th Asian Symposium on Programming Languages and Systems*, volume 8858 of *Lecture Notes in Computer Science*, pages 275–295. Springer, 2014.
5. A. Garcia, C. Laneve, and M. Lienhardt. Static analysis of cloud elasticity. In *Proceedings of PPDP 2015*, 2015.
6. R. Hähnle and E. B. Johnsen. Resource-aware applications for the cloud. *IEEE Computer*, 48(6):72–75, 2015.
7. E. B. Johnsen, K. I. Pun, M. Steffen, S. L. Tapia Tarifa, and I. C. Yu. Meeting deadlines, elastically. In *From Action Systems to Distributed Systems: the Refinement Approach*. CRC Press, 2015. To Appear.
8. E. B. Johnsen, R. Schlatte, and S. L. Tapia Tarifa. Integrating deployment architectures and resource consumption in timed object-oriented models. *Journal of Logical and Algebraic Methods in Programming*, 84(1):67–91, 2015.
9. P. W. Trinder, M. I. Cole, K. Hammond, H. Loidl, and G. Michaelson. Resource analyses for parallel and distributed coordination. *Concurrency and Computation: Practice and Experience*, 25(3):309–348, 2013.