

Alma Mater Studiorum — Università di Bologna

DOTTORATO DI RICERCA IN COMPUTER SCIENCE
AND ENGINEERING

Ciclo: XXIX

Settore Concorsuale di Afferenza: 01/B1
Settore Scientifico Disciplinare: INF/01

Static analysis of concurrent programs based on behavioral type systems

Presentata da: Abel García Celestrín

Coordinatore Dottorato:
Paolo Ciaccia

Relatore:
Cosimo Laneve

Esame finale anno 2017

Abstract

The strength of program static analysis techniques lies on its ability to detect faulty behaviors prior to the execution. This ability requires that the analysis process foresees any possible runtime scenario. A task which is even more complex in the case of concurrent programs, because of the number of alternatives introduced by the usual nondeterminism. In this particular case, some of the most common faulty behaviors are those about erroneous usage of resources, presence of deadlocks and data race conflicts.

Behavioral type systems for programming languages provide a strong mechanism for reasoning on programs actions at static time. In this thesis we discuss two static analysis techniques based on this approach. The first one, targets the resource usage in an ad-hoc language with full-fledged operations for acquiring and releasing virtual machines. The second one, targets the deadlock analysis of Java programs.

In both cases we provide a formal proof of correctness, along with prototype implementations that allow practically to test the feasibility of these solutions. These prototypes have also allowed assessing these techniques against others existing in the literature obtaining very encouraging results.

Acknowledgements

During your life as Ph.D. Student several people crosses paths with you, in such a moment of your life you get to appreciate even the slightest act of support. To all of them I want to express my appreciation and gratefulness.

That said, there are some persons that in one way or another become special by living this experience by your side.

A very deep gratitude goes to my advisor Prof. Cosimo Laneve, I will always take with me the rewarding discussions, and the enormous support in the process of writing and describing scientific results. Of course, I will not forget either the passion for the science, and the stylish approach in presentations (which I have, secretly, tried to imitate a lot).

I also want to thank Elena Giachino for being not only a colleague but also a friend since day one, with whom I have shared amazing moments both professionally and personally speaking.

I want to thank my Ph.D. colleagues: Valeria, Stefano, Francesco, Vincenzo and Vincenzo "the master", and Giacomo, for contributing to create a really nice work environment, and for bringing a nice warm atmosphere to our cold basement. A special gratitude goes to Saverio Giallorenzo for being always two steps forward but showing the rest of us the right way ahead.

I feel very grateful as well to all the people in the Envisage Team. In particular, to Reiner Hanle who invited me to Germany and to his amazing team that make me feel at home there. And to Einar Broch and Elvira Albert who kindly accepted being part of the reviewing committee of this thesis.

Finally, I want to thank Yisleidy Linares (my Yiyi) who is my partner in life but also happens to be a Ph.D. colleague. I don't know how these three years would have resulted without having here supporting me every single day.

Contents

Abstract	i
Acknowledgements	ii
Contents	iii
1 Introduction	1
1.1 Static analysis	1
1.2 Behavioral type systems based analysis	3
1.3 Problems tackled in this work	5
1.4 Aim, objectives, and contribution of this work	6
1.4.1 Objectives	6
1.4.2 Contribution	7
1.5 Document contents	7
I The static resource analysis problem	9
2 Machine usage upper bounds in <code>vm1</code>	11
2.1 Introduction	11
2.2 Motivation and Problem Statement	13
2.3 The language <code>vm1</code>	15
2.3.1 Syntax	15
2.3.2 Semantics	17
2.3.3 Examples	20
2.4 Determinacy of method's arguments release	21
2.5 The behavioral type system of <code>vm1</code>	24
2.6 The analysis of <code>vm1</code> behavioral types	30
2.6.1 Arguments' identities	31
2.6.2 (Re)computing argument's states	32
2.6.3 The translation function	34

2.7	Correctness	39
2.7.1	The extension of the typing to configurations	40
2.7.2	Runtime types evolution	41
2.7.3	The subject reduction theorem	43
2.7.4	Correctness of the cost computation	46
2.7.5	Final demonstration	50
2.8	Related Work	53
2.9	Conclusions	55
3	The SRA tool	57
3.1	Introduction	57
3.2	Type inference	58
3.3	Translation of cost equations	60
3.4	An illustrative example	61
3.5	Related tools and assessments	62
3.6	SRA Deliverable	66
3.7	Conclusions	67
II	The deadlock analysis problem	69
4	Deadlock detection in JVMML	71
4.1	Introduction	71
4.2	Motivation and Problem Statement	74
4.3	The language JVMML _d	76
4.3.1	Syntax	77
4.3.2	Informal semantics	77
4.3.3	JVMML _d deadlock example: the <code>Network</code> class	78
4.4	Preliminaries	80
4.5	A behavioral type system for JVMML _d	86
4.5.1	Typing rules	86
4.5.2	The abstract behavior of the <code>Network</code> class	89
4.6	Analysis	91
4.7	<code>wait</code> , <code>notify</code> and <code>notifyAll</code> methods	92
4.8	Type system extensions	94
4.8.1	Static members	94
4.8.2	Exceptions and exception handlers	96
4.8.3	Arrays	98
4.8.4	Recursive data types	101
4.8.5	Inheritance	103
4.8.6	Alternative concurrent models and native methods	105

4.9	Full proof of correctness	106
4.9.1	Correctness overview	106
4.9.2	The JVMML formal semantics	107
4.9.3	Lam semantics	108
4.9.4	Flattening and circularities	111
4.9.5	Later stage relation	111
4.9.6	Runtime typing	112
4.9.7	Subject Reduction	113
4.9.8	Deadlocks and circularities	120
4.10	Related work	121
4.10.1	Static approaches	121
4.10.2	Non-static approaches	123
4.11	Conclusions	124
5	The JaDA tool	127
5.1	Introduction	127
5.2	Type inference	128
5.3	Analysis of circularities	132
5.4	Related tools and assessments	133
5.5	JaDA Deliverable	135
5.5.1	Customization	137
5.6	Conclusions	138
6	General conclusions	141
6.1	Conclusions	141
6.2	Future work	142
	Bibliography	145

Chapter 1

Introduction

In this Chapter we present a general overview of the concepts and problems treated in this thesis. We start with a brief description of the purpose of program static analysis, and in particular, the behavioral type system based approach. Then, we concisely describe the two program verification problems that are solved by the static techniques presented in this work. Finally, we conclude with the main goals of this thesis and the structure of this document.

1.1 Static analysis

There are two big well-defined groups of approaches for program analysis: static analysis techniques and runtime analysis techniques. In addition, there are as well hybrid techniques that combine the previous two.

The main characteristic of static analysis techniques is that programs are verified without actually being executed. Usually such analyses are performed on the source code or, in other cases, on some form of the *object code* (the intermediate language generated by some compilers like JVMIL in the case of the Java virtual machine). The term *static analysis* is, in general, applied to a verification performed in an automatic way. There are human-aided static analysis tools as well, but these are usually referred in the literature as program understanding, program comprehension, or code review.

The vast majority of the existing programming languages use some sort of static analysis for detecting potential errors prior to the execution time. The simplest examples are perhaps strong-typed languages, that define a type system that is checked at compiling time for ensuring type safety.

However, other kind of properties (i.e. resource usage upper bounds, deadlock freedom, program termination and data race conflicts) may be much more difficult to verify.

The advantage of the static analysis tools is the possibility to reason about the whole set of possible program executions. Theoretically, such checking ensures a certain quality before the software deployment stage. In general, the static analysis is as complex as the properties targeted by the verification. Logically, when these properties become very complex, the main drawbacks of these approaches appear.

The main issues with static analysis techniques is that analyzing all possible program executions takes a considerable amount of time. It is important to consider that since static techniques take place during compilation time they are not as time sensitive as dynamic analysis techniques that take place during the program runtime. Still, even for not very big programs, analyzing the whole set of possible executions may lead to an exponential unmanageable problem. Therefore, for non-trivial programs, some level of abstraction is needed. Abstraction allows to scale the large set of possible program runs into a smaller (analyzable) model. Implicitly, this produces a common trade-off in static analysis techniques between precision and abstraction.

Formal methods. Program verification has been targeted by the Computer Science community since the very beginning. Perhaps one of the very first problems in this sense was Alan Turing's halting problem [57]. To this aim, several formal methods have been developed oriented to the analysis of computer programs. A static analysis technique may be purely based on one of these techniques, or on a combination of them. The following list briefly summarizes some of the most used formal methodologies:

Abstract interpretation: A technique based on the interpretation of an abstraction of the language semantics. This technique often uses over-approximations of the underlying program model in order to considerably reduce the set of possible program executions.

Model checking: This technique is built around the representation of a program by a finite set of states. These states are then analyzed separately and the resulting conclusions are combined to draw the overall verification of the program.

Symbolic execution: This technique tries to fill the gap between the static and dynamic analysis. The main idea is to group all the possible inputs of the program into sets that will produce similar behaviors. Afterwards, a symbolic representation of each set of inputs is evaluated by a symbolic interpreter. The final conclusions are drawn from the results of this symbolic execution.

Dataflow analysis: This is a fix-point based technique. Programs are represented by a control flow graph that corresponds to the source code. The states at each node of the program are then calculated repeatedly until a stable state is reached. Then, the verification is performed on each one of the control flow graph nodes.

1.2 Behavioral type systems based analysis

Behavioral type systems [11] provide a strong mechanism for reasoning on a program possible execution. With respect to standard type systems, the behavioral ones usually enrich standard types by adding some dynamic information of the effects of each expression.

Generic approach

The overall approach can be divided in four stages: *(i)* Parsing of the program, *(ii)* Typing process, *(iii)* Simplification of the program and *(iv)* Verification of the program.

The first step is usually performed after traditional static checks, thus the program is considered syntactically correct and with no obvious semantic errors. The main goal of this process is to obtain a representation of the program that is more suitable for the upcoming processes than the textual representation.

During the typing process the program behavioral type is obtained. This process can be fully automatic, partially automatic or manual. The typing process might not succeed, in that case the program could not be analyzed.

After a behavioral type is obtained, it is derived into an abstract simplified representation of the program. This abstraction drops features that are unrelated to the targeted analysis and focuses on those that are relevant. For example, in deadlock analysis, the abstraction may be an object dependency graph, for resource analysis it may be a set of cost equations.

Finally, the targeted properties are verified over the abstracted program. The verification process is only tied to the abstract program representation easing the use of third-party, and possibly more general, analyzers.

Advantages

The main advantage of this approach is its modularity. The type system is in general bounded to the language targeted by the analysis. The final verification is, in turn, tied to the abstract representation of the program.

This makes possible to apply the same analysis to different programming languages by having proper typing mechanisms producing the same kind of abstract representation.

This idea is illustrated in Figure 1.1. The figure corresponds to the deadlock analysis scheme for three different programming languages **Java**, **Scala** and **ABS**¹. All three languages share the verification stage, since all three share the same simplified representation. However, there are two typing processes, one for **Java** and **Scala**, and the other for **ABS**. The first two languages can use the same inference mechanism since both of them are compiled to **Java** bytecode.




Figure 1.1: General approach of a behavioral type system based analysis

Another rather important advantage is that one can reason on the correctness of the analysis also in a modular way. For example, one can focus on two aspects separately: the typing and the abstraction, and the analysis. In the case of the typing process and its abstract representation, the correctness is usually proved by means of a subject reduction technique [24]. The correctness for the analysis will depend on the analysis technique. For example, in the case of the work described in Part I the correctness of the verification relies on the third party cost equation solver. In the case of the work described in Part II the analysis is based on a formally demonstrated decision algorithm.

¹<http://abs-models.org/>

1.3 Problems tackled in this work

This work presents the theoretical basis and the prototype implementation of two static analysis tools based on behavioral type systems. The two problems targeted by these tools are the analysis of resource usage upper-bounds and the deadlock analysis. In both cases, the targeted language is an object-oriented language that allows concurrency operations.

The description of both tools is presented in the order they were conceived. This is an important remark, because the experience obtained in the development of the first tool has been key in the conception of the second one. In practice, it can be appreciated that the second of these tools presents a much deeper level of complexity. Also, in this second tool the result is more palpable since it can be applied to *real world* programming languages.

The analysis of virtual machines usage. In the first problem we analyze a domain specific language called `vm1`. This is a rather simplistic language but expressive enough to allow the dynamic acquisition (creation) and releasing (deletion) of objects representing virtual machines. In this language, the operations on virtual machines can occur in parallel, and more over, virtual machine objects can run tasks that may, in turn, operate on other virtual machines. This scenario is inspired on the elastic features of the Cloud Computing environments. The purpose of the analysis of this language is to make an (over) estimation of the maximum number of virtual machines that may coexist during the program execution.

The resource analysis problem itself has received lots of attention from the scientific community. However, very few approaches target languages in which are both present, the concurrency feature and the explicit resource removal operation. Moreover, to the best of our knowledge, there are not previous works targeting the static analysis of code exploiting the Cloud elasticity features.

The analysis of deadlocks in Java bytecode. There have been several works targeting the analysis of deadlocks in the Java language and its byte code (JVML). The existing techniques, however, target either a very reduced subset of the language or take little to none consideration about the soundness of their approach. An interesting remark, is also the fact, that to the best of our knowledge none of the existing analyzers reports the analysis of Scala programs, a language that is also compiled into JVML.

1.4 Aim, objectives, and contribution of this work

The main goal of this thesis is to present the strength of the behavioral type system based approach for the static analysis in concurrent programs. To accomplish such purpose, we focus on two program verification problems of high relevance and with application in the newest developments in the Computer Science world.

It is also part of the aim of this work to produce competitive solutions to these problems at the level of other state of the art solutions. We also expect these solutions to be fully automatic and capable to perform with very little human interaction. This is a key feature for the purpose of analyzing non-trivial programs

1.4.1 Objectives

The main objective of this thesis is to provide valid tools for the analysis of the problems discussed in Section 1.3. Technically, the process of building such tools, with the approach followed in this work, involves the following partial objectives:

- **Definition of the behavioral type systems:** To design the static semantic rules that define the type systems for both of the problems tackled in this work.
- **Design and implementation of automatic behavioral type inference:** To implement a fully automatic mechanism for the inference of program types, according to the behavioral type systems designed and the targeted languages.
- **Proofs of correctness of the behavioral type systems:** To demonstrate the correctness of the behavioral type systems in characterizing the underlying programs and producing correct program abstractions without losing key information for the analysis.
- **Design and implementation of the analysis of the behavioral type systems:** To provide an automatic mechanism for drawing conclusions from the abstractions produced during the typing process.
- **Comparison of the result of the proposed tools with existing similar ones:** To demonstrate the effectiveness of the resulting tools by testing them against other related approaches.

1.4.2 Contribution

Two main contributions can be derived from the results of this work:

- i the approach for the static analysis of resource utilization in a concurrent and object oriented language with explicit resource release operations.
- ii the approach for the static analysis of deadlocks in Java bytecode.

Each of these solutions presents a custom and complex behavioral type system. The instrumentation of these type systems as well as the proof of correctness behind them are also a remarkable part of the results obtained in this work.

The design of these solutions cover issues that are typical in the static analysis of programs. The ways in which these issues are resolved constitute also relevant contributions of this work.

Finally, this work also provides a prototype implementation of the two solutions described. These prototypes constitute a simple way to interact with the results of these approaches making easy to compare it with previous and future equivalent techniques.

1.5 Document contents

The current document intends to provide an overall vision of the static analysis techniques based on behavioral type systems, and to present two successful applications of this approach. Following this introduction, the document is divided in two parts of two chapters each:

Part I: The static resource analysis problem

This part is dedicated to the analysis of resources in a concurrent object oriented language. The content is divided in two chapters.

Chapter 2: Machine usage upper bounds in `vm1`, describes a variation of the static resource analysis problem applied to the Cloud Computing scenario. In particular, this problem focuses in the detection of upper bounds for the number of virtual machines necessary to execute a program in this environment. In this chapter, we describe a custom language: `vm1`, which allows modeling the dynamic creation and release operations of the Cloud's elasticity feature. We then discuss a behavioral type system for abstracting virtual machine usage on these programs, and the subsequent

analysis of these types. This chapter concludes with the proof of correctness of the overall process.

Chapter 3: The **SRA** tool, describes a resource analysis tool based on the theory presented in Chapter 2. We discuss the main implementation details of this tool, specially the type inference process and the integration with an external cost equations solver. We present illustrative examples of the technique, and we also review the related literature and compare the tool and its technique with some of the existing approaches.

Part II: The deadlock analysis problem

This part is devoted to the deadlock analysis problem in Java programs. The content is divided in two chapters.

Chapter 4: Deadlock detection in **JVML**, describes the static deadlock analysis problem in **Java** programs. In this chapter, we present a behavioral type system for abstracting dependencies in **JVMLd**, a subset of the **Java** bytecode. We also present the algorithm for reasoning on the presence of circularities in the program dependencies. The chapter concludes with the extension of the behavioral type system to the full set of **Java** features and the proof of correctness of the overall approach.

Chapter 5: The **JaDA** tool, describes the **Java** deadlock analysis tool. We discuss the type inference process and the implementation of the circularities detection algorithm. In this chapter, we also review the related literature and compare the tool and its technique with some of the existing approaches including a commercial grade tool. We conclude with the description of the **JaDA** tool deliverables.

This document concludes with Chapter 6, where we present the main contributions of this thesis and describe some of the future possible lines of research that can be derived from it.

Part I

The static resource analysis problem

Chapter 2

Machine usage upper bounds in vml

Summary

We propose a *static analysis technique* that computes upper bounds of virtual machines usage in a *concurrent* language with explicit acquire and release operations over this kind of objects. In our language, the creation and release operations can occur in separated and possibly concurrent contexts (by passing virtual machine objects as arguments of invocations). Moreover, the objects representing virtual machines are active, that may be asynchronously hosting operations with side effects on the overall resource utilization. Our technique is modular and consists of (i) a type system associating programs with behavioral types that record relevant information for resource usage (creations, releases, and concurrent operations), (ii) a translation function that takes behavioral types and returns cost equations, and (iii) the integration with an automatic off-the-shelf solver for the cost equations.

2.1 Introduction

An accurate assessment of the resource usage in a computer program could reduce, for example, energy consumption and allocation costs. Two criteria that 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.

While it is relatively easy to manually estimate worst-case costs for simple code examples, extrapolating this information for fully real-life complex programs could be cumbersome and highly error-sensitive. The first attempt

about the analysis of resource usage dates back to Wegbreit’s pioneering work in 1975 [58], which develops a technique for deriving closed-form expressions out of programs. The evaluation of these expressions would return upper-bound costs that are parametrized by programs’ inputs.

Wegbreit’s contribution has two limitations: it addresses a simple functional language and it does not formalize the connection between the language and the closed-form expressions. A number of techniques have been developed afterwards to cope with more expressive languages (see [6, 20]) and to make the connection between programs and closed-form expressions precise (see [23, 41]). A more detailed discussion of the related work in the literature is presented in Section 2.8.

The existing cost analysis techniques vary also from the point of view of the resource targeted. The most common targets are probably the memory related resources like the heap space, and the number of steps in the program execution. However, some of the existing techniques also give certain abstraction on the type of the resource analyzed, by allowing code annotations with cost related expressions. Some attempts have even considered the analysis of the time necessary for the program execution [33].

This work targets the analysis of the usage of virtual machines represented by first order objects in a custom domain specific language called `vm1`. The latter is a concurrent object-oriented language that includes explicit acquire and release operations on virtual machines object. These operations can occur in separated and possibly concurrent contexts (by passing virtual machine objects as arguments of invocations).

Our technique is modular and consists of *(i)* a type system associating programs with behavioral types that record relevant information for resource usage (creations, releases, and concurrent operations), *(ii)* a translation function that takes behavioral types and returns cost equations, and *(iii)* the integration with an automatic off-the-shelf solver for the cost equations.

To the best of our knowledge, current cost analysis techniques address (concurrent) languages where resource usage is always cumulative. We have found that only one of the existing techniques considers the scenario where both concurrent and resource removal operations are possible. That is the case of [8], however, our problem settings differ from these in two main aspects: i) the resources targeted by our technique are program objects that may carry, concurrently, operations that acquire and release other resources; and ii) in our case it is possible to have the acquire and release operations of a resource happening in different contexts by passing the objects as arguments of other possibly asynchronous method invocations.

Chapter contents

This Chapter starts by describing, in Section 2.2, the so called *Cloud Computing elasticity* concept that constitutes the inspiration of this work. In Section 2.3, we present `vm1`: a simplified concurrent object-oriented language featuring virtual machine (VM) objects. In Section 2.4 we discuss some technical limitations that simplify, from the practical point of view, the analysis of `vm1` programs. Section 2.5 and 2.6 present the definition of a behavioral type system for `vm1` programs and its further analysis for estimating VM usage upper bounds. We then discuss, in Section 2.7, the correctness of this approach, supported by the formal proof of the soundness of the type system and the corresponding analysis. In Section 2.8 we review some of the main contributions in the literature related to the resource analysis problem. We conclude with the main remarks about this approach.

2.2 Motivation and Problem Statement

Cloud Computing introduces the concept of elasticity, whose main feature is the possibility for the resources, namely virtual machines (VM), to scale according to the software needs. There are two kinds of scaling, vertical and horizontal. Vertical Scaling means changing the deployed VM, adding more memory or disk space or using a more powerful processor, this kind of elasticity is, in general, associated to a disruptive workflow that stops the software execution and restarts it after the changes.

On the other hand, Horizontal Scaling is when the underlying number of VMs changes to fulfill the demand, this kind of scaling is a seamless process for the final user since it does not affect the availability of a service. Horizontal Scaling is the suggested scaling method by main Cloud providers like Amazon, Google and Microsoft Azure. This is appreciated in their pricing models that allow hiring, on demand, instances which are payed for the time they are in use; and also in their APIs that include instructions for requesting and releasing these VM instances on the fly.

In contrast to the elasticity concept, Cloud providers also give to the user the possibility to reserve a number of VM instances for a fixed period of time. This approach is used often by clients that migrate their existing solutions from private data centers to the Cloud. The main problem of this approach is the overcapacity or under-capacity of the allocated resources. The reserved VMs may spend important part of the time in an idle status, or the reserved VMs may be not enough to fulfill the demand. The table in Figure 2.1 shows some of the pricing possibilities given by the main Cloud providers. Notice

that the **on demand** price is a bit higher than the **reserved** alternative. However, the **on demand** model is more convenient when reserved machines may be idle more than the 40% of the time.

Class	Provider	Description	On Demand		1-Year Reserved		Saving
			Hourly Rate	Monthly Rate	Hourly Rate	Monthly Rate	
Micro	Azure	1 Core, 0.75 GB RAM	0.02	14.88			37.5%
	Amazon	1 Core, 1 GB RAM	0.013	9.672	0.009	6.696	30.8%
	Google	1 Core, 0.6 GB RAM	0.012	8.928	0.009	6.696	25.0%
Medium	Azure	1 Core, 3.5 GB RAM, 50 GB SSD	0.085	63.24			37.5%
	Amazon	1 Core, 3.75 GB RAM, 40 GB SSD	0.077	57.288	0.05	37.2	35.1%
	Google	1 Core, 3.57 GB RAM	0.063	46.872	0.045	33.48	28.6%
Large	Azure	16 Core, 112 GB RAM, 800 GB SSD	1.54	1145.76			37.5%
	Amazon	32 Core, 60 GB RAM, 640 GB SSD	1.68	1249.92	1.16	863.04	31.0%
	Google	16 Core, 104 GB RAM	1.184	880.896	0.82	610.08	30.7%

<http://azure.microsoft.com/en-us/pricing/calculator/?scenario=virtual-machines>
<http://aws.amazon.com/ec2/purchasing-options/dedicated-instances/>
<https://cloud.google.com/pricing/>

Figure 2.1: Some pricing options of main Cloud Providers, January 2016

Cloud computing elasticity APIs features acquire and release operations for the management of *On Demand* VMs (see for instance the *Amazon Elastic Compute Cloud* [50] or the *Docker Fiware*) that can be invoked dynamically from running code. The release operation happens to be a very powerful programming artifact, allowing to downscale a program’s resource demand. It is worth noticing that, without a full release operation, the cost of a concurrent program may be modeled by aggregating the sets of operations that can occur in parallel, as in [7]. A full-fledged release operation, however, makes possible to delegate to other methods (possibly concurrent) the release of 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 will depend on the state of the machines in input:

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

- it may be 0 when the two machines have been already released.

In this case, one might over-approximate the cost of `double_release` to 0. However, this leads to disregard releases and makes the analysis (too) imprecise. Moreover, consider the case where the machine objects can be hosting processes that either create or release other virtual machines. Ignoring a machine release may imply to assume the process hosted on it is still running, thus executing in parallel with the rest of the environment. This will have a serious impact in the calculation of the usage of these resources.

Problem statement

Lets us consider the following problem:

Given a pool of on demand VM instances and a computer program that may indistinctly acquire and release these instances, what is the minimum cardinality of the pool that ensures the program runs without interruptions caused by the lack of available VMs? Assume that a VM can be re-acquired if and only if it has been previously released.

A solution to the previous problem has a direct application for both Cloud Providers and Cloud Costumers. For the former, it represents the possibility to know in advance how many real resources to allocate for an specific service. For the latter, it represents the possibility to pay *exactly* for the resources that are needed.

2.3 The language vml

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

2.3.1 Syntax

In `vml` we distinguish between *simple types* T , which are either integers `Int` or virtual machines `VM`, and *future types* `Fut<T>`, which type asynchronous

¹See [35] for more details on a similar language

invocations. The future argument T is instantiated with the return type of the invoked function. We use F to range over simple and future types. The notation $\overline{T x}$ denotes any finite sequence of *name declarations* $T x$ separated by commas. Similarly we write $\overline{T x}$; for a finite sequence of name declaration separated by semicolons.

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

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

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

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

A (*pure*) expression e may be the reserved identifier **this**, a virtual machines identifier or an integer expression. Since our analysis will be parametric with respect to the inputs, we parse integer expressions in a careful way. In particular we split them into *size expressions* se , which are expressions in Presburger arithmetics [18] (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:

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

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

2.3.2 Semantics

The `vml` semantic is defined as a transition relation between *configurations*, noted cn and defined below

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

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 $cn \ cn$; the empty configuration is denoted by ϵ . The transition relation uses two infinite sets of names: *VM names*, ranged over by o, o', \dots and *future names*, ranged over by f, f', \dots . The function `fresh()` returns either a fresh VM name or a fresh future name; the context will disambiguate between the two. The elements of configurations are:

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

Runtime values v are either integers or virtual machines and future names, or \perp , meaning an un-computed value, or an erroneous value err . The fol-

lowing auxiliary functions are used in the semantic rules (we assume a fixed vml program):

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

The transition relation rules are collected in Figure 2.2. The rules are almost standard, except those about the management of virtual machines and method invocations, which we are going to discuss.

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

Rule (ASYNC-CALL) defines asynchronous method invocation $x = e!m(\bar{e})$. This rule creates a fresh future name that is assigned to the identifier x . Rule (BIND-MTD) applies an invocation message by adding to the alive callee virtual machine a process corresponding to the called method. In case the callee virtual machine is not live, either (BIND-MTD-ERR) or (BIND-PARTIAL) is applied, which binds *err* to the future name. Rule (READ-FUT) allows the caller to retrieve the value returned by the callee. It is worth noticing that the semantic of **get** is different from that of **ABS** [46] or of [35] because it is not blocking with respect to other processes waiting to be executed on the

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

Figure 2.2: Semantics of vml.

same machine, see rule (ACTIVATE-GET). In fact, deadlock freedom is out of the scope of this contribution; in any case we refer to [35] for an algorithm (and a prototype) verifying deadlock freedom in ABS.

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

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

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

Problem statement reformulation

After discussing the semantics of vml we can now reformulate the Problem Statement from Section 2.2 in a more formal way. We start by defining the concept of *alive machines*:

Definition 3.1 (Alive machines). *Given a configuration cn , a term $vm(o, \top, p, q) \in cn$ is called alive machine in cn . Let $\mathbf{alive}(cn)$ be the number of different alive machines in cn .*

At this point we can rewrite the objective of this work as:

to define a technique such that, given a configuration cn , it returns a n satisfying, for every $cn \longrightarrow^* cn'$, $\mathbf{alive}(cn') \leq n$ (n is an upper bound to the alive machines in computations rooted at cn).

2.3.3 Examples

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

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

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

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

The analysis of `costly_fact` has been easy because the `release` operation is applied to a locally created virtual machine. Yet, in `vml`, `release` can be also applied to method arguments and the presence of this feature in concurrent codes is a major source of difficulties for the analysis. A paradigmatic example is the `double_release` method discussed in Section 2.2 that may have either a cost of -2 or of -1 or of 0 .

It is worth observing that, while over-approximations (*e.g.* not counting releases) return (too) imprecise costs, under-approximations may return

wrong costs. For example, the following method creates two virtual machines and releases the second one with `this!double_release(x,x)` before the recursive invocation.

```

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

```

The cost of `fake_method(n)` should be n . However this is not the case if `double_release` is under-approximated with cost -2 : In such case, one would wrongly derive a cost 0 of `fake_method(n)`. In Section 3.4 we consider an erroneous `fake_method` that increases the argument of the recursive invocation (instead of decreasing it) and discuss the corresponding cost equations returned by our technique. We notice that, in this case, the amount of virtual machines used by the erroneous method is infinite. The aim of the following sections is to present a technique for determining the cost of method invocations. Such a technique has to be sensible to the identity, and to the state of method's arguments before and after the invocation.

2.4 Determinacy of method's arguments release

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

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

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

```

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

```

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

To analyze such program, one would just have to extend the presented approach with *guarded effects*, i.e., mapping from execution branches to effects, which can precisely identify the effects of each execution branch of the method.

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

```

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

```

the method's effects of `foo2` is not deterministic because `double_release` might still be running *after* `foo2` termination. This means that the termination of `foo2` has no impact on the termination of `double_release`. Therefore, the effect of `foo2` is empty – the caller cannot assume that x and y have been released upon synchronizing with `foo2`. Overall, this simplification supports a more precise analysis. In order to drop this restriction is necessary to add the unsynchronized invocations to the effects of the method. This idea is explained in [35] and also in Section 4.6 The underlying intuition is to pass the effects of unsynchronized invocations to the context, and to consider context effects along with the local effects of each method during the analysis phase.

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

```

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

```

```

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

```

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

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

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

```

VM identity(VM x) { return x; }

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

```

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

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

Simplifications 1, 3, and 4 are enforced by the type system in Section 2.5, in particular simplification 1 by rule (T-METHOD), simplification 3 by rules

\circ	$::=$	$- \mid \alpha$	basic value
\dagger	$::=$	$\alpha \downarrow \mid \partial \mid \perp \mid \top$	vm value
op'	$::=$	$+ \mid - \mid = \mid \leq \mid \geq \mid \wedge \mid \vee$	linear operation
$\mathfrak{r}, \mathfrak{s}$	$::=$	$\circ \mid se$	typing value
\mathfrak{z}	$::=$	$(\circ, \alpha, \mathbf{F}\dagger, \mathbf{R}) \mid \circ$	future value
\mathfrak{x}	$::=$	$- \mid \mathbf{F}\dagger \mid f \mid \mathfrak{z}$	extended value
\mathfrak{a}	$::=$	$0 \mid \nu\alpha \mid \nu f : \mathfrak{m} \alpha(\overline{\mathfrak{s}}) \mid \alpha^\checkmark \mid f_\circ^\checkmark$	atom
\mathfrak{c}	$::=$	$\mathfrak{a} \triangleright \Gamma \mid \mathfrak{a} \ddagger \mathfrak{c} \mid (se)\{\mathfrak{c}\} \mid \mathfrak{c} + \mathfrak{c}$	behavioral type

Figure 2.3: Behavioral Types Syntax

(T-INVOKE) and (T-RELEASE), and simplification 4 by rules (T-INVOKE) and (T-RETURN).

2.5 The behavioral type system of vml

In this context, *Behavioral types* are abstract codes highlighting the features of vml programs that are relevant for the resource cost analysis in Section 2.6. These types support compositional reasonings and are associated to programs by means of a set of type system rules that are defined in this section.

The syntax of behavioral 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 $\mathbf{F}, \mathbf{F}', \dots$. We assume that *err* is a *special* VM name representing the erroneous machine; therefore VM names will also range over *err*. The syntactic rules of these types are presented in Figure 2.3.

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

In order to have a more precise type of continuations, the leaves of behavioural types are labelled with *environments*, ranged over by Γ, Γ', \dots . Environments are maps from method names \mathfrak{m} to terms $\alpha(\overline{\mathfrak{r}}) : \circ, \mathbf{R}$, from names to extended values \mathfrak{x} , from future names to future values, and from VM names to extended values $\mathbf{F}\dagger$, which are called *vm states* in the following. We assume that $\Gamma(\text{err}) = \emptyset \perp$. These environments occurring in the leaves

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

Figure 2.4: Typing rules for pure expressions

are only used in the typing proofs and are dropped in the final types (method types and the main statement type).

VM states $\mathbb{F}\mathfrak{t}$ are a collection \mathbf{F} of future names plus the value \mathfrak{t} of the virtual machines. This \mathbf{F} specifies the set of parallel methods that are going to release the virtual machine; \mathfrak{t} defines whether the virtual machine is alive (\top), or it has been already released (\perp) or, according to scheduler's choices, it may be either alive or released (∂). VM values also include terms α and $\alpha\downarrow$. The value α is given to the argument machines of methods (they will be instantiated by the invocations – see the cost analysis in Section 2.6), the value $\alpha\downarrow$ is given to argument values that are returned by methods *and* can be released by parallel methods ($\alpha\downarrow$ will be also evaluated in the cost analysis). VM values are partially ordered by the relation \leq defined by

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

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

The type system uses judgments of the following form:

- $\Gamma \vdash e : \mathbb{x}$ for pure expressions e , $\Gamma \vdash f : \mathbb{z}$ for future names f , and $\Gamma \vdash \mathbf{m} \alpha(\bar{\mathbb{r}}) : \mathbb{O}, \mathbf{R}$ for methods.
- $\Gamma \vdash_{\mathbf{s}} z : \mathbb{x}, \mathfrak{c} \triangleright \Gamma'$ for expressions z , where \mathbb{x} is the value, \mathfrak{c} is the behavioral type for z and Γ' is the environment Γ *with updates* of names and future names.

- $\Gamma \vdash_{\mathbf{S}} s : \mathfrak{c}$, in this case the updated environments are inside the behavioural type Γ' , in correspondence of every branch of its.

The index \mathbf{S} in the judgments for expressions and statements defines the set of method's arguments – see rule (T-METHOD) – and is used in the rule (T-RETURN) in order to constrain that the returned machine, if any, does not belong to method's arguments (*cf.* Restriction 4 in Section 2.4).

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

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

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

We will also use the operation and notation below:

- $\mathbf{F}\mathfrak{t} \Downarrow$ is defined as follows:

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

and, in Section 2.6, 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 *multihole contexts* $\mathcal{C}[\]$ defined by the following syntax:

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

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

The type system for expressions is reported in Figure 2.4. It is worth noticing that this type system is not standard because (size) expressions containing method's arguments are typed with the expressions themselves. This is crucial in the cost analysis of Section 2.6. It is also worth observing that, by rule (T-PRIMITIVE), $\Gamma \vdash err : err$ (while $\Gamma(err) = \emptyset \perp$). This expedient allows us to save one rule when typing method invocations with either erroneous or released carriers (*cf.* rule (T-INVOKE-BOT)).

The type system for expressions with side effects and statements is reported in Figure 2.5. We discuss rules (T-INVOKE), (T-GET), (T-RELEASE), and (T-RETURN). Rule (T-INVOKE) types method invocations $e!m(\bar{e})$ by using a fresh future names. In particular, in the behavioral type $\nu f : \mathfrak{m} \alpha(\bar{\mathfrak{s}}) \rightarrow \beta$,

the (fresh) future name f is associated to the the method, the VM name of the callee, the arguments *and to the returned value*. This last value is fundamental when the invocation returns a new machine because, in this case, the type acts as a binder of β . Another important remark is about the value of f in the updated environment. This value contains the returned value, the VM name of the callee and its state, and the set of the arguments that the method is going to remove. The VM state of the callee will be used when the method is synchronized to update the state of the returned object, if any (see rule (T-GET)). It is also important to observe that the environment returned by (T-INVOKE) is updated with information about VM names released by the method: every such name will contain f in its state. Let us now discuss the constraints in the second and third lines of the premise of (T-INVOKE). As regards the second line, 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 α is the caller object α' : namely the callee is alive because it has been created by the caller or it is the caller itself,
- (ii) or $\Gamma(\alpha) \neq \emptyset\top$: this case has two subcases, namely either (ii.a) the callee is being released by a parallel method or (ii.b) it is an argument of the caller method – see rule (T-METHOD).

While in (i) we admit that the invoked method releases VM names, in case (ii) we forbid any release, as we discussed in Restriction 3 in Section 2.4.

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

Rule (T-GET) defines the synchronisation with a method invocation that corresponds to a future f . Let $(\diamond, \alpha, \mathbf{F}\mathfrak{t}, \mathbf{R})$ be the value of f in the environment. There are two cases: either (i) $\mathbf{R} \neq \emptyset$ or (ii) $\mathbf{R} = \emptyset$. In case (i), by Restriction 3, $\mathbf{F}\mathfrak{t}$, which is the value of the caller α when the method has been invoked, *is equal* to the value $\Gamma(\alpha)$ (and $\mathbf{F} = \emptyset$). In this case, if the method returns a new machine (*cf.* Restriction 4), its state must be $\emptyset\mathfrak{t}$ (notice that $\mathfrak{t} \sqcap \Gamma(\alpha) \Downarrow = \mathfrak{t}$). The case (ii) is more problematic because the caller may be released by a method running in parallel and, therefore, the possible returned virtual machine must record this information in its state.

$$\begin{array}{c}
\text{(T-ASSIGN-VAR)} \\
\frac{\Gamma(x) = \mathbb{x} \quad \Gamma \vdash_{\mathbb{S}} z : \mathbb{x}', \mathbb{C}}{\Gamma \vdash_{\mathbb{S}} x = z : \mathbb{C}[x \mapsto \mathbb{x}']} \\
\\
\text{(T-INVOKE)} \\
\frac{\Gamma \vdash e : \alpha \quad \Gamma \vdash \bar{e} : \bar{\mathbb{S}} \quad \Gamma \vdash \mathfrak{m} \alpha(\bar{\mathbb{S}}) : \mathbb{O}, \mathbb{R} \quad \Gamma \vdash \mathbf{this} : \alpha' \quad \Gamma(\alpha) \neq \mathbb{F}\perp \\
(\Gamma(\alpha) \neq \mathbb{O}\top \text{ and } \alpha \neq \alpha') \text{ implies } \mathbb{R} = \emptyset \\
\mathbb{R} \cap \left(\{\alpha'\} \cup \{\beta \mid f' \in \text{dom}(\Gamma) \text{ and } \Gamma(f') = (\mathbb{O}', \beta, \mathbb{x}, \mathbb{R}') \text{ and } \mathbb{R}' \neq \emptyset\} \right) = \emptyset \\
f \text{ fresh} \quad \Gamma' = \Gamma[\beta \mapsto (\{f\} \cup \mathbb{F}')\dagger]^{\beta \in \mathbb{R}, \Gamma(\beta) = \mathbb{F}'\dagger}}{\Gamma \vdash_{\mathbb{S}} e \mathfrak{lm}(\bar{e}) : f, \nu f : \mathfrak{m} \alpha(\bar{\mathbb{S}}) \triangleright \Gamma'[f \mapsto (\mathbb{O}, \alpha, \Gamma(\alpha), \mathbb{R})]} \\
\\
\text{(T-INVOKE-BOT)} \\
\frac{\Gamma \vdash e : \alpha \quad \Gamma \vdash \bar{e} : \bar{\mathbb{S}} \quad \Gamma(\alpha) = \mathbb{F}\perp \quad f \text{ fresh}}{\Gamma \vdash_{\mathbb{S}} e \mathfrak{lm}(\bar{e}) : f, \nu f : \mathfrak{m} \alpha(\bar{\mathbb{S}}) \triangleright \Gamma[f \mapsto (\mathit{err}, \alpha, \mathbb{F}\perp, \emptyset)]} \\
\\
\text{(T-GET)} \\
\frac{\Gamma \vdash x : f \quad \Gamma \vdash f : (\mathbb{O}, \alpha, \mathbb{F}\dagger, \mathbb{R}) \\
\mathbb{R}' = f\nu(\mathbb{O}) \setminus \mathbb{R} \quad \dagger' = \dagger \cap (\Gamma(\alpha) \Downarrow) \\
\Gamma' = \Gamma[\beta \mapsto \mathbb{O}\perp]^{\beta \in \mathbb{R}}[\beta' \mapsto \mathbb{O}\dagger']^{\beta' \in \mathbb{R}'}}{\Gamma \vdash_{\mathbb{S}} x.\mathit{get} : \mathbb{O}, f_{\mathbb{O}} \checkmark \triangleright \Gamma'[f \mapsto \mathbb{O}]} \quad \text{(T-GET-DONE)} \\
\frac{\Gamma \vdash x : f \quad \Gamma \vdash f : \mathbb{O}}{\Gamma \vdash_{\mathbb{S}} x.\mathit{get} : \mathbb{O}, \mathbb{O} \triangleright \Gamma} \\
\\
\text{(T-NEW)} \\
\frac{\beta \text{ fresh}}{\Gamma \vdash_{\mathbb{S}} \mathbf{new VM}() : \beta, \nu \beta \triangleright \Gamma[\beta \mapsto \mathbb{O}\top]} \\
\\
\text{(T-RELEASE)} \\
\frac{\Gamma \vdash x : \alpha \quad \Gamma(\alpha) \neq \mathbb{F}\perp \\
\alpha \notin \{\beta \mid f' \in \text{dom}(\Gamma) \text{ and } \Gamma(f') = (\mathbb{O}', \beta, \mathbb{x}, \mathbb{R}') \text{ and } \mathbb{R}' \neq \emptyset\}}{\Gamma \vdash_{\mathbb{S}} \mathbf{release } x : \alpha \checkmark \triangleright \Gamma[\alpha \mapsto \mathbb{O}\perp]} \quad \text{(T-RELEASE-BOT)} \\
\frac{\Gamma \vdash x : \alpha \quad \Gamma(\alpha) = \mathbb{F}\perp}{\Gamma \vdash_{\mathbb{S}} \mathbf{release } x : \mathbb{O} \triangleright \Gamma} \\
\\
\text{(T-IF)} \\
\frac{\Gamma \vdash e : se \quad \Gamma \vdash_{\mathbb{S}} s_1 : \mathbb{C}_1 \quad \Gamma \vdash_{\mathbb{S}} s_2 : \mathbb{C}_2}{\Gamma \vdash_{\mathbb{S}} \mathbf{if } e \{s_1\} \mathbf{else } \{s_2\} : (se)\{\mathbb{C}_1\} + (\neg se)\{\mathbb{C}_2\}} \\
\\
\text{(T-IF-ND)} \\
\frac{\Gamma \vdash e : _ \quad \Gamma \vdash_{\mathbb{S}} s_1 : \mathbb{C}_1 \quad \Gamma \vdash_{\mathbb{S}} s_2 : \mathbb{C}_2}{\Gamma \vdash_{\mathbb{S}} \mathbf{if } e \{s_1\} \mathbf{else } \{s_2\} : \mathbb{C}_1 + \mathbb{C}_2} \\
\\
\text{(T-SEQ)} \\
\frac{\Gamma \vdash_{\mathbb{S}} s_1 : \mathcal{C}[\mathfrak{a}_1 \triangleright \Gamma_1] \cdots [\mathfrak{a}_n \triangleright \Gamma_n] \quad (\Gamma_i \vdash_{\mathbb{S}} s_2 : \mathbb{C}'_i)^{i \in 1..n}}{\Gamma \vdash_{\mathbb{S}} s_1; s_2 : \mathcal{C}[\mathfrak{a}_1 \mathfrak{s} \mathbb{C}'_1] \cdots [\mathfrak{a}_n \mathfrak{s} \mathbb{C}'_n]} \\
\\
\text{(T-RETURN)} \\
\frac{\Gamma \vdash e : \mathbb{O} \quad \Gamma \vdash \mathbf{destiny} : \mathbb{O}' \\
\mathbb{O} \notin \mathbb{S}}{\Gamma \vdash_{\mathbb{S}} \mathbf{return } e : \mathbb{O} \triangleright \Gamma[\mathbb{O}' \mapsto \Gamma(\mathbb{O})]}
\end{array}$$

Figure 2.5: Type rules for expressions and statements.

Let $\mathbb{F}\mathbb{t}$ be the state of the caller (which is recorded in $\Gamma(f)$). We use the operation $\mathbb{t} \sqcap \Gamma(\alpha) \Downarrow$ to this purpose, which means that the returned machine gets the same state of the carrier α if no method is releasing α , otherwise its state is either $\emptyset\partial$ or $\emptyset\alpha \downarrow$, according to the state of α was $\mathbb{F}\top$ or $\mathbb{F}\alpha$.

Rule (T-RELEASE) models the removal of a VM name α . The premise in the second line verifies that the disposal do not address machines that are executing methods, as discussed in Restriction 3 of Section 2.4.

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

The type system of `vml` is completed with the rules for method declarations and programs, given in Figure 2.6.

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 2.6). We also observe that the returned value \circ may be either $_$ or a fresh VM name ($\circ \notin \{\alpha\} \cup \bar{\beta}$) as discussed in Restriction 4 of Section 2.4. The constraints in the third line of the premises of (T-METHOD) implement Restriction 1 of Section 2.4. We also observe that $(\Gamma_i(\gamma) = \Gamma_j(\gamma))^{i,j \in 1..n, \gamma \in \mathbb{S} \cup \text{fv}(\circ)}$ guarantees that every branch of the behavioral type creates a new VM name and, by rule (T-RETURN), the state of the chosen VM name must be always the same.

(T-METHOD)

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

(T-PROGRAM)

$$\frac{\Gamma \vdash \bar{M} : \bar{\mathbb{C}} \quad \Gamma[\mathbf{this} \mapsto \text{start}] \vdash_{\text{start}} s : \mathcal{C}[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n]}{\Gamma \vdash \bar{M} \{ \overline{\mathbf{F}} x ; s \} : \bar{\mathbb{C}}, \mathcal{C}[\mathbf{a}_1 \triangleright \emptyset] \cdots [\mathbf{a}_n \triangleright \emptyset]}$$

Figure 2.6: Behavioral typing rules of method and programs.

We display behavioral types examples by using codes from Sections 2.2 and 2.5. Actually, the following types do not abstract a lot from codes because the programs of the previous sections have been designed for highlighting the issues of our technique. The following listing shows the behavioral type of `double_release` and the step by step rule application for its calculation.

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

```
double_release  $\alpha(\beta, \gamma)$  { // T-Method
   $\beta^\checkmark$ ; // T-Release
   $\gamma^\checkmark$ ; // T-Release, T-Seq
  0 // T-Return, T-Seq
} -, { $\beta, \gamma$ }
```

The behavioral types of (`fact` and `costly_fact` examples from Section 2.5 are the following (we remove tailing 0 when irrelevant):

```
fact  $\alpha(n)$  {
  (n==0){ 0 }
  +(n>0){  $\nu x$ :fact  $\alpha(n-1)$  ;  $x^\checkmark$  }
} -, { }
```

```
costly_fact  $\alpha(n)$  {
  (n==0){ 0 }
  +(n>0){  $\nu\beta$ ;
           $\nu x$ :costly_fact  $\beta(n-1)$ ;
           $x^\checkmark$  ;  $\beta^\checkmark$  }
} -, { }
```

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

2.6 The analysis of vml behavioral types

The types returned by the system in Section 2.5 are used to compute the resource cost of a vml program. This computation is performed by a solver of *cost equations*. These cost equations are terms

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

where m is a (cost) function symbol, exp is an expression that may contain (cost) function symbols applications, and se is a size expression whose variables are contained in \bar{x} (the syntax of exp and se is given in full detail in [28]).

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

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

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

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

2.6.1 Arguments' identities

Consider the following code and the corresponding behavioral types

<pre> Int simple_release(VM x) { release x ; return 0 ; } Int m(VM x, VM y) { Fut<Int> f; f = this!simple_release(x); release y; f.get; return 0; } </pre>	<pre> simple_release $\alpha(\beta)${ β^\checkmark } -, {β} m $\alpha(\beta, \gamma)${ νf : simple_release $\alpha(\beta)$; γ^\checkmark ; f^\checkmark } -, {β, γ} </pre>
---	--

We notice that, in the type of m , there is not enough information to determine whether γ^\checkmark will have a cost equal to -1 or 0. In fact, in the rule (T-METHOD), we assumed that the arguments were pairwise different. However, this is not the case for invocations. For instance, if m is invoked with two arguments that are equal - $\beta = \gamma$ - then γ is going to be released by the invocation $\text{free}(\beta)$ and therefore it counts 0. We solve this problem of arguments' identity in the analysis of behavioral types, in particular in the translation of method types. Namely, the above method m is translated in four cost functions: $m_{\text{peak}}^{\{1\},\{2\}}(x, y)$ and $m_{\text{net}}^{\{1\},\{2\}}(x, y)$, which correspond to the invocations where $x \neq y$, and $m_{\text{peak}}^{\{1,2\}}(x)$ and $m_{\text{net}}^{\{1,2\}}(x)$, which correspond to the

invocations where $x = y$. (The equivalence relation in the superscript never mention **this**, which is also an argument, because, in this case **this** cannot be identified with the other arguments, see below.) Then the translation of an invocation to a method m redirects the invocation to m^Ξ , where Ξ is the equivalence relation expressing the identity of the arguments.

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

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

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

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

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

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

Without loosing in generality, we will always assume that the canonical representative of a set containing 0 is always 0. This index represents the **this** object and we remind that, by Simplification 3 in Section 2.4, such an object cannot be released. This is the reason why, in the foregoing discussion about the method m , we have not mentioned **this**. Additionally, in order to simplify the translation of method invocations, we also assume that the argument **this** is always different from other arguments (the general case just requires more details).

2.6.2 (Re)computing argument's states

In the rules of Figure 2.5, in order to enforce the restrictions in Section 2.4, we have already computed the state of every machine. In this section we

recompute them for a different reason: obtaining a (more) precise cost analysis. Of course one might record the computation of VM states in behavioral types. However, this solution has the drawback that behavioral types become unintelligible because they carry information that is needed at a later stage by the analyzer.

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

- let Ψ be a vm-translation environment. Then

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

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

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

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

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

$$\mathfrak{t} \otimes^\alpha \mathfrak{t}' \stackrel{def}{=} \begin{cases} \perp & \text{if } \mathfrak{t} = \perp \text{ or } \mathfrak{t}' = \perp \\ \alpha & \text{if } \mathfrak{t} \text{ and } \mathfrak{t}' \text{ are variables} \\ \alpha \downarrow & \text{otherwise} \end{cases}$$

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

Then

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

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

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

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

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

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

2.6.3 The translation function

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

1. Δ is the *equivalence relation on formal parameters* identifying those that are equal. We assume that $\Delta(x)$ returns the unique representative of the equivalence class of x . For simplicity we also let $\Delta(x) = x$ for every x that belongs to the local variables. Therefore we can use Δ also as a substitution operation.

2. Ψ is the *translation environment* which stores temporary information about futures that are active (unsynchronized) and about the state of VM names;
3. α is the name of the virtual machine of the current behavioral type;
4. \bar{e} is the sequence of (over-approximated) costs of the current execution branch;
5. the behavioral type being translated; it may be either \mathbf{a} , \mathbf{c} or $\overline{\mathbf{C}}$.

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

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

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

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

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

$$\text{translate}[\Delta, \Psi, \alpha](\bar{e}; e)(\mathbf{a}) =$$

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

In the definition of `translate` we always highlight the last expression in the sequence of costs of the current execution branch (the fourth input).

This is because the cost of the parsed atom applies to it, except for the case of f^\vee . In this last case, let $\bar{e}; e$ be the expression. Since the atom expresses the synchronization of f , $\bar{e}; e$ will have occurrences of f . In this case, the function **translate** has to compute two values: the maximum number of resources used by (the method corresponding to) f during its execution – the *peak cost* used in the substitution σ – and the resources used upon the termination of (the method corresponding to) f – the *net cost* used in the substitution σ' . In particular, this last value has to be decreased by the number of resources released by the method. This is the purpose of the addend $\sum_{\gamma \in \Delta(\mathbf{R}), \Theta(\gamma) = \mathbf{F}\dagger, \mathbf{F} \neq \emptyset} \mathbf{CREL}(\dagger)$ that removes machines that are going to be removed by parallel methods (the constraint $\mathbf{F} \neq \emptyset$) because the other ones have been already counted both in the peak cost and in the net cost. We observe that the instances of the method \mathbf{m}_{peak} and \mathbf{m}_{net} that are invoked are those corresponding to the equivalence relation of the tuple $(\beta, \bar{\beta})$.

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

translate $(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbf{c}) =$

$$\left\{ \begin{array}{ll} \{(se')\{\bar{e}'\}\} & \text{when } \mathbf{c} = \mathbf{a} \triangleright \emptyset \text{ and } \mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbf{a}) = (\Psi', (se')\{\bar{e}'\}) \\ C'' & \text{when } \mathbf{c} = \mathbf{a} \ddagger c' \text{ and } \mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbf{a}) = (\Psi', \{(se')\{\bar{e}'\}\}) \\ & \text{and } \text{dom}(\Psi') \setminus \text{dom}(\Psi) = \mathbf{S} \text{ and } \mathbf{translate}(\Delta \cup \{\mathbf{S}\}, \Psi', \alpha, (se')\{\bar{e}'\}, c') = (\Psi'', C'') \\ C' \cup C'' & \text{when } \mathbf{c} = \mathbf{c}_1 + \mathbf{c}_2 \text{ and } \mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbf{c}_1) = (\Psi', C') \\ & \text{and } \mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, \mathbf{c}_2) = (\Psi'', C'') \\ C' & \text{when } \mathbf{c} = (se')\{c'\} \text{ and } \mathbf{translate}(\Delta, \Psi, \alpha, (se \wedge se')\{\bar{e}\}, c') = (\Psi', C') \\ C' & \text{when } \mathbf{c} = (e')\{c'\} \text{ and } e' \text{ contains } _ \text{ and } \mathbf{translate}(\Delta, \Psi, \alpha, (se)\{\bar{e}\}, c') = (\Psi', C') \end{array} \right.$$

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

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

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

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

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

Then

$\text{translate}(\Xi, m \alpha_1(\bar{x}, \alpha_2, \dots, \alpha_k) \{c\} : \mathbb{D}, \mathbb{R}) =$

$$\left\{ \begin{array}{ll} \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = 0 & [\alpha_1 = \perp] \\ \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{1,1} & [se_1 \wedge \alpha_1 \neq \perp] \\ \vdots & \\ \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_1,1} & [se_1 \wedge \alpha_1 \neq \perp] \\ \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{1,2} & [se_2 \wedge \alpha_1 \neq \perp] \\ \vdots & \\ \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_n,n} & [se_n \wedge \alpha_1 \neq \perp] \\ \mathbb{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = 0 & [\alpha_1 = \perp] \\ \mathbb{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = \mathbb{m}_{\text{peak}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \dots, \alpha_n]) & [\alpha_1 = \partial] \\ \mathbb{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_1,1} & [se_1 \wedge \alpha_1 = \top] \\ \vdots & \\ \mathbb{m}_{\text{net}}^{\Xi}(\bar{x}, \Xi[\alpha_1, \alpha_2, \dots, \alpha_k]) = e_{h_n,n} & [se_n \wedge \alpha_1 = \top] \end{array} \right.$$

Let $(\mathbb{C}_1 \dots \mathbb{C}_n, c)$ be a behavioral type program and let

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

then

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

As an example, we show the output of `translate` when applied to the behavioral type of `double_release` computed in Section 2.5. Since `double_release` has two arguments, we generate two sets of equations, as discussed above. In order to ease the reading, we omit the equivalence classes of arguments that label function names: the reader can grasp them from the number of arguments. For the same reason, we represent a partition $\{\{1\}, \{2\}, \{3\}\}$ corresponding to VM names α_1, α_2 and α_3 by $[\alpha_1, \alpha_2, \alpha_3]$ and $\{\{1\}, \{2, 3\}\}$ by $[\alpha_1, \alpha_2]$ (we write the canonical representatives). For simplicity we do not add the partition to the name of the method.

$$\text{translate}([\alpha_1, \alpha_2, \alpha_3], \text{double_release } \alpha_1(\alpha_2, \alpha_3) \{c\} : -, \{\alpha_2, \alpha_3\}) =$$

$$\left\{ \begin{array}{ll} \text{double_release_peak}(\alpha_1, \alpha_2, \alpha_3) = 0 & [\alpha_1 = \perp] \\ \text{double_release_peak}(\alpha_1, \alpha_2, \alpha_3) = 0 & [\alpha_1 \neq \perp] \\ \text{double_release_peak}(\alpha_1, \alpha_2, \alpha_3) = \text{CREL}(\alpha_2) & [\alpha_1 \neq \perp] \\ \text{double_release_peak}(\alpha_1, \alpha_2, \alpha_3) = \text{CREL}(\alpha_2) + \text{CREL}(\alpha_3) & [\alpha_1 \neq \perp] \\ \\ \text{double_release_net}(\alpha_1, \alpha_2, \alpha_3) = 0 & [\alpha_1 = \perp] \\ \text{double_release_net}(\alpha_1, \alpha_2, \alpha_3) = \text{double_release_peak}(\alpha_1, \alpha_2, \alpha_3) & [\alpha_1 = \emptyset] \\ \text{double_release_net}(\alpha_1, \alpha_2, \alpha_3) = \text{CREL}(\alpha_2) + \text{CREL}(\alpha_3) & [\alpha_1 = \top] \end{array} \right.$$

$$\text{translate}([\alpha_1, \alpha_2], \text{double_release } \alpha_1(\alpha_2) \{c\} : -, \{\alpha_2\}) =$$

$$\left\{ \begin{array}{ll} \text{double_release_peak}(\alpha_1, \alpha_2) = 0 & [\alpha_1 = \perp] \\ \text{double_release_peak}(\alpha_1, \alpha_2) = 0 & [\alpha_1 \neq \perp] \\ \text{double_release_peak}(\alpha_1, \alpha_2) = \text{CREL}(\alpha_2) & [\alpha_1 \neq \perp] \\ \text{double_release_peak}(\alpha_1, \alpha_2) = \text{CREL}(\alpha_2) + \text{CREL}(\perp) & [\alpha_1 \neq \perp] \\ \\ \text{double_release_net}(\alpha_1, \alpha_2) = 0 & [\alpha_1 = \perp] \\ \text{double_release_net}(\alpha_1, \alpha_2) = \text{double_release_peak}(\alpha_1, \alpha_2) & [\alpha_1 = \emptyset] \\ \text{double_release_net}(\alpha_1, \alpha_2) = \text{CREL}(\alpha_2) + \text{CREL}(\perp) & [\alpha_1 = \top] \end{array} \right.$$

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

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

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

```
user1  $\alpha$ ( ) {
   $\nu\beta \ ; \ \nu\gamma \ ; \ \nu f : \text{double\_release } \alpha(\beta, \gamma) \ ; \ f^\checkmark$ 
} - , { }
```

```
user2  $\alpha$ ( ) {
   $\nu\beta \ ; \ \nu f : \text{double\_release } \alpha(\beta, \beta) \ ; \ f^\checkmark$ 
} - , { }
```

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

$$\text{translate}([\alpha_1], \text{user1 } \alpha_1() \{ \mathbb{C}_{\text{user1}} \} : -, \{ \}) =$$

$$\left\{ \begin{array}{ll} \text{user1}_{\text{peak}}(\alpha_1) = 0 & [\alpha_1 = \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = 0 & [\alpha_1 \neq \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) & [\alpha_1 \neq \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{CNEW}(\alpha_1) & [\alpha_1 \neq \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{CNEW}(\alpha_1) + \text{double_release}_{\text{peak}}(\alpha_1, \top, \top) & [\alpha_1 \neq \perp] \\ \text{user1}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{CNEW}(\alpha_1) + \text{double_release}_{\text{net}}(\alpha_1, \top, \top) & [\alpha_1 \neq \perp] \\ \\ \text{user1}_{\text{net}}(\alpha_1) = 0 & [\alpha_1 = \perp] \\ \text{user1}_{\text{net}}(\alpha_1) = \text{user1}_{\text{peak}}(\alpha_1) & [\alpha_1 = \emptyset] \\ \text{user1}_{\text{net}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{CNEW}(\alpha_1) + \text{double_release}_{\text{net}}(\alpha_1, \top, \top) & [\alpha_1 = \top] \end{array} \right.$$

$$\text{translate}([\alpha_1], \text{user2 } \alpha_1() \{ \mathbb{C}_{\text{user2}} \} : -, \{ \}) =$$

$$\left\{ \begin{array}{ll} \text{user2}_{\text{peak}}(\alpha_1) = 0 & [\alpha_1 = \perp] \\ \text{user2}_{\text{peak}}(\alpha_1) = 0 & [\alpha_1 \neq \perp] \\ \text{user2}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) & [\alpha_1 \neq \perp] \\ \text{user2}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{double_release}_{\text{peak}}(\alpha_1, \top) & [\alpha_1 \neq \perp] \\ \text{user2}_{\text{peak}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{double_release}_{\text{net}}(\alpha_1, \top) & [\alpha_1 \neq \perp] \\ \\ \text{user2}_{\text{net}}(\alpha_1) = 0 & [\alpha_1 = \perp] \\ \text{user2}_{\text{net}}(\alpha_1) = \text{user2}_{\text{peak}}(\alpha_1) & [\alpha_1 = \emptyset] \\ \text{user2}_{\text{net}}(\alpha_1) = \text{CNEW}(\alpha_1) + \text{double_release}_{\text{net}}(\alpha_1, \top) & [\alpha_1 = \top] \end{array} \right.$$

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

2.7 Correctness

The correctness of our technique is expressed by the following theorem that relates the `vml` language, its behavioral type system and the *cost* of a `vml` program.

Theorem 7.1 (Correctness). *Let $\overline{M} \{ \overline{Fz} ; s' \}$ be a well-typed program and let $\overline{\mathbb{C}}, \mathbb{c}$ be its behavioral type and cn be its initial configuration. Let also n be a solution of the function $\text{translate}(\overline{\mathbb{C}}, \mathbb{c})$. Then, for every $cn \longrightarrow^* cn'$, $\text{alive}(cn') \leq n$.*

The proof of this theorem, and therefore of the correctness of our approach, consists of two parts. The first part addresses the correctness of the type system in Section 2.5, which is usually expressed by a subject reduction

theorem. This theorem states that if a configuration cn of the operational semantics is well typed and $cn \rightarrow cn'$ then cn' is well-typed as well. It is worth observing that we cannot hope to demonstrate a statement guaranteeing type-preservation because our types are “behavioral” and change during the evolution of the systems. However, it is critical for the correctness of the cost analysis that there exists a relation between the type of cn and the type of cn' .

Therefore, a subject reduction for the type system of Section 2.5 requires

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

2.7.1 The extension of the typing to configurations

When typing configurations there are two kinds of vm names: *static-time VM names* and *runtime VM names*. The former names are ranged over by α, β, \dots ; the latter names, which also include *err*, are ranged over by o, o', \dots . With an abuse of notation, in the following, α and \circ will also range over runtime VM names o and *err*. This abuse will let us to reuse the typing rules in Figure 2.5 (with one exception, see below). The syntax of runtime types is:

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

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

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

- \mathbf{S} is a set of runtime VM names.

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

The runtime type system consists of rules in Figure 2.5 plus the rules in Figure 2.7. The reader may notice that the first rule (BT-RETURN) replaces the definition of the judgment $\Gamma \vdash_{\mathbf{S}} s : \mathbb{c} \triangleright \Gamma'$ for the `return` statement in Figure 2.5. This modification, together with the extensions of atoms, is used to refine the (static-time) typing rules in order to deal with a naming issue in the subject reduction theorem. In particular, when a method returns a freshly created virtual machine, we associate to it a fresh name in (T-INVOKE) that we can use later in (T-GET) and in the rest of the typing of the (static) program. The problem is that, when we type the method invocation (i.e. when we collect the information about the method and its return value in order to type future usage of this value), the identity of the returned virtual machine is still unknown because the virtual machine has not yet been created. In order to solve this problem, we extend atoms with the term $[\circ]$, which retains the returned value and we use the new rule (BT-RETURN) to type the `return` statements.

The other rule of Figure 2.7 that we comment is (BT-VM). This rule allows one to type a VM term containing a number of processes p . To this aim, one has to type every p by means of the judgment $\Gamma[\mathbf{this} \mapsto o] \vdash_{\mathbf{S}} p : (f, \mathbb{c}, \mathbf{S})$; namely the index \mathbf{S} and the third argument of the process type must be the same. In fact, this argument records the actual VM names of the corresponding method invocation – see rule (R-BIND-MTD) in Figure 2.8. In turn, this set is used in (BT-PROCESS) to type the statements of the processes.

Lemma 7.2 (Substitution Lemma). *Let $\Gamma \vdash cn : \mathbb{k}$ and let ι be an injective renaming of names, VM names, and future names. Let $\iota(\Gamma) \stackrel{\text{def}}{=} [\iota(x) \mapsto \iota(\Gamma(x))]^{x \in \text{dom}(\Gamma)}$.*

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

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

2.7.2 Runtime types evolution

As usual with behavioral types, our types change while the programs execute. These changes can be defined and, more importantly, can be related to the

$$\begin{array}{c}
\text{(BT-RETURN)} \\
\frac{\Gamma \vdash e : \Phi \quad \Gamma \vdash \text{destiny} : \Phi' \quad \Phi \notin \mathbf{S}}{\Gamma \vdash_{\mathbf{S}} \text{return } e : [\Phi] \triangleright \Gamma[\Phi' \mapsto \Gamma(\Phi)]}
\end{array}
\qquad
\begin{array}{c}
\text{BT-FUT-RUNNING} \\
\frac{\Gamma(f) = (\Phi, o, \mathbf{Ft}, \mathbf{R})}{\Gamma \vdash \text{fut}(f, \perp) : f(\perp)}
\end{array}$$

$$\begin{array}{c}
\text{BT-FUT-COMPUTED} \\
\frac{\Gamma \vdash v : \Phi \quad \Gamma \vdash f : \Phi}{\Gamma \vdash \text{fut}(f, v) : f(\Phi)}
\end{array}
\qquad
\begin{array}{c}
\text{BT-FUT-UNSYNC} \\
\frac{\Gamma(f) = (\Phi, o, \mathbf{Ft}, \mathbf{R}) \quad v \neq \perp \quad \Gamma \vdash v : \Phi' \quad \alpha \in \mathbf{R}[\Phi'/\Phi] \Rightarrow \Gamma(\alpha) = \emptyset \perp}{\Gamma \vdash \text{fut}(f, v) : f(\Phi')}
\end{array}
\qquad
\begin{array}{c}
\text{BT-EMPTY} \\
\Gamma \vdash \epsilon : \mathbf{0}
\end{array}$$

$$\begin{array}{c}
\text{BT-INVOKE} \\
\frac{\Gamma \vdash \text{mo}(\bar{v}) : \Phi, \mathbf{R} \quad \Gamma(f) = (\Phi, o, \mathbf{Ft}, \mathbf{R}) \quad \Gamma \vdash \bar{v} : \bar{\mathbf{r}}}{\Gamma \vdash \text{invoc}(o, f, \mathbf{m}, \bar{v}) : (f, o \mathbf{m}(\bar{\mathbf{r}}))}
\end{array}$$

$$\begin{array}{c}
\text{BT-VM} \\
\frac{\Gamma \vdash o : \mathbf{Fa} \quad (\Gamma \vdash_{\mathbf{S}_i} p_i : \mathbf{f}_i \quad \mathbf{f}_i = (f_i, \mathbf{c}_i, \mathbf{S}_i))^{i \in 1..n}}{\Gamma \vdash \text{vm}(o, a, p_1; \{p_2 \dots p_n\}) : \text{vm}(o, a \mid \mathbf{f}_1; \{\mathbf{f}_2, \dots, \mathbf{f}_n\})}
\end{array}
\qquad
\begin{array}{c}
\text{BT-PARALLEL} \\
\frac{\Gamma \vdash \text{cn} : \mathbf{k} \quad \Gamma \vdash \text{cn}' : \mathbf{k}'}{\Gamma \vdash \text{cn} \text{ cn}' : \mathbf{k} \parallel \mathbf{k}'}
\end{array}$$

$$\begin{array}{c}
\text{BT-PROCESS} \\
\frac{l = [(x_i \mapsto v_i)^{i \in I}, \text{destiny} \mapsto f] \quad (\Gamma \vdash v_i : \mathbf{x}_i)^{i \in I} \quad \Gamma(f) = (\Phi, o, \mathbf{Ft}, \mathbf{R}) \quad \Gamma[(x_i \mapsto \mathbf{x}_i)^{i \in I}][\text{destiny} \mapsto \Phi] \vdash_{\mathbf{S}} s : \mathcal{C}[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n] \quad \left(\Gamma_i(\gamma) = \Gamma_j(\gamma) \right)_{i,j \in 1..n, \gamma \in \mathbf{S}} \quad \mathbf{R} = (\mathbf{S} \cup \text{fv}(\Phi)) \cap \{\gamma \mid \Gamma_1(\gamma) = \mathbf{F}\perp\}}{\Gamma \vdash_{\mathbf{S}} \{l \mid s\} : (f, \mathcal{C}[\mathbf{a}_1 \triangleright \emptyset] \cdots [\mathbf{a}_n \triangleright \emptyset], \mathbf{S})}
\end{array}$$

Figure 2.7: Runtime Typing rules

cost analysis. The modifications of types, noted \rightsquigarrow , is defined as a reduction relation in Figure 2.8. The rules are mostly an adaptation of the operational semantics of `vm1` at the type level. Therefore, to ease the understanding of the relation \rightsquigarrow , we use similar names for similar rules. Nevertheless, there are few differences from the operational semantics of `vm1`: *i*) rule (ASSIGN) has no corresponding rule in the definition of \rightsquigarrow (assignments are managed at the type system level); *ii*) rule (ACTIVATE) is replaced by the rule (R-GC) (there are no scheduling policies in our runtime types, so the only effect of the rule (ACTIVATE) at the type level is the deletion of a finished process); *iii*) similarly, rule (ACTIVATE-GET) has no corresponding rule in the definition of \rightsquigarrow ; and *iv*) because at runtime all objects identities are known, the rules (READ-FUT), (ASYNC-CALL), (BIND-MTD) and (NEW-VM) have been adjusted in order to replace the names generated a static time with the objects' actual identities.

2.7.3 The subject reduction theorem

Based on the previous construction, we present the subject reduction theorem. Lemma 7.3 relates the static time type system to the runtime type system: this corresponds to the first step of the subject reduction property.

Lemma 7.3. *Let $P = \overline{M} \{ \overline{F} x ; s \}$ be a `vm1` program, and let $\Gamma \vdash P : \overline{\mathbb{C}}, \mathbb{c}$. Then*

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

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

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

Theorem 7.4 (Subject Reduction). *Let $\Gamma \vdash cn : \mathbb{k}$. If $cn \rightarrow cn'$ then there exists Γ' such that $\Gamma' \vdash cn' : \mathbb{k}'$ and $\mathbb{k} \rightsquigarrow^* \mathbb{k}'$.*

Proof. As usual for Subject Reduction theorems we proceed by case on the reduction rule used in $cn \rightarrow cn'$. Since our types are behavioral we cannot assume that the types of a runtime programs state and its reduction are equivalent, instead we prove that this types $(\mathbb{k}, \mathbb{k}')$ are related according to the operational semantics for runtime types, see Figure 2.8. We present the proof of some cases, the rest are either straightforward or similar to those below. Whenever the construction of ι is omitted it is considered as

the identity function. Finally, if not stated otherwise, we implicitly use the notation used in the considered reduction rule:

- Case (ASSIGN).

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

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

- Case (READ-FUT).

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

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

Let $\iota = [\circ''/\circ]$: we have that $\iota(\Gamma) = \Gamma$. Let $\Gamma' = \Gamma[f \mapsto \circ'']$, by lemma 7.2 we have that $\Gamma'[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \circ'][\beta \mapsto \emptyset \perp]^{\beta \in \iota(\mathbf{R})}[\beta' \mapsto \emptyset \mathbb{t}']^{\beta' \in \mathbf{R}''} \vdash s : \iota(c)$ with $\mathbf{R}'' = fv(\circ'') \setminus \iota(\mathbf{R})$. To prove the theorem in this case, we need to show that s can be typed with the environment $\Gamma'[(x_i \mapsto \mathbb{x}_i)^{i \in I}][\text{destiny} \mapsto \circ']$. First, because of the side condition in rule (BT-FUT-UNSYNC), we have that $\Gamma'(\beta) = \emptyset \perp$ for all $\beta \in \iota(\mathbf{R})$. We thus have

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

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

- either \mathbf{R}'' is empty, in which case that substitution is the identity;

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

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

- Case (ASYNC-CALL).

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

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

Let $\iota = [f/f'']$ and lets choose Γ' such that $\Gamma'' = \Gamma'[(x_i \mapsto \mathfrak{x}_i)^{i \in I}][\text{destiny} \mapsto \phi']$, by lemma 7.2 we have that $\iota(\Gamma'[(x_i \mapsto \mathfrak{x}_i)^{i \in I}][\text{destiny} \mapsto \phi']) \vdash \text{vm}(o, \top, \{l \mid x = e!m(\bar{e}); s\}, q) \rightarrow \text{vm}(o, \top, \{l \mid x = f; s\}, q) \text{ invoc}(o', f, m, \bar{v}) : \mathbb{k}'$ with $\mathbb{k}' = \text{vm}(o, \top \mid (f', 0 \ ; \ \mathfrak{c}', \mathbf{S}'), \bar{\mathfrak{f}}) \parallel (f'', m \ o'(\bar{\mathfrak{x}}))$, where $\mathfrak{c}' = \iota(\mathfrak{c})$. Behavioral types \mathbb{k} and \mathbb{k}' are related by (R-ASYNC-CALL).

- Case (RELEASE-VM).

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

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

- Case (RETURN).

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

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

- Case (VM-ERR-RETURN).

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

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

□

2.7.4 Correctness of the cost computation

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

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

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

Definition 7.5. *A runtime type \mathbb{k} is well-formed iff all the names (object, future, etc) it uses are declared and if it contains the object: $vm(start, a \mid (f_{start}, \mathbb{C}, \mathbb{S}); \mathbb{F})$. Given a well-formed runtime type \mathbb{k} , we write: $\mathbf{alive}(\mathbb{k})$ the number of alive VM in \mathbb{k} (i.e., the VM whose state is \top in \mathbb{k}); $\mathbf{mapping}(\mathbb{k})$ the mapping giving the state of VM in \mathbb{k} ; and $\Delta_{\mathbb{k}}$ the set $\{ \{o\} \mid o \text{ is an object of } \mathbb{k} \}$ (i.e., $\Delta_{\mathbb{k}}$ is the identity equivalence relation generated from \mathbb{k}).*

We extend the function `translate` as presented in Figure 2.9 and define:

Definition 7.6. *The cost equations corresponding to a well-formed runtime type \mathbb{k} , written $\mathbf{translate}(\mathbb{k})$ is defined as follows:*

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

Note in Figure 2.9 that we also need to extend the definition of `translate` on atoms: we need to re-define the cost of a `get` operation as we now have not only the actual futures in the runtime type but also the ones coming from the static typing of the program (the ones that belong to Ψ). Moreover the initial definition of `translate` for the `get` included an extra adjustment derived from the fact that the state of some virtual machine were unknown, which is not the case at runtime.

Cost Solution A *cost solution* Σ is a mapping from cost functions (like $m_{\text{net}}^{\Xi}(\bar{v})$ or f_{peak} , where f is a future name) to cost expressions. Basically, this cost solution gives cost expressions of the peak and net cost of all methods in the program, as well as, the peak and net cost of all processes in the runtime type. Following [28, Definition 6], we relate cost solutions and runtime types by assuming the existence of a *cost validation predicate* \models between cost solutions and cost equations as generated by the `translate` function in Section 2.6:

Definition 7.7. *Let C be a set of cost equations $\bigcup_{i=1}^n f_i(\bar{x}_i) = e_i[se_i]$; let Σ be a function that maps all $f_i(\bar{x}_i)$ to cost expressions. Σ is a cost solution of C , noted $\Sigma \models C$, if for all $1 \leq i \leq n$ and all substitutions $[\bar{v}/\bar{x}_i]$ such that $se_i[\bar{v}/\bar{x}_i]$ is true, we have that*

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

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

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

Lemma 7.8. *Let $\overline{M} \{ \overline{Fz} ; s' \}$ be a well-typed program, let $\overline{\mathbb{C}}, \mathbb{c}$ be its behavioral type, cn be its initial configuration and \mathbb{k} the runtime type of cn . The for all cost solution Σ with $\Sigma \models \text{translate}(\overline{\mathbb{C}}, \mathbb{c})$, there exists a cost solution Σ' such that $\Sigma' \models \text{translate}(\mathbb{k})$ and $\Sigma(\text{main}()) = \Sigma'(f_{start_peak}) + \text{alive}(\mathbb{k})$.*

Proof. Let define $\Sigma' = \Sigma[f_{start_peak} \mapsto \Sigma(\text{main}()) - 1]$. By construction (see the construction of the equation page 37 and Figure 2.9), we have that

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

$$\text{translate}(\mathbb{k}) = \left\{ \begin{array}{l} f_{start_peak} = e_{1,1} \quad [se_1] \\ \vdots \\ f_{start_peak} = e_{h_1,1} \quad [se_1] \\ f_{start_peak} = e_{1,2} \quad [se_2] \\ \vdots \\ f_{start_peak} = e_{h_m,m} [se_m] \end{array} \right.$$

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

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

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

Lemma 7.9. *Let $\Gamma \vdash cn : \mathbb{k}$ and assume there is a cost solution Σ such that $\Sigma \models \text{translate}(\mathbb{k})$. Then $\text{alive}(cn) \leq \text{alive}(\mathbb{k}) + \Sigma(f_{\text{start_peak}})$ (recall that $f_{\text{start_peak}}$ is the peak cost of the main process type of \mathbb{k}).*

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

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

Lemma 7.10. *Suppose given two well-formed runtime types \mathbb{k} and \mathbb{k}' such that $\mathbb{k} \rightsquigarrow \mathbb{k}'$. Then, for all cost solution Σ with $\Sigma \models \text{translate}(\mathbb{k})$, there exists Σ' such that: i) $\Sigma' \models \text{translate}(\mathbb{k}')$; and ii) $\text{alive}(\mathbb{k}') + \Sigma'(f_{\text{start_peak}}) \leq \text{alive}(\mathbb{k}) + \Sigma(f_{\text{start_peak}})$.*

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

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

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

- Case (R-BIND-MTD). Given that c is the behavior of $m \ o(\bar{v})$ and that $\text{translate}(\mathbb{k}_1 \parallel \mathbb{k}_2) = \text{translate}(\mathbb{k}_1) \cup \text{translate}(\mathbb{k}_2)$, $\text{translate}(\mathbb{k})$ and $\text{translate}(\mathbb{k}')$ produce the same set of equations up to the substitution of formal parameters. By definition of Σ we have the result for $\Sigma' = \Sigma$.
- Case (R-NEW-VM). Let f_i be all possible branches of f . We define

$$\Sigma' = \Sigma[f_{i,\text{peak}} \mapsto \Sigma(f_{i,\text{peak}}) - 1][f_{i,\text{net}} \mapsto \Sigma(f_{i,\text{net}}) - 1]$$

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

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

□

2.7.5 Final demonstration

Finally, we can combine all the previous lemma to demonstrate the correctness of Theorem 7.1.

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

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

- At step 0, we translate the program into its runtime equivalent cn . By Lemma 7.8, there exists Σ' such that $\Sigma' \models \mathbf{translate}(\mathbb{k})$ and that we have that $\Sigma'(f_{start_peak}) + \mathbf{alive}(\mathbb{k}) = n$. Moreover, by Lemma 7.9 we have that $\mathbf{alive}(cn) \leq \mathbf{alive}(\mathbb{k}) + \Sigma'(f_{start_peak}) = n$.
- for the inductive case, let Γ , a runtime configuration cn' , a runtime type \mathbb{k} and a cost solution Σ' such that $\Gamma \vdash cn : \mathbb{k}$, $\Sigma' \models \mathbb{k}$ and $\Sigma'(f_{start_peak}) + \mathbf{alive}(\mathbb{k}) \leq n$ hold. Then by Theorem 7.4, if $cn \rightarrow cn'$, there exists Γ' and a runtime type \mathbb{k}' such that $\mathbb{k} \rightsquigarrow^* \mathbb{k}'$ and $\Gamma' \vdash cn' : \mathbb{k}'$. By Lemma 7.10, there exist a cost solution Σ'' such that $\Sigma'' \models \mathbb{k}'$ and $\mathbf{alive}(\mathbb{k}') + \Sigma''(f_{start_peak}) \leq \mathbf{alive}(\mathbb{k}) + \Sigma'(f_{start_peak})$. By hypothesis and Lemma 7.9, we that $\mathbf{alive}(cn') \leq \mathbf{alive}(\mathbb{k}') + \Sigma''(f_{start_peak}) \leq n$.

□

$$\begin{array}{c}
\text{R-SKIP} \\
vm(o, \top \mid (f, \mathbf{0}; \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}}) \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}})
\end{array}
\qquad
\begin{array}{c}
\text{R-READ-FUT} \\
vm(o, \top \mid (f, f_{\mathbf{0}}^{\checkmark}; \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}}) \parallel f'(\mathbf{0}') \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{c}[f' / \mathbf{0}], \mathbf{S}); \bar{\mathbf{f}}) \parallel f'(\mathbf{0}')
\end{array}$$

$$\begin{array}{c}
\text{R-ASYNC-CALL} \\
\frac{f'' \text{ fresh}}{vm(o, \top \mid (f, \nu f': \mathbf{m} \ o'(\bar{v}, \bar{\beta}); \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}})} \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{c}[f'' / f'], \mathbf{S}); \bar{\mathbf{f}}) \parallel (f'', \mathbf{m} \ o'(\bar{v}, \bar{\beta}))
\end{array}$$

$$\begin{array}{c}
\text{R-BIND-MTD} \\
\frac{\mathbf{m} \ \alpha(\overline{\text{Int}} \ \bar{x}, \overline{\text{VM}} \ \bar{z}) = \mathbf{c}, \ \mathbf{0}, \ \mathbf{R} \quad \mathbf{S} = \{o, \bar{o}\}}{vm(o, \top \mid \bar{\mathbf{f}}; \bar{\mathbf{f}}) \parallel (f, o \ \mathbf{m}(\bar{v}, \bar{o}))} \\
\rightsquigarrow vm(o, \top \mid \bar{\mathbf{f}}; \bar{\mathbf{f}} \cup \{(f, \mathbf{c}[\bar{v}, \bar{o} / \bar{x}, \bar{z}][o / \alpha], \mathbf{S})\}) \parallel f(\perp)
\end{array}
\qquad
\begin{array}{c}
\text{R-LCHOICE} \\
vm(o, \top \mid (f, \mathbf{c} + \mathbf{c}', \mathbf{S}); \bar{\mathbf{f}}) \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}})
\end{array}$$

$$\begin{array}{c}
\text{R-RCHOICE} \\
vm(o, \top \mid (f, \mathbf{c} + \mathbf{c}', \mathbf{S}); \bar{\mathbf{f}}) \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{c}', \mathbf{S}); \bar{\mathbf{f}})
\end{array}
\qquad
\begin{array}{c}
\text{R-ITE-T} \\
\frac{se = \mathbf{true}}{vm(o, \top \mid (f, (se)\{\mathbf{c}\}, \mathbf{S}); \bar{\mathbf{f}})} \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}})
\end{array}$$

$$\begin{array}{c}
\text{R-ITE-F} \\
\frac{se = \mathbf{false}}{vm(o, \top \mid (f, (se)\{\mathbf{c}\}, \mathbf{S}); \bar{\mathbf{f}})} \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{0}, \mathbf{S}); \bar{\mathbf{f}})
\end{array}
\qquad
\begin{array}{c}
\text{R-NEW-VM} \\
\frac{o' \text{ fresh}}{vm(o, \top \mid (f, \nu \beta; \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}})} \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{c}[o' / \beta], \mathbf{S}); \bar{\mathbf{f}}) \parallel vm(o', \top \mid \emptyset)
\end{array}$$

$$\begin{array}{c}
\text{R-RELEASE-VM} \\
vm(o, \top \mid (f, o'^{\checkmark}; \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}}) \parallel vm(o', a' \mid \bar{\mathbf{f}}'; \bar{\mathbf{f}}') \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}}) \parallel vm(o', \perp \mid \bar{\mathbf{f}}'; \bar{\mathbf{f}}')
\end{array}
\qquad
\begin{array}{c}
\text{R-GC} \\
vm(o, a \mid (f, \mathbf{0} \triangleright \emptyset, \mathbf{S}); \bar{\mathbf{f}} \cup \{\bar{\mathbf{f}}\}) \\
\rightsquigarrow vm(o, a \mid \bar{\mathbf{f}}; \bar{\mathbf{f}})
\end{array}$$

$$\begin{array}{c}
\text{R-RETURN} \\
vm(o, \top \mid (f, [\mathbf{0}], \mathbf{S}); \bar{\mathbf{f}}) \parallel f(\perp) \\
\rightsquigarrow vm(o, \top \mid (f, \mathbf{0} \triangleright \emptyset, \mathbf{S}); \bar{\mathbf{f}}) \parallel f(\mathbf{0})
\end{array}
\qquad
\begin{array}{c}
\text{R-RELEASE-BOT} \\
vm(o, \perp \mid (f, \mathbf{c}, \mathbf{S}); \bar{\mathbf{f}}) \parallel f(\perp) \\
\rightsquigarrow vm(o, \perp \mid (f, \mathbf{0} \triangleright \emptyset, \mathbf{S}); \bar{\mathbf{f}}) \parallel f(err)
\end{array}$$

$$\begin{array}{c}
\text{R-BIND-MTD-ERR} \\
vm(o', \perp \mid \bar{\mathbf{f}}; \bar{\mathbf{f}}) \parallel (f, o' \ \mathbf{m}(\bar{v})) \\
\rightsquigarrow vm(o', \perp \mid \bar{\mathbf{f}}; \bar{\mathbf{f}}) \parallel f(err)
\end{array}
\qquad
\begin{array}{c}
\text{R-BIND-MTD-SERR} \\
(f, err \ \mathbf{m}(\bar{v})) \rightsquigarrow f(err)
\end{array}$$

$$\begin{array}{c}
\text{R-RELEASE-GC} \\
vm(o, \top \mid (f, \dots o'^{\checkmark} \dots, \mathbf{S}), \bar{\mathbf{f}}) \parallel vm(o', \perp \mid \dots) \\
\rightsquigarrow vm(o, \top \mid (f, \dots \mathbf{0} \dots, \mathbf{S}), \bar{\mathbf{f}}) \parallel vm(o', \perp \mid \dots)
\end{array}
\qquad
\begin{array}{c}
\text{R-CONTEXT} \\
\frac{\mathbf{k} \rightsquigarrow \mathbf{k}'}{\mathbf{k} \parallel \mathbf{k}'' \rightsquigarrow \mathbf{k}' \parallel \mathbf{k}''}
\end{array}$$

Figure 2.8: Operational Semantics of Runtime Types

Extension of translate on extended atoms:

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

Extension of translate on process types:

$$\text{translate}(\Delta, \Psi, o, (f, \mathbf{c}, \mathbf{S})) = \begin{cases} f_{\text{peak}} = e_{1,1} & [se_1] \\ \vdots & \\ f_{\text{peak}} = e_{h_n, n} & [se_n] \\ f_{\text{net}} = f_{\text{peak}} & [\alpha = \delta] \\ f_{\text{net}} = e_{h_1, 1} & [se_1 \wedge \alpha = \top] \\ \vdots & \\ f_{\text{net}} = e_{h_n, n} & [se_n \wedge \alpha = \top] \end{cases}$$

$$\text{with } \text{translate}(\Delta, \Psi[o \mapsto \alpha], (\text{true}) 0, \mathbf{c}) = \bigcup_{i=1}^n (se_i)\{e_{1,i}; \dots; e_{h_i,i}\}$$

Extension of translate on runtime types:

$$\text{translate}(\Delta, \Psi, f(\circ)) = \{f_{\text{peak}} = 0, f_{\text{net}} = 0\}$$

$$\text{translate}(\Delta, \Psi, \text{vm}(o, \top \mid \mathbf{f}_1; \mathbf{f}_2 \dots, \mathbf{f}_n)) = \bigcup_{i \in 1..n} \text{translate}(\Delta, \Psi, o, \mathbf{f}_i)$$

$$\text{translate}(\Delta, \Psi, \text{vm}(o, \perp \mid \mathbf{f}; \bar{\mathbf{f}})) = \emptyset$$

$$\frac{\Xi = \text{EqRel}(\bar{v}) \quad \bar{v}' = \Psi(\Xi(\bar{v}))}{\text{translate}(\Delta, \Psi, (f, o \mathbf{m}(\bar{v}))) = \begin{cases} f_{\text{peak}} = \mathbf{m}_{\text{peak}}^{\Xi}(\bar{v}') \\ f_{\text{net}} = \mathbf{m}_{\text{net}}^{\Xi}(\bar{v}') \end{cases}}$$

$$\text{translate}(\Delta, \Psi, \mathbf{k} \parallel \mathbf{k}') = \text{translate}(\Delta, \Psi, \mathbf{k}) \cup \text{translate}(\Delta, \Psi, \mathbf{k}')$$

Figure 2.9: Extension of translate to runtime types

2.8 Related Work

The static analysis of resource utilization dates back to the pioneering work by Wegbreit in 1975 [58] that developed a technique for deriving closed-form expressions expressing program upper bounds. After this, a considerable number of cost analysis techniques have been developed. Those ones that are closely related to this contribution are based either on cost equations (solvers) or on amortized analysis.

Techniques based on cost equations

The techniques based on cost equations address cost analysis in three steps: *(i)* extracting relevant information out of the original programs by abstracting data structures to their size and assigning a cost to every program expression, *(ii)* converting the abstract program into cost equations, and *(iii)* solving the cost equations with an automatic tool. Recent advances have been done for improving the accuracy of the calculation upper-bounds for cost equations [2, 3, 4, 5, 6, 7, 8, 9, 20, 28, 37], we refer to [28] for an in-depth comparison of some of these tools. One of the main advantages of these techniques is the treatment of arithmetical expressions and, in particular, Presburger arithmetic conditionals. This gives the support for a more precise cost analysis of non-deterministic statements during the static analysis. As a downside, techniques that extract control flow graphs and use control flow refinement strategies (such as [6, 7]) have a less precise alias analysis and name identity management. These two features are essential for function or procedure abstraction; in fact, a weak approach would jeopardise compositional reasoning when large programs are considered. On the other hand, one of the more important contributions of these approaches has been the possibility to fully target high-end programming languages like Java [4, 5, 6] and C# [37].

Techniques based on amortized analysis

The techniques based on amortized analysis were introduced by the seminal work of Tarjan [56]. Tarjan's idea was to associate so-called potentials to program expressions by means of type systems (these potentials determine the resources needed for each expression to be evaluated). The connection between the original program and these potentials can be indeed demonstrated by a standard subject-reduction theorem. This technique has received lots of attention by applying it to both functional and imperative languages [13, 23, 43, 44, 45].

The techniques based on types are intrinsically compositional and, more importantly, type derivations can be seen as certificates of abstract descriptions of functions. These methods, in general, do not model the interaction of integer arithmetic with resource usage, which restricts the precision and the domain of analyzable properties. Although, an exception of this is one of the best works based on this approach [41].

There have been few attempts to target high-end programming languages with this technique. That is the case of [13] for Java programs and [41] for C, although in both cases several restrictions apply to the complex features of both languages.

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

A common feature of cost analysis techniques in the literature is that they analyze *cumulative resources*. That is, resources that *do not decrease* during the execution of the programs, such as execution time, number of operations, memory (without an explicit `free` operation). The presence of an *explicit or implicit* release operation considerably entangles the analysis. In [9], a memory cost analysis is proposed for languages with garbage collection. It is worth noticing that the setting of [9] is somewhat less difficult than the one proposed here, because by definition of garbage collection, released memory is always inactive. The impact of the release operation in the cost analysis is thoroughly discussed in [7, 10] by means of the notions of *peak cost* and *net cost* that we have also used in Section 2.6. It is worth noticing that, for cumulative analysis, these two notions coincide while, in non-cumulative analysis (in presence of a release operation), they are different and the *net cost* is key for computing tight upper bounds.

Recently [8] has analysed the cost of a language with explicit releases. We observe that the release operation studied in [8] is used in a very restrictive way: only locally created resources can be released. This constraint guarantees that costs of functions are always not negative, thus permitting the (re)use of non-negative cost models from the cumulative analysis. However, the main difference between our work and the work of [8] is that the target resource in our case are active objects, virtual machines, that may be concurrently hosting operations that, in turn, create and release other virtual

machines.

We conclude by discussing cost analysis techniques for concurrent systems, which are indeed very few [2, 7, 42]. In order to reduce the imprecision of the analysis caused by the nondeterminism, [2, 7] use a clever technique for isolating sequential code from parallel code, called *may-happen-in-parallel* [16]. In contrast, in our work the information of the possibly concurrent operations is carried by the proposed type system, which is a closer idea to what it is proposed in [42]. We notice though, that no one of these contributions considers a concurrent language with a powerful `release` operation that allows one deallocate resources created in outer scopes. In fact, without this operation, one can model the cost by aggregating the sets of operations that can occur in parallel, as in [7], and all the theoretical development is less complex.

2.9 Conclusions

In this Chapter we have presented the `vm1` language: a custom modeling language inspired by the Cloud Computing elasticity feature, where resources (virtual machines) can be dynamically allocated and deallocated. This language also includes features from modern programming languages like recursion and a future based concurrency model. Furthermore, resources representing virtual machines can be *active*, carrying processes that may, in turn, acquire and release other machines asynchronously. These are all features that entangle the complexity of the analysis of `vm1` programs.

Nevertheless, we have presented a technique for the static analysis of `vm1` programs. In particular, we designed a static analysis oriented to the estimation of upper bounds corresponding to the usage of objects representing virtual machines. Our approach is based on a behavioral type system that abstracts the main features of the operations on virtual machines objects. These types are expressive enough to record the effects of complex execution flows like those related to concurrency or recursion.

We demonstrate the soundness of this approach by proving the correctness of the behavioral type system and the translation into cost equations that correctly resemble (an over approximation of) the program execution.

A key detail in this type system is the treatment given to numerical expressions in the management of conditional instructions. This feature may allow achieving a higher precision during the analysis stage when such numerical expressions are expressed in Presburger arithmetics [18]. However, when this is not the case, the cost equations obtained by our technique may be too over approximated. A way to cope with this issue is to include a mech-

anism for adding annotations that can instruct the type system to provide more accurate behaviors [38, 52].

In order to calculate upper bound estimations of the virtual machines usage, the behavioral types of the program are translated into cost equations. The resulting equations can then be analyzed by special solvers oriented to such task. Although, up to date, there have been considered two possible solvers (see Section 3.3), the use of others solvers or even a combination of them could increase the precision of the analysis.

Despite the technical limitations imposed to the language features, we have demonstrated the strength of this approach. We have used the metaphor of virtual machines (VM) usage to illustrate our solution. However it is worth observing that our technique may also be used for the resource analysis of concurrent languages bearing operations of acquire (creation) and release (deletion). For example, our technique, might cover functions such as the classic `malloc` and `free` operations.

Chapter 3

The SRA tool

Summary

We describe the **(S)tatic (R)esource (A)nalysis** tool [30], which is a prototype implementation based on the theory presented in Chapter 2. We discuss a number of technical details involved in the design of this tool, specially the inference of the behavioral types from Section 2.5, and the transformation of the cost equations in Section 2.6 into an input that is adequate for an off-the-shelf solver – **CoFloCo** [28]. In this chapter, we also compare the obtained results with those from similar existing tools. Finally we conclude with the discussion of the main remarks about this solution and the future work.

3.1 Introduction

In this chapter we discuss the **SRA Tool** a prototype implementation of the theory described in Chapter 2. This solution follows the generic approach described in Chapter 1 for the development of a static analysis tool based on behavioral types.

Figure 3.1 particularizes this approach for the case of the VM usage analysis in **VML** programs. The parsing process follows an straightforward implementation based on an automatic parser generation tool. The typing process, is based on the rules defined by a type system that associates abstract behaviors to **vml** expressions (see Section 2.5). The type of each program instruction can be automatically inferred by means of a slight modification in the type system rules. Methods side-effects are also calculated as part of this typing process. Both of these tasks are done by means of a fix-point algorithm. The behaviors obtained by the typing process are then trans-




Figure 3.1: Behavioral types based approach for VML virtual machines usage analysis

lated into cost equations. The main idea underneath is to exploit the power of already existing ad-hoc costs equation solvers (CES). In the case of this solution, two CES can be used: PUBS¹ [2] and CoFloCo² [28].

Chapter contents

In this Chapter we discuss a prototype implementation of the technique described in Chapter 2. After this introduction, Section 3.2 follows with the description of the behavioral type inference for `vm1`. In Section 3.3 we implement the analysis of the virtual machine usage in `vm1` based on the information abstracted by the type system and with the aid of the `CoFloCo` solver. An overall example of our technique is discussed in Section 3.4. In Section 3.5 we compare the results of the `SRA Tool` with those of similar tools in the literature. In Section 3.6 we briefly discuss the `SRA Tool` demo and the third party tools used in its development. We conclude in Section 3.7 with the discussion of the main remarks about this solution and the future work.

3.2 Type inference

The system defined in Section 2.5 is a type-checking process, meaning that it compels the user to write the behavioural types of methods. In order to relieve users from writing verbose annotations, it is possible to turn the type checking system into a type inference. In fact, in the case of `vm1`, this adaptation is not difficult because the only problematic issue is the inference of method's effects R . This issue is easily solved by changing the rules (T-METHOD) and (T-PROGRAM) in Figure 2.6 into those of Figure 3.2. In particular, rule (INF-METHOD) uses a new judgment $\Gamma \vdash T \ m \ (\overline{T} \ x, \overline{Vm} \ z) \{ \overline{F} \ y ; s \} : m \ \alpha(\overline{x}, \overline{\beta}) \ \{ c \} : \emptyset, R'; \Gamma'$, which also re-

¹<http://costa.ls.fi.upm.es/pubs/pubs.php>

²<https://www.se.tu-darmstadt.de/se/group-members/antonio-flores-montoya/cofloco/>

(INF-METHOD)

$$\frac{\begin{array}{c} \Gamma(\mathbf{m}) = \alpha(\bar{x}, \bar{\beta}) : \mathbb{O}, \mathbf{R} \quad \mathbf{S} = \{\alpha\} \cup \bar{\beta} \quad \mathbb{O} \notin \mathbf{S} \\ \Gamma[\mathbf{this} \mapsto \alpha][\mathbf{destiny} \mapsto \mathbb{O}][\bar{x} \mapsto \bar{x}][\bar{z} \mapsto \bar{\beta}][\alpha \mapsto \emptyset\alpha][\bar{\beta} \mapsto \emptyset\bar{\beta}] \vdash_{\mathbf{S}} s : \mathcal{C}[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n] \\ \left(\Gamma_i(\gamma) = \Gamma_j(\gamma) \right)_{i,j \in 1..n, \gamma \in \mathbf{S} \cup \text{fv}(\mathbb{O})} \quad \mathbf{R}' = (\mathbf{S} \cup \text{fv}(\mathbb{O})) \cap \{\gamma \mid \Gamma_1(\gamma) = \mathbf{F}\perp\} \end{array}}{\Gamma \vdash T \mathbf{m} (\overline{\text{Int } x}, \overline{\text{Vm } z}) \{ \overline{F} y ; s \} : \mathbf{m} \alpha(\bar{x}, \bar{\beta}) \{ \mathcal{C}[\mathbf{a}_1 \triangleright \emptyset] \cdots [\mathbf{a}_n \triangleright \emptyset] \} : \mathbb{O}, \mathbf{R}'; \Gamma[\mathbf{m} \mapsto \alpha(\bar{x}, \bar{\beta})] : \mathbb{O}, \mathbf{R}'}$$

(INF-PROGRAM)

$$\frac{\Gamma \vdash \overline{M} : \overline{\mathbb{C}}, \Gamma \quad \Gamma[\mathbf{this} \mapsto \mathbf{start}] \vdash_{\mathbf{start}} s : \mathcal{C}[\mathbf{a}_1 \triangleright \Gamma_1] \cdots [\mathbf{a}_n \triangleright \Gamma_n]}{\Gamma \vdash \overline{M} \{ \overline{F} x ; s \} : \overline{\mathbb{C}}, \mathcal{C}[\mathbf{a}_1 \triangleright \emptyset] \cdots [\mathbf{a}_n \triangleright \emptyset]}$$

Figure 3.2: Behavioural inference rules of method and programs.

turns a new environment Γ' and where \mathbf{R}' may be different from the method's effects stored in $\Gamma(\mathbf{m})$. The main difference between (T-METHOD) and (INF-METHOD) is that the latter one lets Γ' record method's effects computed in the premise. The key constraint in Figure 3.2 is the left premise of (INF-PROGRAM). This premise imposes method definitions to be typed with the *same* input and output environment. That is, the environment Γ is a *fixpoint* of the rule.

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

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

It is easy to verify that 1 holds. Henceforth, the correctness of the inference process. For example, let

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

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

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

3.3 Translation of cost equations

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

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

$\text{eq}(\text{CNEW}(A), 0, [], [A = 3]).$ $\text{eq}(\text{CNEW}(A), 1, [], [A < 3]).$	$\text{eq}(\text{CREL}(A), -1, [], [A = 1]).$ $\text{eq}(\text{CREL}(A), 0, [], [A > 1]).$
--	---

In order to illustrate the encoding in CoFloCo of the cost equations in Section 2.6, we discuss the equations of the two factorial programs `fact` and `cheap_fact` in Section 2.3. The CoFloCo equations are reported in Figure 3.3.

These equations do not require any comments because they exactly correspond to the output of the `translate` function in Section 2.6. We just observe that, in addition to these equations, one needs to specify a so-called *entry point*, which corresponds to the main body. This entry point has the format

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

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

```

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

```

Figure 3.3: Cost equations of `fact` and `costly_fact` in CoFloCo format

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

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

3.4 An illustrative example

Consider the following wrong version of the `fake_method` example from Section 2.3, where the programmer has erroneously written “ $n + 1$ ” instead of “ $n - 1$ ” in the recursive invocation (henceforth the recursion is unbounded because the base case is never met). As a result one can expect that the number of created virtual machines is infinite because, at every step, only one of the two created machines is released.

```

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

```

The equations generated by the analysis are shown in Figure 3.4. The results of the analysis of these equations produce a cost of "infinity" for the `main` method, as suspected. However, let's suppose the situation in which (by changing `double_release(x,x)` to `double_release(x,y)`) both VMs are released. In this case the set of equations of `fake_method` remains almost identical with the only difference that the element (`crel(1)` in equations 10, 12, 14 and 18, now appears two times). If we re-run the cost analysis with the new equations we obtain, instead of infinity, a maximum cost of two VMs for the `main` method. Again this is the expected behavior because even though the program runs indefinitely there are no more than two VMs (in addition to the main one) co-existing at the same time.

There is an issue about Figure 3.3 that deserves to be commented. `CoFloCo` outputs an error when it is fed with that set of equations. This is because instructions 8 and 21 are mutually recursive with instructions 5 and 17-18, respectively, and mutual recursion is banned by the analyser. It is worth noticing that 8 and 21 correspond to invocations of `fact_net` and `costly_fact_net` when the carrier is ∂ . Intuitively, the idea behind these equations is that when the VM state is unknown the analyzer should assume the worst and consider the highest possible cost (the peak cost) as the net cost. The workaround to solve this problem is to remove the mutual recursion by replacing *in-place* the internal function invocation by its body.

3.5 Related tools and assessments

There are very few tools targeting resource analysis in programs that feature concurrency and negative costs due to explicit release operations. To the best of our knowledge, the most relevant tool in the literature is `SACO`³, which is based on the results of [7] and is able to address a setting similar to our one. We also compare `SRA` with `C4B`⁴ [22], a tool based on amortized analysis

³Web site at <http://ei.abs-models.org:8082/clients/web/>

⁴Web site at <http://www.cs.yale.edu/homes/qcar/aaa/>

```

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

```

Figure 3.4: Cost equations of `fake_method` and `double_release` in CoFloCo format

and using the results in [40]. Notwithstanding C4B targets only sequential programs, it may be considered as the state-of-the-art of *certified tools* for the automatic static analysis of resource usage bounds.

Some remarks on the tool comparison are in order:

- The three tools target three different languages: C4B targets C, SACO targets ABS, and SRA targets vml. Therefore, in order to compare the three tools, every test program has been rewritten according to the corresponding language of the tool. The sources of these tests are available in github⁵.
- Every tool uses different metrics for quantitative analysis. SRA quantifies the resource usage by means of the operations `new VM()` and `release`. C4B has two metrics: the (i) back-edge metrics that assigns a cost of 1 to every back edge in the control-flow graph, thus not allowing negative costs, and (ii) the tick metric that uses a special operation

⁵<https://github.com/abelunibo/SRA-Tests-Files>

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

Figure 3.5: Comparison of C4B, SACO, and SRA for some simple programs

`tick(n)` that has a cost of n where n can be either a positive or negative number. SACO, on the other hand, considers both of these kinds of metrics and, in addition, it also support other ones that are specific to the domain of the targeted language.

- Both SACO and SRA separate *peak costs* and *net costs*. This distinction is not made in C4B because it does not address concurrency. Therefore we restrict the results of Table 3.5 only to the *net costs* of the corresponding programs.
- Some programs are *not analyzable* with some of the tools, this is either because (i) the programs include some features that are not covered by the analyzer (i.e. concurrency in the case of the C4B tool) or (ii) the underlying solver cannot deal with the resulting equations (i.e. non polynomial equations in the case of SRA or non monotonic equations in the case of SACO)

Table 3.5 reports the results of our analysis. It is worth observing that SRA returns better results than both C4B and SACO for the program `betterThanAmortized`. This program, reported in Figure 3.6, defines a method `betterThanAmortised(n,m,x)` that recursively invokes itself and, when ($n > m$) every invocation has cost 1, otherwise it has cost 2. Therefore, the cost of `betterThanAmortised(2*n,n,x)` is $3*n$ and this is the result of SRA. However, the other techniques do not recognize that the most costly branch is executed only half of the times and return $2*(2*n)=4*n$. Such precision is achieved in SRA by means of the treatment given to numerical expressions, and by the use of a solver (as [28]) capable to deal with conditional branches accurately. This precision is usually lost in techniques based on type systems because of the unifications that are necessary to enforce on the branches of conditionals, in order to type the continuations.


```

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

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

```

Figure 3.6: The program `betterThanAmortized`

Our behavioral type system, instead, allows the necessary flexibility to overcome this issue. We also observe that `betterThanAmortized` is not an exotic, ad-hoc program but a rather common programming pattern (a sequence of nested conditional statements).

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

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

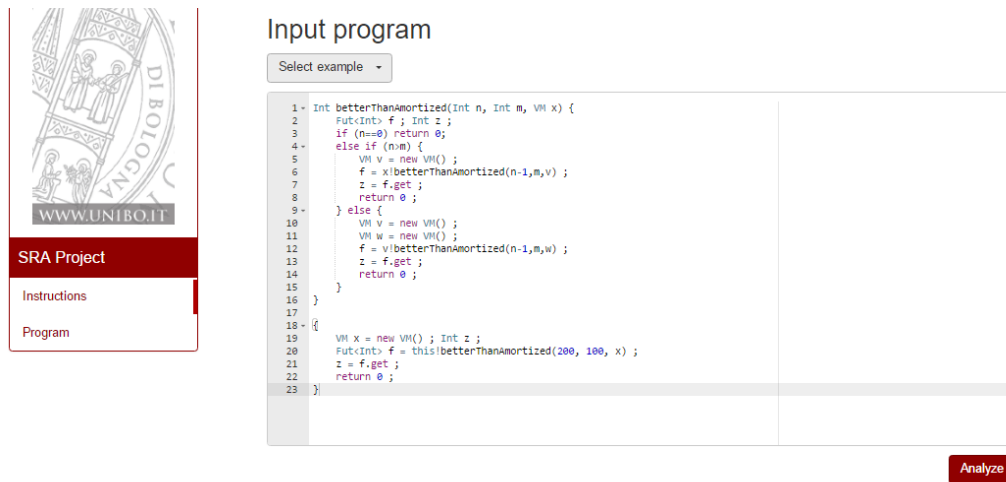


Figure 3.7: SRA tool web demo screen shot

3.6 SRA Deliverable

The SRA tool is available through a dedicated website [30]. The demo page provides a set of default `vm1` program examples that are designed to highlight the main features of the tool. The demo page also allows the user to insert custom programs and analyze them.

The prototype components match the general description given in Section 1.2. There are four main components: (i) A `vm1` parser, (ii) an inference mechanism, (iii) a finite state representation of the program and (iv) a cost analysis process.

The Figure 3.7 shows the demo page of the SRA tool.

Third-party tools involved The orchestration of the SRA Tool is implemented in Java and it benefits from two out-of-the-box solutions.

The parsing process is done with the help of the ANTLR tool⁶, this tool provides automatic code generation for taking care of the syntactic analysis of the program as well as the abstract implementation of a semantic checking process that can be later on customized.

There are a couple of solvers capable to solve the cost equations resulting from the translation step in Section 3.3, PUBS⁷ [2] and CoFloCo⁸ [28].

⁶<http://www.antlr.org/>

⁷<http://costa.ls.fi.upm.es/pubs/pubs.php>

⁸<https://www.se.tu-darmstadt.de/se/group-members/antonio-flores-montoya/cofloco/>

3.7 Conclusions

To the best of our knowledge the tool here described presents the first static analysis technique that computes upper bounds of virtual machines usage in concurrent programs, that may create and, more importantly, may release such machines.

Our analysis consists of (i) a type system that extracts relevant information about resource usages in programs, called behavioral types; (ii) an automatic translation that transforms these types into cost expressions; and (iii) the application of automatic solvers, like `CoFloCo` or `PUBS`, to compute upper bounds of the usage of virtual machines in the original program.

A relevant property of our technique is its modularity. In this solution we have applied the technique to a small, ad-hoc, language. However, a properly extension of the type system, would allow to apply the analysis to other more complex languages. In addition, by changing the translation algorithm, it is possible to target other solvers that may compute, for certain patterns, better upper bounds. For example, those based on multivariate amortized analysis [40] or in theorem provers [1].

As we have stated there are not currently other static resource analyzers with the same target of the `SRA Tool`. We have, though, assessed the performance of our analyzer by comparing with two state-of-the-art level analyzers: `C4B` and `SACO`. The tested examples demonstrated the programming patterns on which `SRA` excels, like the handling of conditional blocks, the deallocation of resources and the handling of combinations of recursive and concurrent behaviors.

Part II

The deadlock analysis problem

Chapter 4

Deadlock detection in JVMML

Summary

This Chapter presents a new solution for the deadlock detection in Java programs at static time. This technique targets specifically the Java intermediate language – the Java bytecode – which has the additional value of allowing the analysis of other languages compiled to the same bytecode, like Scala. This technique uses behavioral types to extract abstract models out of the bytecode instructions. These models are subsequently analyzed by means of a fix-point decision algorithm which detects the presence of circular dependencies. The correctness of this technique is demonstrated for a subset of the Java bytecode instructions.

4.1 Introduction

In concurrent languages, a *deadlock* is a circular dependency between a set of threads, each one waiting for an event produced by another thread in the set. In the Java programming language, deadlocks are usually *resource-related*, namely they are caused by operations ensuring different threads the exclusive access to a set of resources. Java programs may also lead to the so-called *communication-related* deadlocks, which are common in network based systems. These deadlocks, which are thoroughly studied in [34, 48], are out of the scope of this work.

As a way to ensure consistency, the Java Language Specification provides different mechanisms to enforce a *should-happen-before* relationship on memory operations over shared variables [36, Chapter 17].

Our analysis focuses on some of these enforcing mechanisms, in particular: the *synchronized constructions*, of both code blocks and method blocks, and

the `Thread.start()` and `Thread.join()` methods. There are also other mechanisms that are not currently handled by the technique here described, such as, the use of `volatile` variables and the higher-level synchronization API defined on package `java.util.concurrent`.

The dependencies enforced by this *should-happen-before* relation may, in some times, produce a circularity. Whenever the execution flow may reach a state with such a circularity, we are in presence of a necessary condition for a deadlocked execution. However, because of the scheduling process this is not a sufficient condition. The Figure 4.1 shows (a time line representation of) some examples of deadlocked programs.

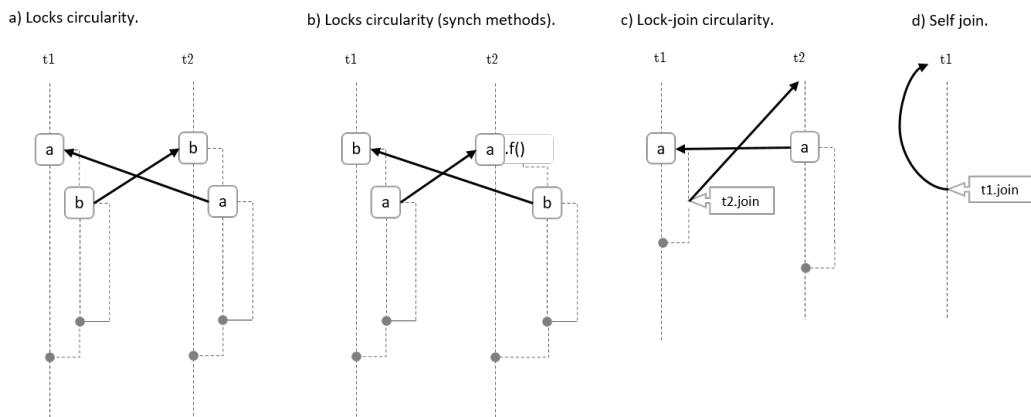


Figure 4.1: Cases of dependency circularities that may lead to deadlocks. (Lock acquisitions are represented with squares, the corresponding release is marked with a circle)

On the other hand, the mere presence of a circularity may be sometimes misleading. This is the case when other *should-happen-before* relations in the program ensure that a state, in which such circularity is verified, will not be reached by the execution flow. See Figure 4.2.

Targeting JVM, the Java bytecode

The solution here proposed addresses the compilation target of every Java application: the *Java Virtual Machine Language*, JVM, also called Java bytecode.

The decision of addressing JVM instead of Java was motivated by two reasons: Java is syntactically too complex and it has no reference semantics. On the contrary, JVM a rather simple stack language – it has 198 instructions




Figure 4.2: Cases of dependency circularities that **do not produce deadlocks**. (Lock acquisitions are represented with squares, the corresponding release is marked with a circle)

– and has a reference semantics that is defined by the behavior of the *Java Virtual Machine* (JVM) [36, Chapter 6].

This opinion also seems supported by a number of type systems that have been defined in the past for demonstrating the correctness of the Java bytecode Verifier [29, 49, 55]. It is worth to observe that reasoning on Java bytecode does not narrow our original goal, because as the Java compilation target its behavior matches exactly the one intended by the original source. Furthermore, this also does not simplify too much the analysis because it retains the same complex features of Java.

Analyzing JVMIL has also other relevant advantages: addressing programming languages that are compiled to the same bytecode, such as *Scala* [53], and the possibility to analyze proprietary software whose sources are not available.

We have defined an inference system that extracts abstract models out of Java bytecode instructions. This inference system consists of a number of rules, identical for most of the instructions except for those that have some effect in the synchronization process: invocations, locks acquisitions and releases, object manipulation and control flow operations.

Chapter contents

In this Chapter we discuss the main theoretical components of our deadlock analysis solution. A motivational analysis of the deadlock problems in programming languages and specially in `Java` is presented in Section 4.2. To ease the presentation of the theoretical aspects of this work we propose the language `JVMLd` in Section 4.3, which is a subset of `JVML` that include basics constructs and instructions for concurrent operations. `JVMLd` is simpler than `JVML`, it leaves out some complex features of the original language like exception handling, inheritance, arrays manipulation, etc. However, `JVMLd` itself still poses a difficult challenge for deadlock analysis, with features like recursive invocations or fields manipulation. We present a behavioral type system rules for `JVMLd` in Section 4.5. The ulterior analysis of these types for the detection of circularities is presented in Section 4.6. We also discuss, in Section 4.7, one more type of deadlock in `Java` and a possible alternative to extend the theory here described in order to detect it. In Section 4.8 we discuss the extensions of the behavioral type system in order to cope with full `JVML` programs. The full proof of the correctness of the presented behavioral type system and its analysis is discussed in Section 4.9. In Section 4.10 we review the existing works in the literature related to the deadlock analysis and to the static analysis of `Java` programs. Finally this chapter concludes with the main remarks about the presented technique.

4.2 Motivation and Problem Statement

Deadlocks are a common threat of concurrent programs, which occur when a set of threads are blocked, due to each attempting to acquire a lock held by another. Such errors are difficult to detect or anticipate, since they may not occur during every execution, and may have catastrophic effects for the overall functionality of the software system. In 1994, a deadlock flaw in the Automated Baggage System of Denver's was one of the causes of losses for more than 100 million dollars¹.

At the time of writing this document, the *Oracle Bug Database*² reports more than 40 unresolved bugs due to deadlocks, while the *Apache Issue Tracker*³ reports around 400 unresolved deadlock bugs. These two databases refer to programs written in `Java`, a mainstream programming language in a lot of domains, such as web and Cloud applications, user applications and

¹<http://ntl.bts.gov/DOCS/rc9535br.html>

²<http://bugs.java.com/>

³<https://issues.apache.org/jira>

mobile applications.

The objective of our research is to design and implement a technique capable of detecting potential deadlock bugs of Java programs *at compilation time* and, to this aim, we propose an end-to-end automatic static analysis tool. Our implementation handles most of the features of Java, including threads, synchronization operations, exceptions, static members, arrays, interfaces, polymorphism, and recursive data types.

Deadlock detection of Java programs is a difficult and dubious task. It is difficult because concurrency in Java is modeled by threads that may perform read/write operations over shared variables. The order in which concurrent operations are performed depends on the scheduling strategy implemented in the Java Virtual Machine (JVM). Therefore deadlocks may not occur during every execution. Additionally, Java is a full-fledged programming language, with many features, and with many third-part libraries that are written directly in machine code. It is dubious because Java has no reference formal semantics, therefore it is not possible to deliver any soundness proof.

Problem statement

Precise deadlock analysis is, in general, an undecidable problem. This can be proved in the following way (by contradiction):

Undecidability of static deadlock detection. Suppose that for all programs P and for all of its possible inputs X_P there is an algorithm $\text{hasDeadlock}(P, X_P)$ such that:

$\text{hasDeadlock}(P, X_P)$ returns true \iff $P(X_P)$ reaches a deadlocked state

Then consider the following program Z :

```
Z (input_program){
  if(hasDeadlock(Z, input_program))
    return;
  else
    causeDeadlock();
}
```

Consider now running the program Z with itself as argument: $Z(Z)$. There are two choices, (i) program $Z(Z)$ returns (line 2 is reached), then there is a contradiction because $\text{hasDeadlock}(Z, Z)$ returns true only if the program deadlocks; and (ii) $Z(Z)$ causes a deadlock (line 4 is reached),

then there is also a contradiction because a deadlock is only reached when `hasDeadlock(Z,Z)` returns false. \square

In this work we focus on an alternative, but also very challenging, goal:

Given a program P to have an algorithm `mayDeadlock(P)` such that: `mayDeadlock(P)` returns true \implies P may reach a deadlocked state

Notice that if proved sound an algorithm like `mayDeadlock(P)` may have a great value, since it is able to verify programs as deadlock free. Of course, some precision from this algorithm is required. A naive "return true" implementation is sound, although not useful at all.

This leads also to an implicit problem:

Given the fact that a program may reach a deadlocked state: to produce the trace that reaches this state.

This information is a key clue for either resolving the deadlock problem or discarding a false positive.

4.3 The language JVMML_d

In JVMML_d, a program is a collection of *class files* whose methods have bodies written in JVMML_d bytecode. This bytecode is a partial map from *addresses* ADDR to instructions. Addresses, which are ranged over L, L', \dots are intended to be nonnegative integers and we use the function $L + 1$ that returns the least address that is strictly greater than L . When P is a bytecode, we write $dom(P)$ to refer to its domain (the set of addresses) and we assume that $0 \in dom(P)$ for every bytecode P .

Class files use a number of *names*: for classes, ranged over by C, D, \dots , for fields, ranged over by f, f', \dots , for methods, ranged over by m, m', \dots , and for local variables, ranged over by x, y, \dots . A possible empty sequence of names or syntactic categories of the following grammar is written by over-lining the name or the syntactic category, respectively. For instance a sequence of local variables is written \bar{x} . Class files CF are defined by the grammar:

$$\begin{aligned}
 CF & ::= \text{class } C \{ \overline{fields} : \overline{FD} \ \overline{methods} : \overline{MD} \} \\
 FD & ::= C.f : T \\
 MD & ::= T m (C, \overline{T}) P \mid \text{synchronized } T m (C, \overline{T}) P \\
 T & ::= \top \mid \text{int} \mid C
 \end{aligned}$$

where \top is a special type that include all the other types (any value of any type has also type \top). This type will represent values that are unusable in our static semantics. The type name C represents a class type. In JVM_D, classes are *not recursive*.

4.3.1 Syntax

Instructions *Instr* of JVM_D bytecode are of the following form:

$$\begin{aligned} \text{Instr} ::= & \text{inc} \mid \text{pop} \mid \text{push} \mid \text{load } x \mid \text{store } x \mid \text{if } L \mid \text{goto } L \\ & \mid \text{new } C \mid \text{putfield } C.f : T \mid \text{getfield } C.f : T \\ & \mid \text{invokevirtual } C.m(\bar{T}) \\ & \mid \text{monitorenter} \mid \text{monitorexit} \mid \text{start } C \mid \text{join } C \\ & \mid \text{return} \end{aligned}$$

4.3.2 Informal semantics

The semantics of JVM_D follows exactly the semantics of the corresponding instructions in JVM described in [36, Chapter 6]. A full formal description of the operational semantics can be found in Section 4.9.2). The informal meaning of some of these instructions is as follows:

- **inc** increments the content of the stack; **pop** and **push**, respectively, pops and pushes the integer 0 on the stack; **load** x and **store** x respectively loads the value of x on the stack and pops the top value of the stack by storing it in x ; **if** L pops the top value of the stack and either jumps to the instruction at address L , if it is nonzero, or goes to the next instruction; **goto** L is the unconditional jump;
- **new** C allocates a new object of type C , initializes it and pushes it on top of the stack; **putfield** $C.f : T$ pops the value on the stack and the underlying object value, and assigns the former to the field f of the latter; **getfield** $C.f : T$ pops the object on the stack and pushes the value in the field f of that object;
- **invokevirtual** $C.m(T_1, \dots, T_n)$ pops n values from the stack (the arguments of the invocation) and dispatches the method m on the object on top of the stack; when the method terminates, the returned value is pushed on the stack;
- **monitorenter**, **monitorexit** are the synchronization primitives that pop the object on the stack and respectively lock and unlock it;

- **start** *C* creates and starts a new thread for the object on top of the stack; **join** *C* pops the thread on top of the stack and joins it with the current one. These two operations correspond to `invokevirtual java/lang/Thread/start()` and `invokevirtual java/lang/Thread/join()` on a thread of class *C* in JVM, respectively. We separate them from `invokevirtual` in order to provide more structure to our semantics (they have an effect on the set of threads – see the operational semantics in Section 4.9.2);
- **return** terminates program execution.

4.3.3 JVM_d deadlock example: the Network class

Figure 4.3 reports a Java class called `Network` and some of its bytecode. The corresponding `main` method creates a network of `n` threads by invoking `buildNetwork` – say t_1, \dots, t_n – that are all potentially running in parallel with the caller – say t_0 . Every two adjacent threads share an object, which are also created by `buildNetwork`.

Every thread t_i locks the two adjacent objects, that are passed as (implicit) arguments of the thread, and terminates – this is performed by the method `takeLocks`. It is well-known that, if the network is circular – the thread t_n is sharing one of its objects with t_0 and if all the thread have the strategy of locking objects then a deadlock may occur. On the contrary, if the network is not circular, no deadlock will ever occur.

The `buildNetwork` method will produce a deadlock depending on its actual arguments: it is deadlock-free when it is invoked with two different objects, otherwise it may deadlock (if also $n > 0$). Therefore, in the case of Figure 4.3, the program is deadlocked, while it is deadlock free if we comment the instruction `buildNetwork(n,x,x)` and uncomment `buildNetwork(n,x,y)`.

The problematic issue of `Network` is that the number of threads is not known statically – `n` is an argument of `main`. This is displayed in the bytecode of `buildNetwork` in Figure 4.3 by the instructions at addresses 30 and 37 that respectively created a new thread and start it, and by the recursive invocation at instruction 47.

Our technique is powerful enough to cope with such problems and predict the faulty behavior in case of the invocation `buildNetwork(n,x,x)` and the correct behavior if the invocation is `buildNetwork(n,x,y)`.

```

class Network{

public void main(int n){
  Object x = new Object();
  Object y = new Object();

  // deadlock
  buildNetwork(n, x, x);

  // no deadlock
  //buildNetwork(n, x, y);
}

public void buildNetwork(int n,
Object x, Object y){
  if (n==0) {

    takeLocks(x,y) ;

  } else {

    final Object z = new Object() ;

    //anonymous Thread child class
    Thread thr = new Thread(){
      public void run(){
        takeLocks(x,z) ;
      };

      thr.start();
      this.buildNetwork(n-1,z,y) ;
    }
  }

public void takeLocks(Object x, Object y){
  synchronized (x) {
    synchronized (y) {

    }
  }
}
}

public void buildNetwork(int n, Object x, Object y)
0 iload_1          //n
1 ifne 13
4 aload_0          //this
5 aload_2          //x
6 aload_3          //y
7 invokevirtual 24 //takeLocks(x, y):void
10 goto 50
13 new 3
16 dup
17 invokespecial 8 //Object()
20 astore 4        //z
22 new 26
25 dup
26 aload_0         //this
27 aload_2         //x
28 aload 4         //z
30 invokespecial 28 //Network1(this, x, z)
33 astore 5        //thr
35 aload 5         //thr
37 invokevirtual 31 //start():void
40 aload_0         //this
41 iload_1         //n
42 iconst_1
43 isub
44 aload 4         //z
46 aload_3         //y
47 invokevirtual 36 //buildNetwork(n-1, z, y):void
50 return

public void takeLocks(Object x, Object y)
0 aload_1;        //x
1 dup;
2 astore_3;
3 monitorenter;  //acquires x
4 aload_2;        //y
5 dup;
6 monitorenter;  //acquires y
7 monitorexit;  //releases y
8 aload_3;
9 monitorexit;  //releases x
16 return;

```

Figure 4.3: Java Network program and corresponding bytecode of methods buildNetwork and takeLocks. Comments in the bytecode give information of the objects used and/or methods invoked in each instruction

4.4 Preliminaries

Type values and flattened record types. Our technique traces dependencies between objects and it is therefore critical that the static semantics keeps track of object identities. let OBJ be the set of *object names*, ranged over by a, b, \dots . OBJ includes the two special names *null* and *void* and the *thread names* (in **Java**, threads are objects as well), which are ranged over by t, t', \dots . We also use x, y, z, \dots to range over (generic) *names* and X, Y, Z, \dots to range over *variable names* to be used in the typing of methods and to be instantiated when methods are invoked.

Let (\mathbf{h}, \leq) be the lattice of *accesses*, where $\mathbf{h} = \{-, \mathbf{r}, \mathbf{w}\}$ – with $-$ meaning no access, \mathbf{r} meaning read access, \mathbf{w} meaning write (and possible read) access – and $- \leq \mathbf{r} \leq \mathbf{w}$. Since (\mathbf{h}, \leq) is a lattice, we will use the operation of least upper bound \sqcup . Let

– *type values* χ and *flattened record types* ς be the terms

$$\chi ::= \top \mid \text{int} \mid X \mid a \qquad \varsigma ::= ([\mathbf{f}^{\mathbf{h}} : \chi], \mathbf{C})$$

It is worth to remark that flattened record types also display the kind of access to fields and bear their class set, which is a singleton in JVML_d (but not in full **Java**, where sets are not singleton because of inheritance, see Section 4.8.5). We always shorten \mathbf{f}^- into \mathbf{f} ; therefore record types are record types with accesses where the accesses to fields are always $-$. We use $\mathcal{T}, \mathcal{T}', \dots$ for denoting sets of types $a[\mathbf{f}^{\mathbf{h}} : \chi]$.

Let

$$\mathbf{f}^{\mathbf{h}} : \chi \oplus \mathbf{f}^{\mathbf{h}'} : \chi' \stackrel{\text{def}}{=} \begin{cases} \mathbf{f}^{\mathbf{h} \sqcup \mathbf{h}'} : \chi & \text{if } \mathbf{h}, \mathbf{h}' \leq \mathbf{r} \text{ and } \chi = \chi' \\ \mathbf{f}^{\mathbf{w}} : \text{int} & \text{if } \mathbf{h} \sqcup \mathbf{h}' = \mathbf{w} \text{ and } \chi = \text{int} = \chi' \\ \mathbf{f}^{\mathbf{w}} : \top & \text{otherwise} \end{cases}$$

It is worth to notice that $\mathbf{f}^{\mathbf{h}} : X \oplus \mathbf{f}^{\mathbf{h}'} : Y$ is \top when $X \neq Y$. Let also

$$([\dots, \mathbf{f}_i^{\mathbf{h}_i} : \chi_i, \dots], \mathbf{C}) \oplus ([\dots, \mathbf{f}_i^{\mathbf{h}'_i} : \chi'_i, \dots], \mathbf{C}) \stackrel{\text{def}}{=} ([\dots, \mathbf{f}_i^{\mathbf{h}_i} : \chi_i \oplus \mathbf{f}_i^{\mathbf{h}'_i} : \chi'_i, \dots], \mathbf{C})$$

(it is assumed that the fields of the records are the same – because they have the same class type – and the operation is performed on every field).

Environments. The foregoing static semantics uses abstract heaps, called *environments* Γ , which map names to type values or to flattened record types. There are two basic operations on environments: one for *sequential*

composition – the *update* $\Gamma[\Gamma']$ – and one for *parallel composition* – the *merge* $\Gamma \oplus \Gamma'$. They are defined as follows

$$\Gamma[\Gamma'](a) = \begin{cases} \Gamma(a) & \text{if } a \in \text{dom}(\Gamma) \setminus \text{dom}(\Gamma') \\ \Gamma'(a) & \text{if } a \in \text{dom}(\Gamma') \setminus \text{dom}(\Gamma) \\ ([\dots, \mathbf{f}_i^{\mathbf{h}_i \cup \mathbf{h}'_i} : \chi'_i, \dots], \mathbf{C}) & \text{if } \Gamma(a) = ([\dots, \mathbf{f}_i^{\mathbf{h}_i} : \chi_i, \dots], \mathbf{C}) \\ & \text{and } \Gamma'(a) = ([\dots, \mathbf{f}_i^{\mathbf{h}'_i} : \chi'_i, \dots], \mathbf{C}) \end{cases}$$

$$(\Gamma \oplus \Gamma')(a) = \begin{cases} \Gamma(a) & \text{if } a \in \text{dom}(\Gamma) \setminus \text{dom}(\Gamma') \\ \Gamma'(a) & \text{if } a \in \text{dom}(\Gamma') \setminus \text{dom}(\Gamma) \\ \Gamma(a) \oplus \Gamma'(a) & \text{otherwise} \end{cases}$$

For example, let $\Gamma = a \mapsto ([\mathbf{f}^{\mathbf{h}} : b], \mathbf{C})$ and $\Gamma' = a \mapsto ([\mathbf{f}^{\mathbf{w}} : c], \mathbf{C})$. Then $\Gamma[\Gamma'] = a \mapsto ([\mathbf{f}^{\mathbf{w}} : c], \mathbf{C})$, which is the standard update of an environment plus the recording of the writing operation, while $\Gamma \oplus \Gamma' = a \mapsto ([\mathbf{f}^{\mathbf{w}} : \top], \mathbf{C})$, namely the field value is undefined because a *race condition* on the field \mathbf{f} of a caused by the parallel execution of two threads.

In the static semantics, methods have types that are different from type values. This is because, in methods types, the (tree) structure of methods' arguments must be completely described. For this reason and also because class types in JVML_d are not recursive, we use (unflattened) record types ϑ , which are terms

$$\vartheta ::= \top \mid \mathbf{int} \mid X \mid (a[\overline{\mathbf{f}^{\mathbf{h}} : \vartheta}], \mathbf{C}).$$

Record types that have fields with accesses \mathbf{f}^- (which are always shortened into \mathbf{f}) are ranged over by the metavariables τ, τ', \dots . We shorten $a[]$ into a ; in particular $\text{null}[]$ and $\text{void}[]$ are shortened into null and void , respectively. Let $\text{root}(\cdot)$ be the function defined as follows: $\text{root}((\text{void}, \mathbf{C})) = \varepsilon$, $\text{root}((a[\overline{\mathbf{f}^{\mathbf{h}} : \vartheta}], \mathbf{C})) = a$, $\text{root}(\cdot)$ is the identity otherwise. We notice that the sequence $\text{root}(\text{void}) \cdot S$ is equal to S .

The operation \oplus is extended to record types as follows

$$\vartheta \oplus \vartheta' \stackrel{\text{def}}{=} \begin{cases} (a[\dots, \mathbf{f}^{\mathbf{h}} : \vartheta_{\mathbf{f}} \oplus \mathbf{f}^{\mathbf{h}'} : \vartheta'_{\mathbf{f}}, \dots], \mathbf{C}) & \text{if } \vartheta = (a[\dots, \mathbf{f}^{\mathbf{h}} : \vartheta_{\mathbf{f}}, \dots], \mathbf{C}) \\ & \text{and } \vartheta' = (a[\dots, \mathbf{f}^{\mathbf{h}'} : \vartheta'_{\mathbf{f}}, \dots], \mathbf{C}) \\ \top & \text{if } \text{root}(\vartheta) \neq \text{root}(\vartheta') \end{cases}$$

There is a straightforward way to get an environment from a record type and, conversely, to transform an environment and an object name into its (unflattened) record type. These two functions, called $\text{flat}(\cdot)$ and $\text{mk_tree}(\cdot)$, are defined as follows:

$$\text{flat}(\vartheta) = \begin{cases} \emptyset & \text{if } \vartheta \in \{\top, \mathbf{int}, X\} \\ a \mapsto ((\overline{\mathbf{f}^{\mathbf{h}} : \text{root}(\vartheta)}], \mathbf{C}) \oplus (\bigoplus_{\vartheta' \in \overline{\vartheta}} \text{flat}(\vartheta')) & \text{if } \vartheta = (a[\overline{\mathbf{f}^{\mathbf{h}} : \vartheta}], \mathbf{C}) \end{cases}$$

$$mk_tree(\Gamma, \chi) = \begin{cases} \chi & \text{if } \chi \in \{\top, \text{int}, X\} \\ (a[\overline{\mathbf{f}} : mk_tree(\Gamma, \chi)], \mathbf{C}) & \text{if } \chi = a \text{ and } \Gamma(a) = ([\overline{\mathbf{f}^h} : \chi], \mathbf{C}) \end{cases}$$

We let

$$flat(\{\vartheta_1, \dots, \vartheta_n\}) \stackrel{def}{=} \bigcup_{1 \leq i \leq n} flat(\vartheta_i)$$

and

$$mk_tree(\Gamma, (\chi_1, \dots, \chi_k)) \stackrel{def}{=} (mk_tree(\Gamma, \chi_1), \dots, mk_tree(\Gamma, \chi_k))$$

we finally define $typeof(\Gamma, \chi) = \mathbf{C}$ if $\Gamma(\chi) = ([\overline{\mathbf{f}^h} : \chi'], \mathbf{C})$, it is $typeof(\Gamma, \chi) = \text{int}$ if $\Gamma(\chi) = \text{int}$, undefined otherwise.

Lams. Behavioral types are *lams* [34], which express object dependencies and method invocations. The syntax of lams, noted ℓ , is the following one:

$$\ell ::= \mathbf{0} \mid (a, b)_t \mid \mathbf{C.m}(a[\overline{\mathbf{f}} : \tau], \overline{\tau}) \rightarrow \tau \mid (\nu a)\ell \mid \ell \& \ell \mid \ell + \ell$$

The type $\mathbf{0}$ is the empty type; $(a, b)_t$ specifies a dependency between the object a and the object b that has been created by the thread t ; $\mathbf{C.m}(a[\overline{\mathbf{f}} : \tau], \overline{\tau}) \rightarrow \tau''$ defines the invocation of $\mathbf{C.m}$ with carrier $a[\overline{\mathbf{f}} : \tau]$, with arguments $\overline{\tau}'$ and with returned record type τ'' . The last two elements of the tuple $\overline{\tau}'$ record the thread t that performed the invocation and the last object name b whose lock has been acquired by t . These two arguments will be used by our analyzer to build the right dependencies between callers and callees. The operation $(\nu a)\ell$ creates a new name a whose scope is the type ℓ ; the operations $\ell \& \ell'$ and $\ell + \ell'$ define the conjunction and disjunction of the dependencies in ℓ and ℓ' , respectively.

The operators $+$ and $\&$ are associative and commutative. As usual, we shorten $\ell_1 + \dots + \ell_n$ and $\ell_1 \& \dots \& \ell_n$ into $\sum_{i \in 1..n} \ell_i$ and $\&_{i \in 1..n} \ell_i$, respectively.

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

$$\mathbf{C.m}(a[\overline{\mathbf{f}} : \tau], \overline{\tau}', t, b) \rightarrow \tau'' = \ell_{\mathbf{C.m}}$$

with $\ell_{\mathbf{C.m}}$ being the *body* of $\mathbf{C.m}$, and ℓ is the *main lam*, the lam corresponding to the method in the role of program entry point. The type τ'' is considered an argument of the lam function as well.

Behavioural Class Table. A *behavioural class table*, noted BCT , is a map from pairs $\mathbf{C.m}$ to *method types*:

$$\mathbf{C.m} : (\bar{\tau}, t, a) \rightarrow (\nu \bar{c}) \langle \vartheta, \mathcal{T}, K, \bar{\vartheta}, \ell \rangle$$

where

- $\bar{\tau}$ is the tuple of record types of the carrier and of the arguments; t is the thread name of the caller and a is the last lock taken (and not released) by the caller;
- $(\nu \bar{c})$ are the names that have been created by the method (and occur in the part in angular brackets);
- ϑ is the record type of the returned value; this type may contain fresh object names that are created within the method body;
- \mathcal{T} is the set of thread record types that have been created by the method and that survive (have not been synchronized) after the method execution; (this set only contains record types where root names are fresh);
- K is the set of thread names that occur in the arguments (in $\bar{\tau}$) and that have been synchronized within the body of $\mathbf{C.m}$;
- $\bar{\vartheta}$ is the record type highlighting the accesses and the changes to the arguments performed by the invocation;
- ℓ is the lam of the method $\mathbf{C.m}$.

Let $\text{BCT}(\mathbf{C.m}) = (\bar{\tau}, t, a) \rightarrow (\nu \bar{c}) \langle \vartheta, \mathcal{T}, K, \bar{\vartheta}, \ell \rangle$; we define the *instance of a method type*:

let σ be a substitution for object names and variable names occurring in $\bar{\tau}, t, a$ such that $(\bar{\tau}, t, a)\sigma = (\bar{\tau}', t', a')$; let also \bar{b} be a tuple of names such that there is no clash with names in $\text{dom}(\sigma)$. Then

$$\text{BCT}(\mathbf{C.m})(\bar{\tau}', t', a')(\bar{b}) \stackrel{\text{def}}{=} (\langle \vartheta, \mathcal{T}, K, \bar{\vartheta}, \ell \rangle \{\bar{b}/\bar{c}\})\sigma .$$

We notice that this instantiation operation requires that $(\bar{\tau}, t, a)\sigma = (\bar{\tau}', t', a')$, namely actual arguments $\bar{\tau}'$ *must be more specific* than formal parameters patterns $\bar{\tau}$. For instance, it is not possible that $\bar{\tau}' = \top, \bar{\tau}''$ because the first element of $\bar{\tau}$ is never a \top value (it has at least a root object name).

Typing judgments. Let P be a JVML_d bytecode. P will be typed by the judgment:

$$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \ell; \Psi$$

where:

- BCT is the behavioural class table that maps every method of the program to its lam;
- Γ, F, S, Z, T, K are *vectors* indexed by the addresses of P ;
- the environment Γ_i maps object names to record types at address i ;
- the map F_i takes local variables and returns type values at address i ; let F_\top be the map such that $F(x) = \top$, for every x ;
- the stack S_i returns a sequence of value types at address i ;
- Z_i is the *sequence of object names* locked at address i ;
- T_i is the *set of thread names* that are alive at instruction i and that have been created locally either by P or by a method invoked by P ;
- K_i returns the *set of thread names* that have been synchronized at instruction i ;
- t is the thread name where P is executed;
- ℓ is the behavioral type of P ;
- Ψ may be either empty, in this case the judgment is shortened into $\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \ell$, or a tuple $(\vartheta, Z_i, T_i, K_i, \Gamma_i)$. We let \sqcup be the commutative and associative operator defined as follows

$$(\vartheta, Z_i, T_i, K_i, \Gamma_i) \sqcup \emptyset \stackrel{\text{def}}{=} (\vartheta, Z_i, T_i, K_i, \Gamma_i)$$

$$(\vartheta, Z_i, T_i, K_i, \Gamma_i) \sqcup (-, Z_j, T_j, K_j, \Gamma_j) \stackrel{\text{def}}{=} (\vartheta, Z_i, T_i \cup T_j, K_i \cup K_j, \Gamma_i \oplus \Gamma_j)$$

$$(\vartheta, Z_i, T_i, K_i, \Gamma_i) \sqcup (\vartheta', Z_j, T_j, K_j, \Gamma_j) \stackrel{\text{def}}{=} (\vartheta \oplus \vartheta', Z_i, T_i \cup T_j, K_i \cup K_j, \Gamma_i \oplus \Gamma_j)$$

The symbol $-$ represents the absence of a value – this happens in case of abrupt termination, see rules in Figure 4.8. We remark that the above operation is only defined on tuples with the element Z_i equal.

Let $Z_i = a_1 \cdots a_n$ be the sequence of (locked) object names at instruction i – we assume that the leftmost object name is the more recent one that has been taken. We define $[Z_i]$ and \widehat{Z}_i^t as follows

$$\begin{aligned} [Z_i] &= a_1 \\ \widehat{Z}_i^t &= \&_{j \in 2..n} (a_j, a_{j-1})_t \end{aligned}$$

The sequence Z_i may contain twice the same object name; this means that the thread has acquired twice the corresponding lock. For instance $Z_i = a \cdot a$. In this case $\widehat{Z}_i^t = (a, a)_t$, which is not a circular dependency – this is the reason why we index dependencies with thread names. Let $lock_i$ be special object name that is paired with the thread t . The lam \widehat{T}_i^Γ is the term

$$\widehat{T}_i^\Gamma = \&_{t \in T_i} \text{typeof}(\Gamma_i, t).run(\text{mk_tree}(\Gamma, t), t, lock_i)$$

where `run` is the method that is invoked by the Java instruction `t.start()`, where t is of type `thread`.

The function $names(\cdot)$. The function $names(i)$ takes an address i and returns a tuple of names whose length depends on the address.

This function is critical when methods are recursive because it allows us to keep the set \mathcal{T} finite. We remind that the type of a method has shape $(\bar{\tau}, t, a) \rightarrow (\nu \bar{a}') \langle \vartheta, \mathcal{T}, K, \vartheta, \ell \rangle$. The bound names \bar{a}' correspond to (renamings of) a subset of names returned by $names(\cdot)$ in the method's body – see rules for method definitions in Figure 4.5.

For example, consider a method that starts exactly a new thread t (without joining it) and, in a successive instruction at address i , invokes itself recursively (this happens for instance in method `BuildNetwork` of Figure 4.3). The set \mathcal{T} of this method will be a singleton x and $names(i)$ will return a name t' for instantiating x . If $t \neq t'$, we will get $T_{i+1} = \{t, t'\}$ and we will reach a `return` instruction with this same set. But then, by the rules for method definitions, the unique way to type the recursive method is that $t = t'$, henceforth this enforces $names(i) = t$. (Actually, our solution is more tricky than this because the new thread is a record type ϑ and we require that a same record type already exists in \mathcal{T}' .)

4.5 A behavioral type system for JVM_d

4.5.1 Typing rules

The type rules corresponding to the behavioral type system for JVM_d programs are reported in Figures 4.4 and 4.5. The most important type rules for the deadlock analysis are commented below.

The rule for `invokevirtual` computes the instance of $\text{BCT}(\mathbf{C.m})$ according to the arguments of the invocation. Let $\langle \vartheta, \mathcal{T}, K', \bar{\vartheta}, \ell \rangle$ be such instance. First of all we notice that the lam ℓ is not used in the rule. This is key for ensuring a compositional typing process, in other words, the type of a method does not depend on the behavior of other methods. Since `invokevirtual` is a sequential invocation, its effects are reported in the environment Γ_{i+1} that types the next instruction, notice the environment update operation. There are several effects: (i) the returned value, because it may be an object created by the method; (ii) the updates of the arguments of the invocation – the tuple $\bar{\vartheta}$ –; (iii) the threads that have been spawned by `C.m` and not synchronized – the set \mathcal{T} . We notice that the updates of elements in K' are not reported in Γ_{i+1} because they are already included in (ii). Let us comment on the premise $T_{i+1} = (T_i \setminus K') \cup \text{root}(\mathcal{T})$. The map T keeps track of the alive threads, whose invocations are reported in the lams – see the term $\widehat{T}_i^{\Gamma_i}$ in the conclusion. It is important to keep the set T_i as small as possible because this impacts on the precision of our analysis. Therefore we remove the threads that have been synchronized in `C.m` – the set K' – and we add those that have been created by the invocation – the set $\text{root}(\mathcal{T})$. The lam in the conclusion is the conjunction of the invocation to `C.m`, the invocations in $\widehat{T}_i^{\Gamma_i}$ and the dependencies of the sequence of locks in Z_i – the term \widehat{Z}_i^t .

Rules for `start` and `join` are also method invocations: the combination of `start` and `join` on the same thread actually corresponds to the rule for `invokevirtual`. The rule for `start` computes the instance of $\text{BCT}(\mathbf{C.run})$ where \mathbf{C} is the type of the thread t' on top of the stack S_i . Let ϑ be the effect of this method on t' (this records updates of fields, for example). The environment Γ_{i+1} is therefore the *merge* of ϑ with Γ_i because the new thread t' , being in parallel with the current one t , may produce effects at any point in the future. For similar reasons, Γ_{i+1} also contains the threads spawned by `C.run` – the set \mathcal{T} . Correspondingly, the set T_{i+1} is augmented with the thread that has been just spawned and those spawned by it. The rule for `join` updates the environment Γ_{i+1} with the whole effects of $\text{BCT}(\mathbf{C.run})$, that is the updates of ϑ and the threads \mathcal{T} that have been spawned by it. The maps T and K are modified in a similar way to `invokevirtual` (the set \mathcal{T} is

Method invocations

$$\begin{array}{l}
P[i] = \text{invokevirtual } \mathbf{C.m}(\mathbf{T}_0, \dots, \mathbf{T}_k) \quad i+1 \in \text{dom}(P) \quad S_i = \chi_k \cdots \chi_0 \cdot S' \\
\text{typeof}(\Gamma_i, \chi_0) = \mathbf{C} \quad \text{mk_tree}(\Gamma_i, (\chi_0, \dots, \chi_k)) = (\tau_0, \dots, \tau_k) \\
\bar{b} = \text{names}(i) \quad \text{BCT}(\mathbf{C.m})(\tau_0, \dots, \tau_k, t, [Z_i])(\bar{b}) = \langle \vartheta, \mathcal{T}, K', (\vartheta_0, \dots, \vartheta_k), \ell \rangle \\
\Gamma_{i+1} = \Gamma_i[\text{flat}(\vartheta) \oplus (\bigoplus_{i \in 0..k} \text{flat}(\vartheta_i)) \oplus \text{flat}(\mathcal{T})] \quad S_{i+1} = \text{root}(\vartheta) \cdot S' \\
F_{i+1} = F_i \quad Z_{i+1} = Z_i \quad K_{i+1} = K_i \cup K' \quad T_{i+1} = (T_i \setminus K') \cup \text{root}(\mathcal{T})
\end{array}$$

$$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t \ \& \ \mathbf{C.m}(\tau_0, \dots, \tau_k, t, [Z_i]) \rightarrow \vartheta$$

$$\begin{array}{l}
P[i] = \text{start } \mathbf{C} \quad i+1 \in \text{dom}(P) \quad S_i = t' \cdot S_{i+1} \\
\text{typeof}(\Gamma_i, t') = \mathbf{C} \quad \text{mk_tree}(\Gamma_i, t') = \tau \quad \bar{b} = \text{names}(i) \\
\text{BCT}(\mathbf{C.run})(\tau, t', \text{lock}_{t'}) (\bar{b}) = \langle \text{void}, \mathcal{T}, K', \vartheta, \ell \rangle \\
\Gamma_{i+1} = \Gamma_i \oplus \text{flat}(\vartheta) \oplus \text{flat}(\mathcal{T}) \quad F_i = F_{i+1} \quad Z_i = Z_{i+1} \\
K_{i+1} = K_i \quad T_{i+1} = T_i \cup \{t'\} \cup \text{root}(\mathcal{T})
\end{array}$$

$$\begin{array}{l}
P[i] = \text{join } \mathbf{C} \quad i+1 \in \text{dom}(P) \quad S_i = t' \cdot S_{i+1} \\
\text{typeof}(\Gamma_i, t') = \mathbf{C} \quad \text{mk_tree}(\Gamma_i, t') = \tau \quad \bar{b} = \text{names}(i) \\
\text{BCT}(\mathbf{C.run})(\tau, t', \text{lock}_{t'}) (\bar{b}) = \langle \text{void}, \mathcal{T}, K', \vartheta, \ell \rangle \quad \Gamma_{i+1} = \Gamma_i \\
F_i = F_{i+1} \quad Z_i = Z_{i+1} \quad K'' = K_i \cup K' \quad T'' = T_i \setminus K' \\
K_{i+1}, T_{i+1} = \begin{cases} K'', T'' \setminus \{t'\} & \text{if } t' \in T_i \\ K'' \cup \{t'\}, T'' & \text{if } t' \notin T_i \end{cases}
\end{array}$$

$$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t$$

$$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t \ \& \ ([Z_i], \text{lock}_{t'})_t$$

Locking instructions

$$\begin{array}{l}
P[i] = \text{monitorenter} \quad i+1 \in \text{dom}(P) \\
\Gamma_{i+1} = \Gamma_i \quad F_i = F_{i+1} \quad S_i = a \cdot S_{i+1} \\
K_i = K_{i+1} \quad T_i = T_{i+1} \quad Z_{i+1} = a \cdot Z_i
\end{array}$$

$$\begin{array}{l}
P[i] = \text{monitorexit} \quad i+1 \in \text{dom}(P) \\
\Gamma_{i+1} = \Gamma_i \quad F_i = F_{i+1} \quad S_i = a \cdot S_{i+1} \\
K_i = K_{i+1} \quad T_i = T_{i+1} \quad Z_{i+1} = Z_i \setminus a
\end{array}$$

$$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t$$

$$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t$$

Object manipulation instructions

$$\begin{array}{l}
P[i] = \text{new } \mathbf{C} \quad i+1 \in \text{dom}(P) \quad a = \text{names}(i) \quad \text{fields}(\mathbf{C}) = \bar{\mathbf{f}} \quad \Gamma_{i+1} = \Gamma_i[a \mapsto ([\bar{\mathbf{f}} : \bar{\mathbf{T}}], \mathbf{C})] \\
F_i = F_{i+1} \quad S_{i+1} = a \cdot S_i \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \quad T_i = T_{i+1}
\end{array}$$

$$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t$$

$$\begin{array}{l}
P[i] = \text{putfield } \mathbf{C.f} : \mathbf{T} \quad i+1 \in \text{dom}(P) \\
S_i = \chi \cdot a \cdot S_{i+1} \quad \Gamma_i(a) = [\mathbf{f}^h : \chi', \bar{\mathbf{f}}^h : \bar{\chi}'] \\
\Gamma_{i+1} = \Gamma_i[a \mapsto [\mathbf{f}^w : \chi, \bar{\mathbf{f}}^h : \bar{\chi}']] \\
F_i = F_{i+1} \quad Z_i = Z_{i+1} \quad T_i = T_{i+1} \quad K_i = K_{i+1}
\end{array}$$

$$\begin{array}{l}
P[i] = \text{getfield } \mathbf{C.f} : \mathbf{T} \quad i+1 \in \text{dom}(P) \\
S_i = a \cdot S' \quad \Gamma_i(a) = [\mathbf{f}^h : \chi, \bar{\mathbf{f}}^h : \bar{\chi}] \\
\Gamma_{i+1} = \Gamma_i[a \mapsto [\mathbf{f}^{hw} : \chi, \bar{\mathbf{f}}^h : \bar{\chi}]] \\
S_{i+1} = \chi \cdot S' \quad F_i = F_{i+1} \quad Z_i = Z_{i+1} \\
K_i = K_{i+1} \quad T_i = T_{i+1}
\end{array}$$

$$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t$$

$$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t$$

Figure 4.4: Type rules for JVM_{L_D} programs – Part I

Stack manipulation instructions

$P[i] = \text{inc} \quad i + 1 \in \text{dom}(P)$ $\frac{\Gamma_i = \Gamma_{i+1} \quad F_i = F_{i+1} \quad S_i = \text{int} \cdot S' = S_{i+1} \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \quad T_i = T_{i+1}}{\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t}$	
$P[i] = \text{pop} \quad i + 1 \in \text{dom}(P)$ $\frac{\Gamma_{i+1} = \Gamma_i \quad F_i = F_{i+1} \quad S_i = \chi \cdot S_{i+1} \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \quad T_i = T_{i+1}}{\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t}$	$P[i] = \text{push} \quad i + 1 \in \text{dom}(P)$ $\frac{\Gamma_i = \Gamma_{i+1} \quad F_i = F_{i+1} \quad \text{int} \cdot S_i = S_{i+1} \quad Z_i = Z_{i+1} \quad T_i = T_{i+1} \quad K_i = K_{i+1}}{\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t}$
$P[i] = \text{load } x \quad i + 1 \in \text{dom}(P)$ $\frac{\Gamma_i = \Gamma_{i+1} \quad F_i = F_{i+1} \quad S_{i+1} = F_i(x) \cdot S_i \quad K_i = K_{i+1} \quad T_i = T_{i+1} \quad Z_i = Z_{i+1}}{\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t}$	$P[i] = \text{store } x \quad i + 1 \in \text{dom}(P)$ $\frac{\Gamma_i = \Gamma_{i+1} \quad F_{i+1} = F_i[x \mapsto \chi] \quad S_i = \chi \cdot S_{i+1} \quad Z_i = Z_{i+1} \quad T_i = T_{i+1} \quad K_i = K_{i+1}}{\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t}$

Control flow instructions

$P[i] = \text{goto } L \quad L \in \text{dom}(P)$ $\frac{\Gamma_i = \Gamma_L \quad F_i = F_L \quad S_i = S_L \quad K_i = K_L \quad T_i = T_L \quad Z_i = Z_L}{\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t}$	$P[i] = \text{if } L \quad i + 1, L \in \text{dom}(P)$ $\frac{\Gamma_{i+1} = \Gamma_i = \Gamma_L \quad F_i = F_{i+1} = F_L \quad S_i = \text{int} \cdot S_{i+1} \quad S_{i+1} = S_L \quad K_i = K_{i+1} = K_L \quad T_i = T_{i+1} = T_L \quad Z_i = Z_{i+1} = Z_L}{\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t}$
---	--

Return

$P[i] = \text{return} \quad S_i = \chi \cdot S'$
$\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \widehat{T}_i^{\Gamma_i} \ \& \ \widehat{Z}_i^t ; (mk_tree(\Gamma_i, \chi), \emptyset, Z_i, T_i, K_i, \Gamma_i)$

Method definitions

$\text{BCT}(\mathbf{C.m}) = (\tau_0, \dots, \tau_k, t, a) \rightarrow (\nu \bar{a}') \langle \vartheta, \mathcal{T}, K', \bar{\vartheta}, \ell \rangle$ $F_1 = F_{\top}[0 \mapsto \text{root}(\tau_0), \dots, k \mapsto \text{root}(\tau_k)]$ $\Gamma_1 = \bigoplus_{i \in 0..k} flat(\tau_i)$ $S_1 = \varepsilon \quad Z_1 = a \quad T_1 = \emptyset \quad K_1 = \emptyset$ $\left(\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \ell_i ; \Psi_i \right)_{i \in \text{dom}(P)}$ $\ell = \sum_{i \in \text{dom}(P)} \ell_i \quad \bigsqcup_{i \in \text{dom}(P)} \Psi_i = (\vartheta, a, \mathcal{T}, K', \Gamma')$ $\bar{\vartheta} = mk_tree(\Gamma', (\text{root}(\tau_0), \dots, \text{root}(\tau_k)))$ $\bar{a}' = fn(\vartheta, \mathcal{T}, K', \bar{\vartheta}) \setminus fn(\tau_0, \dots, \tau_k, t, a)$ <hr/> $\text{BCT} \vdash \mathbf{T C.m} (\mathbf{T}_1, \dots, \mathbf{T}_k) P$	$\text{BCT}(\mathbf{C.m}) = (a' [\bar{\mathbf{f}} : \bar{\tau}], \tau_1, \dots, \tau_k, t, a) \rightarrow (\nu \bar{a}'') \langle \vartheta, \mathcal{T}, K', \bar{\vartheta}, \ell \rangle$ $F_1 = F_{\top}[0 \mapsto a', 1 \mapsto \text{root}(\tau_1), \dots, k \mapsto \text{root}(\tau_k)]$ $\Gamma_1 = flat(a' [\bar{\mathbf{f}} : \bar{\tau}]) \oplus \left(\bigoplus_{j \in 1..k} flat(\tau_j) \right)$ $S_1 = \varepsilon \quad Z_1 = a' \cdot a \quad T_1 = \emptyset \quad K_1 = \emptyset$ $\left(\text{BCT}, \Gamma, F, S, Z, T, K, i \vdash_t P : \ell_i ; \Psi_i \right)_{i \in \text{dom}(P)}$ $\ell = \sum_{i \in \text{dom}(P)} \ell_i \quad \bigsqcup_{i \in \text{dom}(P)} \Psi_i = (\vartheta, a' \cdot a, \mathcal{T}, K', \Gamma')$ $\bar{\vartheta} = mk_tree(\Gamma', (a', \text{root}(\tau_1), \dots, \text{root}(\tau_k)))$ $\bar{a}'' = fn(\vartheta, \mathcal{T}, K', \bar{\vartheta}) \setminus fn(a' [\bar{\mathbf{f}} : \bar{\tau}], \tau_1, \dots, \tau_k, t, a)$ <hr/> $\text{BCT} \vdash \text{synchronized } \mathbf{T C.m} (\mathbf{T}_0, \dots, \mathbf{T}_k) P$
---	---

Figure 4.5: Type rules for JVM_L_d programs – Part II

not considered because it has been added in the corresponding `start`. Here we also distinguish the two cases (a) the thread has been spawned locally – $t' \in T_i$ – and (b) the thread has been spawned by the caller – $t' \notin T_i$.

The rules for the locking instructions `monitorenter` and `monitorexit` update the map Z as expected. We notice that this update has an effect on the lam $\widehat{Z_{i+1}}^t$ by augmenting or reducing the dependencies, respectively.

There are two rules for method declarations `C.m`, depending on whether the method is `synchronized` or not. We discuss the latter idiom, the former one being similar. The rule verifies whether what is declared in `BCT(C.m)` does match with what is checked by the judgments of the instructions in its body. Assuming that t is the name of the current thread and a is the last lock it has acquired then we constrain the first instruction of `C.m` body to be typed with F_1 such that the first k local variables are set to the arguments of the invocation and Γ_1 defining such arguments. As regards S_1 , T_1 and K_1 , they are all empty, while $Z_1 = a$. The matching between `BCT(C.m)` and what is checked by the judgments is performed by collecting the tuples $(\chi, Z_i, T_i, K_i, \Gamma_i)$ of the `return` statements and merging the results. Notice that, if the body of `C.m` does not perform any concurrent operation (`start`, `join`, `monitorenter`, `monitorexit`) then every T_i is empty and every Z_i is a . Therefore, the lam of every instruction is always 0.

4.5.2 The abstract behavior of the Network class

Figure 4.6 details the behavioral type of the `Network` class in Figure 4.3. The types have been simplified for easing the readability: the actual JaDA types are more complex and verbose. Comments (in gray) explain the side effects of invocations, other comments (in green) correspond to the lines that are commented in Figure 4.3.

The behavior of `main` begins by calling the constructor of the class `Object`. Notice that, after such invocation, the structure of `x` and `y` is known. Then the type reports the invocation to `buildNetwork`.

The behavior of `takeLocks` is the parallel composition of two dependencies corresponding to the acquisition of the locks of `x` and `y`. Every dependency is formed by the last held lock and the current element. Notice that every method receives an extra argument corresponding to the last acquired lock at the moment of the invocation, in this case that argument is `u`. Dependencies are labeled with the identifier of the current thread, this identifier is also an extra argument, `t` in this case.

The behavior of `buildNetwork` has five states: (i) the invocation to `takeLocks`, (ii) the creation and initialization of the object `z`, (iii) the cre-

```

main(this,t,u) : T{thr} =                                --> returns unsync thread: thr
  Object.init(x,t,u) + Object.init(y,t,u) +              --> structure of x and y: x:x[], y:y[]
  //deadlock
  buildNetwork(this,int,x,x,t,u)                        --> creates unsync thread: thr
  //no deadlock
  //buildNetwork(this,int,x,y,t,u)

takeLocks(this,x,y,t,u) = (u,x)t & (x,y)t

buildNetwork(this,x,y|t,u) : T{thr} =                  --> returns unsync thread: thr
  takeLocks(this,x,y,t,u) +
  Object.init(z,t,u) +                                  --> z:z[]
  Network$1.init(thr,this,x,z,t,z) +                  --> thr:thr[this$0:this[], val$x:x[], val$z: z[]]
  Network$1.run(thr,thr,uthr) +
  Network$1.run(thr,thr,uthr) &                        --> recursive invocation goes in parallel
  buildNetwork(this,int,z,y,t,u)                      with the created thread

Object.init(this,t,u) : this[] = 0                    -->no side effects

Thread.init(this,t,u) : this[] = 0                   -->no side effects

Network$1.init(this,x1,x2,x3,t,u) : this[this$0 : x1,val$x : x2,val$z : x3] =
  Thread.init(this,t,u)

Network$1.run(this[this$0 : x1,val$x : x2,val$z : x3],t,u) =
  takeLocks(x1,x2,x3,t,u)                             --> run methods are treated like any other

```

Figure 4.6: BuildNetwork's lams

ation and initialization of the thread `thr`, (*iv*) the spawn of `thr`, (*v*) and the recursive invocation (in parallel with the spawned thread `thr`). The `buildNetwork` method also reports one spawned thread as side effect. This may appear contradictory (because `buildNetwork` spawns n threads). However, in this case JaDA is able to detect that `thr` is the only thread (from those created) that may be relevant (for the deadlock analysis) in an outer scope. This deduction is done considering the objects in the record structure of `thr`.

The constructors of `Object` and `Thread` have an empty behavior. On the contrary, the constructor of the class `Network$1` is more complex (`Network$1` is the name the JVM automatically gives to the anonymous `Thread` child class⁴ instantiated inside the method `buildNetwork` of the class `Network`). Being defined as an inner class, `Network$1` has access to the local variables in the scope in which it has been created, namely the variables `x`, `z` and the `this` reference to the container instance. The JVM addresses this by passing these variables to the constructor of the class and assigning them to internal fields, named also with the special character '\$' to avoid name collisions: `val$x`,

⁴<https://docs.oracle.com/javase/tutorial/java/javaOO/anonymousclasses.html>

`val$z` and `this$0`. Notice that the behavior of the constructor keeps track of two important things: the invocation to the constructor of the parent class `Thread.init` and the changes in the carrier object which goes from `this` to `this[this$0:x1, val$x:x2, val$z:x3]` where x_i are the formal arguments.

Finally, the behavior of the `run` method from the class `Network$1` contains only the invocation to the `takeLocks` method. Notice that `run` method assumes a certain structure from the carrier object. We remark that the body of this special method is typed exactly like any other method's body.

4.6 Analysis

Once behavioral types have been computed for the whole $JVML_d$ program, we can analyze the type of the `main` method. The analysis here proposed is based on the algorithm defined in [34, 48] that we briefly overview in this section. This section is complemented by Section 5.3, where it is reported the pseudo-code of the algorithm.

First of all, the semantics of lams is very simple: it amounts to unfolding method invocations. The critical points are that (i) every invocation may create new fresh names and (ii) the method definitions may be recursive. These two points imply that a lam model may be infinite state, which makes any analysis nontrivial. We recall that the states of lams are conjunctions ($\&$) of dependencies and method invocation (because types with disjunctions $+$ are modeled by sets of states with conjunctive dependencies).

The results of [34, 48] allow us to reduce to models of lams that are *finite*, i.e. finite disjunctions of finite conjunctions of dependencies. In turn, this finiteness makes possible to decide the presence of a so-called *circularity*, namely terms such as $(a, b)_t \& (b, a)_{t'}$. The reader is referred to [34, 48] for the details about the algorithm.

Actually, dependencies in [34, 48] are not indexed by thread names; here we use more informative dependencies in order to cope with `Java` reentrant locks. In particular $(a, b)_t \& (b, a)_t$ is not a circularity and, when $t \neq t'$, we carefully separate it from $(a, b)_t \& (b, a)_{t'}$. The main difference between the algorithm used in our tool and the one in [34, 48] is the definition of *transitive closure*, which is the base of the notion of circularity. Let $t \neq t'$ and let \checkmark be a special object name. Let also ℓ be a conjunction $\&$ of dependencies.

- The *transitive closure* of ℓ , noted ℓ^+ , is the least conjunction of either dependencies or function invocations that contains ℓ and such that if $(a, b)_t \& (b, c)_{t'}$ is a sub term of ℓ^+ then either (i) $(a, c)_{\checkmark}$ is a sub term of ℓ^+ , if $t \neq t'$, or (ii) $(a, c)_t$ is a sub term of ℓ^+ , if $t = t'$.

- We say that a lam has a *circularity* if it (or a lam equivalent to it when function invocations have been removed) has a sub term ℓ such that $(a, a)_{\checkmark}$ is in ℓ^+ .

For example $\ell = (a, b)_t \& (b, a)_t \& (b, c)_{t'}$ has no circularity because $\ell^+ = (a, b)_t \& (b, a)_t \& (a, a)_t \& (b, b)_t \& (b, c)_{t'} \& (a, c)_{\checkmark}$ does not contain any pair $(a, a)_{\checkmark}$ (this symbol \checkmark is a special thread name indicating that the dependency is due to the contributions of two or more threads).

Actually, the notion of transitivity that we use in JaDA is a refinement of the one described above. In particular, dependencies are labeled by the sequence of threads that contribute to them. These more informative labels allows us to compute transitive closures of terms like $(a, b)_t \& (b, c)_t \& (a, c)_{t'} \& (c, b)_{t'}$ in a more precise way. In fact, in this case, the mutual exclusion on the initial object a makes the concurrent presence of $(b, c)_t$ and $(c, b)_{t'}$ not possible.

In addition, using sequences of thread names as indexes of dependencies allows us to reconstruct, at least in part, the computation that produce the error. This is crucial for detecting false positives.

4.7 wait, notify and notifyAll methods

This section corresponds to the most recent progress in development of our solution. The current idea is still being analyzed and may not be exempt of mistakes.

We discuss a possible extension to cover the synchronization mechanism enforced by the `wait`, `notify` and `notifyAll` methods. These methods enable programmers to access the JVM's support for expressing thread coordination. They are public and final methods of the class `Object`, so they are inherited by all classes and cannot be modified. The invocations to `wait`, `notify` and `notifyAll` succeed for threads that already hold the lock of the object a on the stack. In this case, the `wait` instruction moves its own thread to the *wait set* of the object a and the object is relinquished by performing as many unlock operations as the integer stored in the lock field of a . The instructions `notify` and `notifyAll` respectively wake up one thread and all the threads in the wait set of a . These threads are re-enabled for thread scheduling, which means competing for acquiring the lock of a . The winner will lock the object a as many unmatched `monitorenter` it did on a before the `wait`-operation.

The *wait-notify* relation between threads can easily lead to deadlocks. There are two possible scenarios:

1. the `o.wait()` operation in t does not *happens-before* a matching `o.notify()` in t' ;
2. a lock on an object (different than `o`) held by t is blocking the execution of t' , thus preventing the invocation of `o.notify()` to happen.

The solution we are investigating uses *lams* with special couples $(a, a^w)_t$ and $(a, a^n)_t$ for representing the wait-notify dependencies. For example, for the wait operation, $(a, a^w)_t$ can be read as “thread t has the lock of `a` and has invoked method `wait` on it”. Similarly $(a, a^n)_t$ means “thread t has the lock of `a` and has invoked method `notify` on it”. However, these dependencies are not sufficient for deadlock analysis. In fact, the mere presence of a wait couple indicates a potential deadlock. For example a lam like

$$\ell = (a, a^w)_t \ \& \ (a, a^n)_{t'}$$

has to be considered as a deadlock, since nothing ensures that the notification will happen after the wait invocation (in this case it will depend on the scheduler). However, this would be a huge over-approximation. The precision can be enhanced if it is guaranteed that the wait operation *happens-before* the matching notification. For this purpose we keep track of one more type of lam couples: $(a, t')_t$ to be read as “thread t' was spawned by thread t while it was holding the lock on `a`”. Next, consider the following states:

$$\ell_1 = (a, a^w)_t \ \& \ (a, t')_t \ \& \ (a, a^n)_{t'} \qquad \ell_2 = (b, a)_t \ \& \ (a, a^w)_t \ \& \ (b, a)_{t'} \ \& \ (a, a^n)_{t'}$$

We can draw the following conclusions. First, ℓ_1 is not deadlocked, when t' started, t is holding the lock of `a`, so the notification will not occur until t releases the lock of `a`, which happens exactly when the `wait` is invoked. Second, ℓ_2 is deadlocked, the situation is similar to the previous one, but after releasing `a` the thread t still holds the lock on `b`. The thread t' , on the other hand, cannot grab the lock on `a` until it can grab the lock on `b`. This implies that the notification is never performed.

Following this analysis, we can define a strategy for deciding whether a wait couple is matched by a notify couple:

a lam that contains $(a, a^w)_t \ \& \ (a, a^n)_{t'}$ is safe whenever (i) t' is spawned by t while holding the lock on `a` and (ii) whenever `a` is the only object in common in the lock chains acquired by t and t' before the wait and the notify operations, respectively.

We are currently studying the soundness of this solution and we will report our results in a future work.

The bug8012326 As part of the experimentation process of JaDA we have analyzed a number of the deadlock-related bugs in the Oracle Bug Database ⁵. One of the analyzed errors is the bug number 8012326, described as: “Deadlock occurs when `Charset.availableCharsets()` is called by several threads at the same time”. Our analysis revealed that this deadlock happens after the invocation of one native function (that checks for the existing *char sets* in the OS). In particular, one execution path within this function suspends the calling thread while it holds locks on JVM objects that are blocking parallel threads. A clever solution would be to manually type the corresponding native function with a wait couple, in this way the deadlock could be predicted with the idea previously described. We remark the flexibility of our technique in allowing simple modifications that can solve new problems.

4.8 Type system extensions

In order to keep the presentation simple, the language $JVML_d$ and the theory developed in Sections 4.5 and 4.6 partially cover JVM. In this section we discuss how we address features as static fields, recursive data types, exception handling, inheritance and dynamic dispatch, native method calls, and alternative concurrency models. In particular we will discuss the problems we found and our solutions for extending the basic technique from the previous sections.

4.8.1 Static members

In Java it is possible to create fields and methods that belong to the class, rather than to an instance of the class. This is accomplished with the `static` modifier: fields and methods that have the static modifier, called *static fields* and *static methods*, respectively can be accessed from anywhere in the code, without creating an instance of the class. The bytecode retains ad-hoc operations for static members – `putstatic`, `getstatic` and `invokestatic`. In order to adequately deal with static members, we need (i) to extend the arguments in the method types to include the static fields accessed by the method and the corresponding effects, and (ii) to extend the environment Γ_i with definitions of static fields. Point (i) is not problematic, as regards (ii), we admit environments that are defined on class names and return record types containing static fields only. In this way we can address static fields as standard ones – see the first three rules in Figure 4.7.

⁵<http://bugs.java.com/>

Method invocations

$$\begin{array}{c}
P[i] = \text{invokestatic } \mathbf{C.m} (\mathbf{T}_1, \dots, \mathbf{T}_k) \quad i+1 \in \text{dom}(P) \quad S_i = \chi_k \cdots \chi_1 \cdot S' \\
mk_tree(\Gamma_i, (\mathbf{C}, \chi_1, \dots, \chi_k)) = (\tau_0, \tau_1, \dots, \tau_k) \quad \text{BCT}(\mathbf{C.m})(\tau_0, \dots, \tau_k, t, [Z_i]) = \langle \vartheta, \mathcal{T}', \mathcal{R}', K', (\vartheta_0, \vartheta_1, \dots, \vartheta_k), \ell \rangle \\
\Gamma_{i+1} = \Gamma_i[\text{flat}(\vartheta) \oplus (\bigoplus_{i \in 0..k} \text{flat}(\vartheta_i))] \quad S_{i+1} = \text{root}(\vartheta) \cdot S' \quad F_{i+1} = F_i \quad Z_{i+1} = Z_i \quad K_{i+1} = K_i \cup K' \\
\mathcal{T}_{i+1}, \mathcal{R}_{i+1} = \begin{cases} (\mathcal{T}_i \setminus K') \cup \mathcal{T}', \mathcal{R}_i \cup \mathcal{R}' & \text{if } i \text{ is not recursive and } (\mathcal{T}_i \cup \mathcal{R}_i) \cap (\mathcal{T}' \cup \mathcal{R}') = \emptyset \\ \mathcal{T}_i \setminus K', (\mathcal{R}_i \setminus K') \cup \mathcal{R}' \cup \mathcal{T}' & \text{otherwise} \end{cases} \\
\hline
\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i^{\Gamma_i} \& \widetilde{\mathcal{R}}_i \& \widehat{Z}_i^t \& \mathbf{C.m}(\tau_0, \dots, \tau_k, t, [Z_i]) \rightarrow \vartheta^\nabla \\
\\
P[i] = \text{putstatic } \mathbf{C.f} : \mathbf{T} \quad i+1 \in \text{dom}(P) & P[i] = \text{getstatic } \mathbf{C.f} : \mathbf{T} \quad i+1 \in \text{dom}(P) \\
S_i = \chi \cdot S_{i+1} \quad \Gamma_i(\mathbf{C}) = [\mathbf{f}^h : \chi', \overline{\mathbf{f}^h} : \chi'] & \Gamma_i(\mathbf{C}) = [\mathbf{f}^h : \chi, \overline{\mathbf{f}^h} : \chi] \quad \Gamma_{i+1} = \Gamma_i[\mathbf{C} \mapsto [\mathbf{f}^{\text{h}\cup\text{r}} : \chi, \overline{\mathbf{f}^h} : \chi]] \\
\Gamma_{i+1} = \Gamma_i[\mathbf{C} \mapsto [\mathbf{f}^w : \chi, \overline{\mathbf{f}^h} : \chi']] & S_{i+1} = \chi \cdot S_i \quad F_i = F_{i+1} \quad Z_i = Z_{i+1} \\
F_i = F_{i+1} \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \quad \mathcal{T}_i = \mathcal{T}_{i+1} \quad \mathcal{R}_i = \mathcal{R}_{i+1} & K_i = K_{i+1} \quad \mathcal{T}_i = \mathcal{T}_{i+1} \quad \mathcal{R}_i = \mathcal{R}_{i+1} \\
\hline
\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i^{\Gamma_i} \& \widetilde{\mathcal{R}}_i \& \widehat{Z}_i^t & \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i^{\Gamma_i} \& \widetilde{\mathcal{R}}_i \& \widehat{Z}_i^t \\
\\
P[i] = \text{invokestatic } \mathbf{C.m} (\mathbf{T}_1, \dots, \mathbf{T}_k) \quad i+1 \in \text{dom}(P) \quad S_i = \chi_k \cdots \chi_1 \cdot S' \quad \mathbf{C} \notin \text{dom}(\Gamma_i) \\
mk_tree(\Gamma_i, (\mathbf{C}_1, \dots, \mathbf{C}_h)) = (\tau'_1, \dots, \tau'_h) \quad \text{BCT}(\mathbf{C}.\langle \text{cinit} \rangle)(\tau'_1, \dots, \tau'_h, t, [Z_i]) = \langle _, \mathcal{T}', \mathcal{R}', K', \vartheta', \vartheta'_1, \dots, \vartheta'_h, \ell' \rangle \\
\Gamma' = \Gamma_i[\text{flat}(\vartheta') \oplus (\bigoplus_{i \in 0..h} \text{flat}(\vartheta'_i))] \quad mk_tree(\Gamma', (\mathbf{C}, \chi_1, \dots, \chi_k)) = (\tau_0, \tau_1, \dots, \tau_k) \\
\text{BCT}(\mathbf{C.m})(\tau_0, \dots, \tau_k, \text{root}(\vartheta'_1), \dots, \text{root}(\vartheta'_h), t, [Z_i]) = \langle \vartheta, \mathcal{T}'', \mathcal{R}'', K'', (\vartheta_0, \vartheta_1, \dots, \vartheta_k, \vartheta''_1, \dots, \vartheta''_h), \ell \rangle \\
\Gamma_{i+1} = \Gamma'[\text{flat}(\vartheta) \oplus (\bigoplus_{i \in 0..k} \text{flat}(\vartheta_i)) \oplus (\bigoplus_{i \in 1..h} \text{flat}(\vartheta''_i))] \\
S_{i+1} = \text{root}(\vartheta) \cdot S' \quad F_{i+1} = F_i \quad Z_{i+1} = Z_i \quad K_{i+1} = K_i \cup K' \cup K'' \\
\mathcal{T}_{i+1} = \mathcal{T}_i \setminus (K' \cup K'') \cup \mathcal{T}' \cup \mathcal{T}'' \quad \mathcal{R}_{i+1} = \mathcal{R}_i \setminus (K' \cup K'') \cup \mathcal{R}' \cup \mathcal{R}'' \\
\hline
\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i \& \widetilde{\mathcal{R}}_i \& \widehat{Z}_i^t \& (\mathbf{C}.\langle \text{cinit} \rangle(\tau'_1, \dots, \tau'_h, t, [Z_i]) \rightarrow _ + \mathbf{C.m}(\tau_0, \dots, \tau_k, t, [Z_i]) \rightarrow \vartheta^\nabla)
\end{array}$$

Figure 4.7: Type rules for static members

There are Java static members that may have subtle effects on deadlocks. These are the *static constructors* of classes. These constructors, called `<cinit>` in the bytecode, are executed by means of an implicit mechanism of the JVM in correspondence of the first access to the classes, *without any explicit instruction in the bytecode!* To address deadlocks caused by these methods, we assume that classes with static constructors *are not added* to the initial environment. This means that, if the `invokestatic`, `putstatic` and `getstatic` are the first operations accessing to the class \mathbf{C} , then no-one of the first three rules of Figure 4.7 can be applied because $\Gamma_i(\mathbf{C})$ is not defined and, therefore, $mk_tree(\Gamma_i, \mathbf{C})$ will fail. In this case, our system uses alternative rules: Figure 4.7 shows the rule for `invokestatic`. In this rule we are invoking a static method $\mathbf{C.m}$ in an environment Γ_i that has no binding for \mathbf{C} – namely the static constructor of \mathbf{C} has still not been run. This means that the `<cinit>` method of \mathbf{C} must be type-checked beforehand and the method $\mathbf{C.m}$ must be analyzed in the resulting environment. We notice

that the method type of $\mathbf{C}.<\mathbf{cinit}>$ is

$$(\tau'_1, \dots, \tau'_h, t, b) \rightarrow (\nu \bar{a})\langle -, \mathcal{T}', \mathcal{R}', K', \vartheta', \vartheta'_1, \dots, \vartheta'_h, \ell' \rangle$$

that is, in this case, the carrier of the method – the class \mathbf{C} – is not mentioned in the arguments, but the effects of the invocation are reported in the type ϑ' . We notice that the method $\mathbf{C}.<\mathbf{cinit}>$ has $h + 2$ arguments: the first h documents are classes with static fields that are accessed by this constructor. In the rule of Figure 4.7 we are assuming that the corresponding constructors of τ'_1, \dots, τ'_h have already been analyzed. Otherwise *we use a different type* for $\mathbf{C}.<\mathbf{cinit}>$ where the arguments contain exactly the classes whose static constructor has been already analyzed – therefore there are 2^h different types of $\mathbf{C}.<\mathbf{cinit}>$ and, consequently, 2^h different typing rules. The environment resulting from the static constructor is then used to type the invocation of $\mathbf{C}.m$ in a similar way to the first rule of Figure 4.7.

An unsound alternative to the case explosion generated by this solution is to assume that all static constructors are executed prior to the `main` method. This approach is unsound, because it does not actually resembles the semantics of the JVM. However, under the assumption that no concurrent operations take place in these constructors this solution does not affect the correctness of the results.

4.8.2 Exceptions and exception handlers

When an exception occurs during the execution of a program, the JVM removes from the stack all the elements, except the topmost one (the exception object, which is of a subclass of `Throwable`), and looks for a handler in the corresponding exception table. The exception table is a sequence of tuples (`from`, `to`, `target`, `type`), where the first three components are addresses and the last one is the type of the exception. So, if the exception occurred at j , the handler is defined by the first tuple $(i, i', i'', \mathbf{type})$ in the exception table such that $i \leq j < i'$ and the type of the exception matches `type`. If a tuple $(i, i', i'', \mathbf{type})$ is found, the control is given to the instruction at i'' ; otherwise the exception is re-thrown to the caller. If no handler is found for that exception, the program terminates (in an abrupt way).

In order to model exceptions at static time, we assume that method bodies P also include an exception table that is accessed by means of the notation $P.ET$. The table $P.ET$ takes in input a pair (i, \mathbf{C}) , namely the instruction index i and the type \mathbf{C} of the exception, and returns an instruction index k .

Because a method may return an exception object, we need to extend the behavioral class table, in order to cope with abrupt terminations. Therefore

BCT maps $\mathbf{C.m}(\bar{\tau}, t, a)$ to a tuple

$$(\nu \bar{a}') \langle \vartheta, \bar{\vartheta}_e, \mathcal{T}', \mathcal{R}', K', \bar{\vartheta}, \ell \rangle$$

where $\bar{\vartheta}_e$ is a finite sequence of record types that is linear with respect to the classes in sequence (pairwise different types have pairwise different subclasses of **Throwable**) that is used to select the right handler in the exception table. In case of abrupt termination, see rules in Figure 4.8 – the returned value ϑ is absent. We model the absence of ϑ with the symbol $-$. This requires an extension of the \sqcup operation in Section 4.6 for tuples $(\vartheta, Z_i, \mathcal{T}_i, \mathcal{R}_i, K_i, \Gamma_i)$:

$$(\vartheta, Z_i, \mathcal{T}_i, \mathcal{R}_i, K_i, \Gamma_i) \sqcup (-, Z_j, \mathcal{T}_j, \mathcal{R}_j, K_j, \Gamma_j) \stackrel{def}{=} (\vartheta, Z_i, \mathcal{T}_i \cup \mathcal{T}_j, \mathcal{R}_i \cup \mathcal{R}_j, K_i \cup K_j, \Gamma_i \oplus \Gamma_j)$$

The rule for **invokevirtual** in presence of exceptions is defined in Figure 4.8. The first six lines of this rule are similar to those of Figure 4.4, except for the return tuple of the invoked method that also contains $\bar{\vartheta}_e$. This tuple of record types is used to select the right handler according to the definitions in *P.ET*.

In particular, for every $(a[\overline{\mathbf{f}^h : \vartheta'}], \{\mathbf{C}_e\}) \in \bar{\vartheta}_e$, we select the index $P.ET(i, \mathbf{C}_e)$, if any. Let this index be h ; we type the instruction at h by constraining S_h to be a single element stack and K_h and \mathcal{T}_h to be over-approximations of those used for normal termination. Then K_h is taken to be equal to K_i – in case of abrupt termination we assume that no thread in input is joined by the called method – and, correspondingly, \mathcal{T}_h is $\mathcal{T}_i \cup \text{root}(\mathcal{T})$ – we assume that every thread created in case of normal termination has been already created. If the exception is not handled by the exception table then we report all the informations in Ψ in order to be lifted to the method type of P . Figure 4.8 also defines the rule for the static correctness of **athrow**. This rule constraints the state of the initial instruction of the corresponding handler, if any, otherwise it stores the record and its type in Ψ , in order to be returned to the caller.

Method invocations with exceptions and athrow

$$\begin{array}{l}
P[i] = \text{invokevirtual } \mathbf{C.m}(\mathbf{T}_0, \dots, \mathbf{T}_k) \quad i + 1 \in \text{dom}(P) \quad S_i = \chi_k \cdots \chi_0 \cdot S' \\
\text{typeof}(\Gamma_i, \chi_0) = \mathbf{C} \quad \text{mk_tree}(\Gamma_i, (\chi_0, \dots, \chi_k)) = (\tau_0, \dots, \tau_k) \\
\text{BCT}(\mathbf{C.m})(\tau_0, \dots, \tau_k, t, [Z_i]) = \langle \vartheta, \overline{\vartheta_e}, \mathcal{T}', \mathcal{R}', K', (\vartheta_0, \dots, \vartheta_k), \ell \rangle \\
\Gamma_{i+1} = \Gamma_i[\text{flat}(\vartheta) \oplus (\bigoplus_{i \in 0..k} \text{flat}(\vartheta_i))] \quad S_{i+1} = \text{root}(\vartheta) \cdot S' \quad F_{i+1} = F_i \\
Z_{i+1} = Z_i \quad K_{i+1} = K_i \cup K' \\
\mathcal{T}_{i+1}, \mathcal{R}_{i+1} = \begin{cases} (\mathcal{T}_i \setminus K') \cup \mathcal{T}', \mathcal{R}_i \cup \mathcal{R}' & \text{if } i \text{ is not recursive} \\ & \text{and } (\mathcal{T}_i \cup \mathcal{R}_i) \cap (\mathcal{T}' \cup K') = \emptyset \\ \mathcal{T}_i \setminus K', (\mathcal{R}_i \setminus K') \cup \mathcal{T}' \cup \mathcal{R}' & \text{otherwise} \end{cases} \\
\left(\begin{array}{l} P.\text{ET}(i, \mathbf{C}_e) = h \\ \text{implies} \\ \Gamma_h = \Gamma_{i+1}[\text{flat}((a[\overline{\mathbf{f}^h} : \vartheta']), \{\mathbf{C}_e\})] \quad S_h = a \\ F_h = F_i \quad Z_h = Z_i \quad K_h = K_i \quad \Psi_{\mathbf{C}_e} = \emptyset \\ \mathcal{T}_h, \mathcal{R}_h = \begin{cases} \mathcal{T}_i \cup \mathcal{T}', \mathcal{R}_i \cup \mathcal{R}' & \text{if } i \text{ is not recursive} \\ & \text{and } (\mathcal{T}_i \cup \mathcal{R}_i) \cap (\mathcal{T}' \cup K') = \emptyset \\ \mathcal{T}_i, \mathcal{R}_i \cup \mathcal{T}' \cup \mathcal{R}' & \text{otherwise} \end{cases} \\ \\ (i, \mathbf{C}_e) \notin \text{dom}(P.\text{ET}) \\ \text{implies} \\ \Gamma' = \Gamma_i[\text{flat}((a[\overline{\mathbf{f}^h} : \vartheta']), \{\mathbf{C}_e\})] \\ \Psi_{\mathbf{C}_e} = \begin{cases} (-, [\mathbf{C}_e \mapsto a], Z_i, \mathcal{T}_i \cup \mathcal{T}', \mathcal{R} \cup \mathcal{R}', K_i, \Gamma') & \text{if } i \text{ is not recursive} \\ & \text{and } (\mathcal{T}_i \cup \mathcal{R}_i) \cap (\mathcal{T}' \cup K') = \emptyset \\ (-, [\mathbf{C}_e \mapsto a], Z_i, \mathcal{T}_i, \mathcal{R} \cup \mathcal{T}' \cup \mathcal{R}', K_i, \Gamma') & \text{otherwise} \end{cases} \end{array} \right)^{(a[\overline{\mathbf{f}^h} : \vartheta'], \{\mathbf{C}_e\}) \in \overline{\vartheta_e}}
\end{array}$$

$$\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t \ \& \ \mathbf{C.m}(\tau_0, \dots, \tau_k, t, [Z_i]) \rightarrow \vartheta^\nabla, \bigsqcup_{\mathbf{C}_e \in \text{typeof}(\overline{\vartheta_e})} \Psi_{\mathbf{C}_e}$$

$$\begin{array}{l}
P[i] = \text{athrow} \quad S_i = \chi \cdot S' \quad \text{typeof}(\Gamma_i, \chi) = \mathbf{C} \\
\left(\begin{array}{l} P.\text{ET}(i, \mathbf{C}) = h \quad \text{implies} \quad \Gamma_h = \Gamma_i \quad S_h = \chi \quad F_h = F_i \quad Z_h = Z_i \\ K_h = K_i \quad \mathcal{T}_h = \mathcal{T}_i \quad \mathcal{R}_h = \mathcal{R}_i \quad \Psi = \emptyset \\ \\ (i, \mathbf{C}) \notin \text{dom}(P.\text{ET}) \quad \text{implies} \quad \Psi = (-, [\mathbf{C} \mapsto \chi], Z_i, \mathcal{T}_i, \mathcal{R}_i, K_i, \Gamma_i) \end{array} \right) \\
\hline
\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t, \Psi
\end{array}$$

Figure 4.8: The type rules for for managing exceptions

4.8.3 Arrays

In Java, arrays are objects, are dynamically created, and may be assigned to variables of type `Object`. In JVM there are two instructions for creating single-dimension arrays `newarray`, for arrays containing primitive data types, and `anewarray`, for arrays containing objects. (The operation `multinewarray` for creating multi-dimensional arrays is handled by our prototype in a similar way to the previous ones.) In this section we focus on single-dimensional arrays containing objects because the other ones are not problematic for deadlock analysis.

The typical pattern for creating arrays objects in Java, where the elements store pairwise different new objects (of class `Object`) is

```

L1: bipush 10
    anewarray java/lang/Object
    astore 1
L2: iconst_0
    istore 2
L3: iload 2
    bipush 9
    if_icmpgt L4
    aload 1
    iload 2
    new java/lang/Object
    dup
    invokespecial java/lang/Object.<init> ()
    aastore
    iload 2
    iinc
    istore 2
    goto L3
L4: return

```

Namely, an array of ten elements is created and the store operation initializing the element is performed in a cycle starting at L3.

Our technique does not identify integers; henceforth it is not possible to distinguish different indexes and the elements stored in the corresponding locations. Therefore, arrays are over-approximated by assuming that they *store the same element at every index*. Because of this, we are able to analyze those cases where the elements of the array are “compatible” (see below). For instance, our analysis covers arrays that are initialized with *pairwise different new objects*.

Arrays of class `C` are modeled by structures retaining single values, namely records ($[\text{idx} : \chi], [\mathbf{C}]$). Let $\text{Env}(\chi, \mathbf{C}, \overline{b}_\$)$ be an environment defining χ of type `C` and, whenever the value is a record, every name occurring in the fields of χ , and so on recursively. That is, the function $\text{Env}(\chi, \mathbf{C}, \overline{b}_\$)$ returns a *representative* for arrays’ values that stores in the possible fields either \top or `int` or a variable or an object name already defines in Γ_i (this environment is left implicit in the notation) or a name that is subscribed with $\$$. (Names indexed by $\$$ are managed in an ad-hoc way by the analyzer, see below.)

Let Eq be a mapping from names $a_\$$ to names. Let $\chi' \triangleright_{\Gamma_i}^{\text{Eq}} \chi$, say χ' is *Eq-compatible with χ* , if one of the following holds:

1. $\chi' = \chi$,

Array manipulation instructions

$$\begin{array}{c}
P[i] = \text{anewarray } \mathbf{C} \quad i + 1 \in \text{dom}(P) \\
(a, \overline{b}_s) = \text{names}(i) \quad \Gamma_{i+1} = \Gamma_i[a \mapsto ([\text{idx}^w : a_s], [\mathbf{C}])] + \text{Env}(a_s, \mathbf{C}, \overline{b}_s) \\
F_i = F_{i+1} \quad S_{i+1} = a \cdot S_i \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \quad \mathcal{T}_i = \mathcal{T}_{i+1} \quad \mathcal{R}_i = \mathcal{R}_{i+1} \\
\hline
\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t
\end{array}$$

$$\begin{array}{c}
P[i] = \text{aastore} \quad i + 1 \in \text{dom}(P) \\
S_i = \chi \cdot \text{int} \cdot a \cdot S_{i+1} \quad \text{there is Eq: } \chi \triangleright_{\Gamma_i}^{\text{Eq}} a_s \\
\Gamma_{i+1} = \Gamma_i[a \mapsto ([\text{idx}^w : a_s], [\mathbf{C}])][\chi \mapsto_{\Gamma_i} a_s] \\
F_i = F_{i+1} \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \\
\mathcal{T}_i = \mathcal{T}_{i+1} \quad \mathcal{R}_i = \mathcal{R}_{i+1} \\
\hline
\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i^{\Gamma_i} \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t
\end{array}
\qquad
\begin{array}{c}
P[i] = \text{aload} \quad i + 1 \in \text{dom}(P) \\
S_i = \text{int} \cdot a \cdot S' \quad \Gamma_i(a) = ([\text{idx}^h : a_s], [\mathbf{C}]) \\
\Gamma_{i+1} = \Gamma_i[a \mapsto ([\text{idx}^{h \cup r} : a_s], [\mathbf{C}])] \\
S_{i+1} = a_s \cdot S' \quad F_i = F_{i+1} \quad Z_i = Z_{i+1} \\
K_i = K_{i+1} \quad \mathcal{T}_i = \mathcal{T}_{i+1} \quad \mathcal{R}_i = \mathcal{R}_{i+1} \\
\hline
\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i^{\Gamma_i} \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t
\end{array}$$

$$\begin{array}{c}
P[i] = \text{putfield } \mathbf{C} \cdot \mathbf{f} : \mathbf{T} \quad i + 1 \in \text{dom}(P) \\
S_i = \chi \cdot a_s \cdot S_{i+1} \quad \Gamma_i(a_s) = ([\mathbf{f}^h : \chi', \overline{\mathbf{f}^h} : \chi'], \mathbf{C}) \quad \text{there is Eq: } \chi \triangleright_{\Gamma_i}^{\text{Eq}} \chi' \\
\Gamma_{i+1} = \Gamma_i[a_s \mapsto ([\mathbf{f}^w : \chi', \overline{\mathbf{f}^h} : \chi'], \mathbf{C})][\chi \mapsto_{\Gamma_i} \chi'] \\
F_i = F_{i+1} \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \quad \mathcal{T}_i = \mathcal{T}_{i+1} \quad \mathcal{R}_i = \mathcal{R}_{i+1} \\
\hline
\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i^{\Gamma_i} \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t
\end{array}$$

Figure 4.9: Type rules for array related instructions

2. $\chi' = \text{Eq}(\chi)$,
3. $\chi' = \text{null}$,
4. $\Gamma_i(\chi') = [\dots \mathbf{f}_j^{\mathbf{h}'_j} : \chi'_j \dots]$ and $\Gamma_i(\chi) = [\dots \mathbf{f}_j^{\mathbf{h}_j} : \chi_j \dots]$ and, for every j , $\mathbf{h}'_j \leq \mathbf{h}_j$ and $\chi'_j \triangleright_{\Gamma_i}^{\text{Eq}} \chi_j$.

Whenever there is Eq such that $\chi' \triangleright_{\Gamma_i}^{\text{Eq}} \chi$ we also define $[\chi' \mapsto_{\Gamma_i} \chi]$ as follows:

- it is \emptyset whenever $\chi' = \text{int} = \chi$ or $\chi' = \top = \chi$ or $\Gamma_i(\chi') = \Gamma_i(\chi)$,
- it is $\Gamma_i(\chi)$ whenever $\Gamma_i(\chi') = \text{null}$,
- it is $[\mathbf{f}_1^w : \chi_1, \dots, \mathbf{f}_n^w : \chi_n][(\chi'_j \mapsto_{\Gamma_i} \chi_j)^{j \in 1..n}]$ whenever $\Gamma_i(\chi') = [\mathbf{f}_1^{\mathbf{h}'_1} : \chi'_1, \dots, \mathbf{f}_n^{\mathbf{h}'_n} : \chi'_n]$ and $\Gamma_i(\chi) = [\mathbf{f}_1^{\mathbf{h}_1} : \chi_1, \dots, \mathbf{f}_n^{\mathbf{h}_n} : \chi_n]$.

The JVMML rules for manipulating single-dimensional arrays storing objects are defined in Figure 4.9.

The rule for `anewarray` defines *at the same time* a fresh name for the array, called a , the corresponding name a_s for the field value and a *representative* that defines its fields. The rule for `aastore` checks whether the object

χ to be stored in a is actually in the relation $a_{\$} \triangleright_{\Gamma_i}^{\text{Eq}} \chi$, for some **Eq**. This means that $a_{\$}$ is a representative of χ and this last object can be safely dismissed. The resulting environment Γ_{i+1} records that the array has been accessed in writing (**w**) and *updates the names in $\Gamma_i(\chi)$ with the corresponding values in $\Gamma_i(a_{\$})$* . The instruction **aastore** has a catastrophic side-effect on the argument χ . In fact, since names in $\Gamma_i(a_{\$})$, become aliases of the corresponding names in $\Gamma_i(\chi)$, these aliases may modify the latter ones in an unpredictable way. Therefore, we need to *update* all the tree structures of $\Gamma_i(\chi)$ – *c.f.* the operation $[\chi \mapsto_{\Gamma_i} a_{\$}]$. A similar situation occurs when one wants to modify a field of name $a_{\$}$. For this reason we also report in Figure 4.9 the rule of the **putfield** instruction when the name of the object to modify is indexed by a $\$$. The rule is different from the **putfield** in Figure 4.4.

There is another subtlety to examine before discussing the management of $\$$ -indexed names in our algorithm for circularities in lams. This subtlety is about the merge operation \oplus in Section 4.6 is made more restrictive when record types are arrays. Let

$$\mathbf{f}^{\mathbf{h}} : \chi \oplus_{\text{aarr}} \mathbf{f}^{\mathbf{h}'} : \chi' \stackrel{\text{def}}{=} \begin{cases} \mathbf{f}^{\mathbf{h} \sqcup \mathbf{h}'} : \chi & \text{if } \chi = \chi' \\ \mathbf{f}^{\mathbf{w}} : \top & \text{otherwise} \end{cases}$$

Then, the operation $\mathbf{f}^{\mathbf{h}} : \chi \oplus \mathbf{f}^{\mathbf{h}'} : \chi'$ is refined into $\mathbf{f}^{\mathbf{h}} : \chi \oplus_{\text{aarr}} \mathbf{f}^{\mathbf{h}'} : \chi'$ whenever either χ or χ' are $\$$ -indexed names. In particular:

$$([\text{idx}^{\mathbf{h}} : \chi], [\mathbf{C}]) \oplus ([\text{idx}^{\mathbf{h}'} : \chi'], [\mathbf{C}]) \stackrel{\text{def}}{=} ([\text{idx}^{\mathbf{h}} : \chi \oplus_{\text{aarr}} \text{idx}^{\mathbf{h}'} : \chi'], [\mathbf{C}])$$

Therefore, the access in parallel of two threads to a same array of objects is admitted only when the *two threads access to the same representative*. It is worth to notice that, representatives act as place holders – they may be instantiated by whatever name.

In fact, these $\$$ -indexed names are managed in ad-hoc ways by the analyzer. In particular, a term like $(b, c)_t \& (c, a_{\$})_{t'}$ is a circularity because $a_{\$}$ may be replaced by *every name of the same type*, including b , if its type is the same. Similarly, $(a_{\$}, b)_t \& (b, c)_t \& (a_{\$}, c)_{t'} \& (c, b)_{t'}$ is a circularity because the dependencies starting with $a_{\$}$ might concern different objects.

4.8.4 Recursive data types

Recursive data types are a standard feature of programming languages that allows programmers to define and use lists, trees, etc. These types are finite but not (statically) bounded and, as we did with arrays, we use finite representations.

A representative of a recursive type value is built by unfolding the recursive types (exactly) up to those nodes containing a name of a class already present in the tree. Nodes inside the tree are labeled by new names, nodes in the leaves are labeled either (for non recursive types) with \top or `int` or with names already present in the environment or (for recursive types) with names subscribed by a $\$,$ such as $r_\$,$ where r is the label of the node of the class that is already present in the tree. By construction, these structures are finite. Names $r_\$$ have the same structure of r except that every name of a recursive type is indexed by $\$$ (and the access tag is set to `w`). For instance, if C is a class whose type is `[val : int, next : C]` then, in correspondence of a `new C` instruction, we produce an environment

$$r \mapsto [\text{val}^w : \text{int}, \text{next}^w : r_\$], \quad r_\$ \mapsto [\text{val}^w : \text{int}, \text{next}^w : r_\$]$$

(the access tags are all `w` because we are initializing the values).

We use the predicate “ C is recursive” when C has a field whose values contain other values of class C or a subclass of its. Rules of `new` and `putfield` in Figure 4.4 are rewritten as follows when the class is recursive:

$$\frac{\begin{array}{l} P[i] = \text{new } C \quad C \text{ is recursive} \quad i+1 \in \text{dom}(P) \\ (a, \overline{a_\$}) = \text{names}(i) \quad \Gamma_{i+1} = \Gamma_i \oplus \mathbf{Env}(\chi, C, \overline{a_\$}) \\ F_i = F_{i+1} \quad S_{i+1} = \chi \cdot S_i \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \quad \mathcal{T}_i = \mathcal{T}_{i+1} \quad \mathcal{R}_i = \mathcal{R}_{i+1} \end{array}}{\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t}$$

$$\frac{\begin{array}{l} P[i] = \text{putfield } C.f : T \quad C \text{ is recursive} \quad i+1 \in \text{dom}(P) \\ S_i = \chi' \cdot \chi \cdot S_{i+1} \quad \Gamma_i(\chi) = ([\mathbf{f}^h : \chi_f, \overline{\mathbf{f}^h} : \chi_f], C) \quad \text{there is Eq: } \chi' \triangleright_{\Gamma_i}^{\text{Eq}} \chi_f \\ \Gamma_{i+1} = \Gamma_i[\chi \mapsto ([\mathbf{f}^w : \chi_f, \overline{\mathbf{f}^h} : \chi_f], C)][\chi' \mapsto_{\Gamma_i} \chi_f] \\ F_i = F_{i+1} \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \quad \mathcal{T}_i = \mathcal{T}_{i+1} \quad \mathcal{R}_i = \mathcal{R}_{i+1} \end{array}}{\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i^{\Gamma_i} \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t}$$

We omit any comment because the arguments are similar to those for arrays.

The presence of recursive types makes the definitions of the functions $\text{flat}(\cdot)$ and $\text{mk_tree}(\cdot)$ in Section 4.4 inadequate. Indeed, in the full language, the invocations to $\text{flat}(\cdot)$ and $\text{mk_tree}(\cdot)$ in the rules of Figures 4.4, 4.5, 4.7, and 4.8 are replaced by $\text{flat}_\emptyset(\cdot)$ and $\text{mk_tree}_\emptyset(\cdot)$, which are defined below:

$$\begin{aligned}
flat_A(\vartheta) &= \begin{cases} \emptyset & \text{if } \vartheta \in \{\top, \mathbf{int}, X\} \\ \emptyset & \text{if } root(\vartheta) = a_{\S} \text{ and } a_{\S} \in A \\ flat_{(A \setminus a) \cup \{a_{\S}\}}(\vartheta) & \text{if } root(\vartheta) = a_{\S} \text{ and } a \in A \\ a \mapsto (\overline{[\mathbf{f}^h : root(\vartheta)]}, \mathbf{C}) \oplus (\bigoplus_{\vartheta' \in \bar{\vartheta}} flat_A(\vartheta')) & \text{if } \vartheta = (a \overline{[\mathbf{f}^h : \vartheta]}, \mathbf{C}) \\ & \text{and } \mathbf{C} \text{ is not recursive} \\ a \mapsto (\overline{[\mathbf{f}^h : root(\vartheta)]}, \mathbf{C}) \oplus (\bigoplus_{\vartheta' \in \bar{\vartheta}} flat_{A \cup \{a\}}(\vartheta')) & \text{if } \vartheta = (a \overline{[\mathbf{f}^h : \vartheta]}, \mathbf{C}) \\ & \text{and } \mathbf{C} \text{ is recursive and } a \notin A \end{cases} \\
mk_tree_A(\Gamma, \chi) &= \begin{cases} \chi & \text{if } \chi \in \{\top, \mathbf{int}, X\} \text{ or } (\chi = a_{\S} \text{ and } a_{\S} \in A) \\ (a \overline{[\mathbf{f} : mk_tree_A(\Gamma, \chi)]}, \mathbf{C}) & \text{if } \chi = a \text{ and } \Gamma(a) = (\overline{[\mathbf{f}^h : \chi]}, \mathbf{C}) \\ & \text{and } \mathbf{C} \text{ is not recursive} \\ (a \overline{[\mathbf{f} : mk_tree_{A \cup \{a\}}(\Gamma, \chi)]}, \mathbf{C}) & \text{if } \chi = a \text{ and } \Gamma(a) = (\overline{[\mathbf{f}^h : \chi]}, \mathbf{C}) \\ & \text{and } a \notin A \text{ and } \mathbf{C} \text{ is recursive} \\ (a_{\S} \overline{[\mathbf{f} : mk_tree_{(A \setminus a) \cup \{a_{\S}\}}(\Gamma, \chi)]}, \mathbf{C}) & \text{if } \chi = a_{\S} \text{ and } \Gamma(a) = (\overline{[\mathbf{f}^h : \chi]}, \mathbf{C}) \\ & \text{and } a \in A \text{ and } \mathbf{C} \text{ is recursive} \end{cases}
\end{aligned}$$

4.8.5 Inheritance

JVML_d does not admit to derive classes from other classes. As a consequence, when a method is invoked, it is possible to uniquely locate the method definition. This feature is expressed in the type rule for `invokevirtual` in Figure 4.4 by the constraint that $typeof(\Gamma_i, \chi_0)$ is always a single element.

On the contrary, Java, as many object-oriented programming languages, admits subclasses, thereby allowing redefinitions of methods in the subclasses. In this context, locating a method definition is performed at run-time through a mechanism called dynamic dispatch that uses the actual value of the called object. Clearly, this mechanism has a significant impact on deadlock analysis because dispatching an invocation to a definition rather than another may change the overall result. It is also clear that, since our approach is static, our analyzer cannot resolve the dynamic dispatch in a precise way.

The type system discussed in this solution is over-constraining on the relations between consecutive environments Γ_i and Γ_{i+1} (or Γ_L , in case of jumps) because it imposes an equality relation between them. For instance, let D and E be subclasses of C, which has a method `foo`, which is overridden both in D and E. and consider the following code

```

C w;

if (z)
  w = new D x;
else
  w = new E y;

w.foo();

```

where the invocation `w.foo()` will be dispatched to the method `D.foo` or `E.foo`, according to the value of `z`. This code does not type with the system in Figures 4.4 and 4.5 because the classes of `x` and `y` are different (assuming that the function $names(\cdot)$ at instructions corresponding to `w = x;` and to `w = y;` returns the same name). There is a standard solution to this problem. This solution relies on subtypes and relaxes the equality constraint of types of successor instructions in Figures 4.4 and 4.5. This solution is described in more detail in Section 5.2 as it is closely related to the inference process. Here we focus on the case of inheritance, examining the solution for programs like the foregoing one.

Let $fields(C)$ be the refinement of the function in Section 4.5 that returns a pair of tuples of field names \bar{f}, \bar{g} , where \bar{f} are the field names of the class C and \bar{g} are the names *defined* in one of C_1, \dots, C_h (and not inherited). Let \perp be a new type value that means “*undefined*” (while \top is the greatest value, \perp is the least one). In our prototype, whenever an object is created, the corresponding name is bound to a record that contains the fields of the class *plus every field defined in its subclasses* (and not inherited). In particular, the rule that we use for `new C` is

$$\frac{
\begin{array}{l}
P[i] = \text{new } C \quad i+1 \in dom(P) \quad a = names(i) \quad fields(C) = \bar{f}, \bar{g} \\
\Gamma_i[a \mapsto ([\bar{f} : \top, \bar{g} : \perp], C)] <: \Gamma_{i+1} \\
F_i <: F_{i+1} \quad a \cdot S_i <: S_{i+1} \quad Z_i = Z_{i+1} \quad K_i = K_{i+1} \quad \mathcal{T}_i = \mathcal{T}_{i+1} \quad \mathcal{R}_i = \mathcal{R}_{i+1}
\end{array}
}{
\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i^{\Gamma_i} \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t
}$$

where fields of class C are set to \top while the other ones are set to \perp . We notice that the constraint between environments, fields and stacks is “ $<:$ ” rather than “ $=$ ”. This $<:$ relation makes it possible the analysis of programs like the above one. For example, if C has exactly one field f , D extends C with the field g and E extends C with the field g' , w at `w.foo()` is typed with $([f : \top, g : \top, g' : \top], C \cdot D \cdot E)$, where $C \cdot D \cdot E$ denotes a set of classes.

The rule for `invokevirtual` that deals with dynamic dispatch on a set

of classes $C_1 \cdots C_h$ is (we are using again “=” rather than “<”):

$$\begin{aligned}
P[i] &= \text{invokevirtual } C_{j.m}(\tau_0, \dots, \tau_k) \quad i+1 \in \text{dom}(P) \quad S_i = \chi_k \cdots \chi_0 \cdot S' \\
&\text{typeof}(\Gamma_i, \chi_0) = \{C_1, \dots, C_h\} \quad \text{mk_tree}(\Gamma_i, (\chi_0, \dots, \chi_k)) = (\tau_0, \dots, \tau_k) \\
&\left(\text{BCT}(C_{j.m})(\tau_0, \dots, \tau_k, t, [Z_i]) = \langle \vartheta^j, \mathcal{T}^j, \mathcal{R}^j, K^j, (\vartheta_0^j, \dots, \vartheta_k^j), \ell^j \rangle \right)^{j \in 1..h} \\
&\mathcal{T}' = \bigcup_{j \in 1..h} \mathcal{T}^j \quad \mathcal{R}' = \bigcup_{j \in 1..h} \mathcal{R}^j \quad K' = \bigcap_{j \in 1..h} K^j \\
\vartheta' &= \vartheta^1 \oplus \dots \oplus \vartheta^h \quad (\vartheta_0', \dots, \vartheta_k') = (\vartheta_0^1, \dots, \vartheta_k^1) \oplus \dots \oplus (\vartheta_0^h, \dots, \vartheta_k^h) \\
\Gamma_{i+1} &= \Gamma_i[\text{flat}(\vartheta') \oplus (\bigoplus_{i \in 0..k} \text{flat}(\vartheta_i'))] \quad S_{i+1} = \text{root}(\vartheta') \cdot S' \\
F_{i+1} &= F_i \quad Z_{i+1} = Z_i \quad K_{i+1} = K_i \cup K' \\
\mathcal{T}_{i+1}, \mathcal{R}_{i+1} &= \begin{cases} (\mathcal{T}_i \setminus K') \cup \mathcal{T}', \mathcal{R}_i \cup \mathcal{R}' & \text{if } i \text{ is not recursive} \\ & \text{and } (\mathcal{T}_i \cup \mathcal{R}_i) \cap (\mathcal{T}' \cup \mathcal{R}') = \emptyset \\ \mathcal{T}_i \setminus K', (\mathcal{R}_i \setminus K') \cup \mathcal{T}' \cup \mathcal{R}' & \text{otherwise} \end{cases}
\end{aligned}$$

$$\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P : \widehat{\mathcal{T}}_i \ \& \ \widetilde{\mathcal{R}}_i \ \& \ \widehat{Z}_i^t \ \& \ \sum_{j \in 1..h} C_{j.m}(\tau_0, \dots, \tau_k, t, [Z_i]) \rightarrow (\vartheta^j)^\nabla$$

In particular, for every class in the type of the carrier, we compute the corresponding lam and the overall type is the disjunction of the possible invocations – this is our modeling of *overridden methods* (we are assuming that a method that is not present in a superclass has empty behavior). It is worth to observe that the *cumulative effect* of all the overridden methods will impact on the judgment for the next instruction – see the definitions of Γ_{i+1} , F_{i+1} , S_{i+1} , K_{i+1} , \mathcal{T}_{i+1} , and \mathcal{R}_{i+1} .

The above rule assumes that the method m is *not overloaded* in any class. In case of overloading, we also need to consider all the possible instances caused by the types of the arguments.

Interfaces. In Java it is possible to define abstract classes, called *interfaces*, whose methods have empty behaviors. The language constrains the programmer to implement these interfaces by means of classes. We deal with interfaces as they were standard classes and with their implementations as they were subclasses of them. Henceforth `invokeinterface` instructions are typed as the `invokevirtual` instruction in the above rule.

4.8.6 Alternative concurrent models and native methods

During its evolution Java has been bringing different mechanisms to write concurrent programs. The current proposal targets the (classic way of) concurrency expressed by `Thread` objects and `synchronized` blocks (the `wait-notify` mechanism is not covered by this work and by our current prototype – see Section 5.4 for a discussion of a possible solution). There are as well many third party components that allow abstracting threads manipulation

in general. Our approach allows the user to manually instruct the analysis to deal with these scenarios. This is done by the addition of manual entries in the BCT, in this way method declarations with specific behaviors are not sent through the inference algorithm.

This approach makes also possible to deal with method implementations that do not comply with the restrictions described in this solution, such as native methods.

4.9 Full proof of correctness

In this section we develop the formal proof of the correctness of the type system presented in Section 4.5 and of its corresponding analysis presented in Section 4.6. The first epigraph of this section presents a short less formal overview of correctness demonstration. The full technical details of the demonstration can be found in the rest of the section.

4.9.1 Correctness overview

The first part of the proof addresses the soundness of the type system in Section 4.6. This requires

1. an operational semantics of JVMML; we consider the one defined in Section 4.9.2;
2. an extension of the type system to handle JVM runtime configurations, defined in Section 4.9.6.

Then, as usual with type systems, the soundness is represented by a subject reduction theorem expressing that, if a JVM configuration cn has lam ℓ and cn reduces to a configuration cn' then (i) cn' is also well-typed and (ii) if ℓ' is the type of cn' then ℓ and ℓ' are in a relation called *later stage* (Section 4.9.5) and noted $\ell \succcurlyeq \ell'$.

The second part of the proof is about the correctness of the later stage relation with respect to the deadlock analysis. We first demonstrate that the lam of a deadlocked configuration always contains a *circular dependency* between names. Then we demonstrate that, if two types ℓ and ℓ' are in the later stage relationship, namely $\ell \succcurlyeq_{\text{BCT}} \ell'$, with respect to a given BCT, and ℓ has no circular dependency then also ℓ' has no circular dependency. The details of this proof can be found in Section 4.9.7.

As a byproduct of the above results we get that, if the lam of a JVMML program has no circularity then the JVMML program is deadlock-free. It is worth

to notice that our lams have infinite-state models. Nevertheless, in [34] has been demonstrated the decidability of whether the type will manifest a circular dependency in one of its states or not, here we extend that demonstration in Section 4.9.8.

4.9.2 The JVML formal semantics

In this section we present the operational semantics of JVML_d . Each instruction performs a transformation of machine states, that are configurations

$$P \Vdash_H \langle (pc^{\mathbf{C}^m}, f_1, s_1) \cdot \varphi_1, z_1 \rangle_{t_1} \cdots \langle (pc^{\mathbf{D}^m}, f_n, s_n) \cdot \varphi_n, z_n \rangle_{t_n} \dagger W$$

where each tuple $\langle \varphi, z \rangle_t$ represents an active thread t , such that:

- φ is the stack of activation records (or *frames*), where
 - pc is the *program counter* which contains the address of the instruction to be executed
 - f is a total map from the set of local variables to the set of values
 - s is the *operand Stack* of values
- z is the set of objects locked by the thread,

and W is the set of waiting threads (due to joins).

The *Heap* contains

- all objects created by `new`. Each object contains:
 - fields: the value of field a of object o is accessed with $H(o).a$. The update is noted $H(o).a \mapsto v$.
 - the class of the object, returned by $H(o).class$
 - a *wait set* containing the threads which are waiting for a thread to finish (namely they invoked `join` on o).
 - a counter F which tracks the number of locks acquired by a thread on the object

Moreover we use $H(o) \mapsto (\rho_{\perp}^{\mathbf{C}}, \mathbf{C})$ to allocate heap space for the object o of class \mathbf{C} , by assigning a default value to the fields.

The *waiting set* $W = \{W_{o_1}, \dots, W_{o_n}\}$, where W_{o_i} is the set of threads that performed a `join` on the object (of class `Thread`) o_i .

The rules are presented in Figures 4.10 and 4.11.

$$\begin{array}{c}
\text{(S-POP)} \\
\frac{P[pc^{c.m}] = \text{pop}}{P \Vdash_H \langle (pc^{c.m}, f, v \cdot s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (pc^{c.m} + 1, f, s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-INC)} \\
\frac{P[pc^{c.m}] = \text{inc}}{P \Vdash_H \langle (pc^{c.m}, f, n \cdot s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (pc^{c.m} + 1, f, (n + 1) \cdot s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-STORE)} \\
\frac{P[pc^{c.m}] = \text{store } x}{P \Vdash_H \langle (pc^{c.m}, f, v \cdot s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (pc^{c.m} + 1, f[x \mapsto v], s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-PUTFIELD)} \\
\frac{P[pc^{c.m}] = \text{putfield } D.a \ \tau \quad H' = H(o).a \mapsto e}{P \Vdash_H \langle (pc^{c.m}, f, e \cdot o \cdot s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_{H'} \langle (pc^{c.m} + 1, f, s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-NEW)} \\
\frac{P[pc^{c.m}] = \text{new } D \quad o \notin \text{dom}(H) \quad H' = H[o \mapsto (\rho^D, D)]}{P \Vdash_H \langle (pc^{c.m}, f, s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_{H'} \langle (pc^{c.m} + 1, f, o \cdot s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-IF-TRUE)} \\
\frac{P[pc^{c.m}] = \text{if } L^{c.m} \quad n \neq 0}{P \Vdash_H \langle (pc^{c.m}, f, n \cdot s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (L^{c.m}, f, s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-IF-FALSE)} \\
\frac{P[pc^{c.m}] = \text{if } L^{c.m}}{P \Vdash_H \langle (pc^{c.m}, f, 0 \cdot s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (pc^{c.m} + 1, f, s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-PUSH)} \\
\frac{P[pc^{c.m}] = \text{push}}{P \Vdash_H \langle (pc^{c.m}, f, s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (pc^{c.m} + 1, f, 0 \cdot s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-LOAD)} \\
\frac{P[pc^{c.m}] = \text{load } x}{P \Vdash_H \langle (pc^{c.m}, f, s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (pc^{c.m} + 1, f, f(x) \cdot s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-GOTO)} \\
\frac{P[pc^{c.m}] = \text{goto } L^{c.m}}{P \Vdash_H \langle (pc^{c.m}, f, s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (L^{c.m}, f, s) \cdot \varphi, z \rangle_t} \\
\\
\text{(S-GETFIELD)} \\
\frac{P[pc^{c.m}] = \text{getfield } D.a \ \tau \quad H(o).a = e}{P \Vdash_H \langle (pc^{c.m}, f, o \cdot s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (pc^{c.m} + 1, f, e \cdot s) \cdot \varphi, z \rangle_t}
\end{array}$$

Figure 4.10: Reduction rules: Basic operators

4.9.3 Lam semantics

Let \mathcal{A} be the (infinite) set of *object names*, ranged over by a, b, c, \dots . In the syntax of ℓ , see Section 4.4, the operations “&” and “+” are associative, commutative with 0 being the identity on &, and definitions and lams are equal up-to alpha renaming of bound names. Namely, if $a \notin \text{var}(\ell)$, the

$\frac{\text{(S-INVK)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{invokevirtual D.m}'(\text{T}_1, \dots, \text{T}_n) \\ \text{synchronized} \notin \text{mod}(\text{D.m}) \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f_1, v_n \cdots v_1 \cdot o \cdot s_1 \rangle \cdot \varphi, z \rangle_t \rightarrow \\ \Vdash_H \langle \langle \text{1}^{\text{D.m}'}, f_2[0 \mapsto o, 1 \mapsto v_1, \dots, n \mapsto v_n], \epsilon \rangle \cdot \\ (\text{pc}^{\text{C.m}} + 1, f_1, \bullet \cdot s_1) \cdot \varphi, z \rangle_t}$	$\frac{\text{(S-INVK-SYNCH-0)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{invokevirtual D.m}'(\text{T}_1, \dots, \text{T}_n) \\ \text{synchronized} \in \text{mod}(\text{D.m}') \\ H(o).F = 0 \quad H^i = H(o).F \mapsto 1 \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f_1, v_n \cdots v_1 \cdot o \cdot s_1 \rangle \cdot \varphi, z \setminus \{o\} \rangle_t \rightarrow \\ \Vdash_H \langle \langle \text{1}^{\text{D.m}'}, f_2[0 \mapsto o, 1 \mapsto v_1, \dots, n \mapsto v_n], \epsilon \rangle \cdot \\ (\text{pc}^{\text{C.m}} + 1, f_1, \bullet \cdot s_1) \cdot \varphi, z \rangle_t}$
$\frac{\text{(S-INVK-SYNCH-N)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{invokevirtual D.m}'(\text{T}_1, \dots, \text{T}_n) \\ \text{synchronized} \in \text{mod}(\text{D.m}') \\ H(o).F = n, n > 0 \quad H' = H(o).F \mapsto n + 1 \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f_1, v_n \cdots v_1 \cdot o \cdot s_1 \rangle \cdot \varphi, z \cup \{o\} \rangle_t \rightarrow \\ \Vdash_{H'} \langle \langle \text{1}^{\text{D.m}'}, f_2[0 \mapsto o, 1 \mapsto v_1, \dots, n \mapsto v_n], \epsilon \rangle \cdot \\ (\text{pc}^{\text{C.m}} + 1, f_1, \bullet \cdot s_1) \cdot \varphi, z \cup \{o\} \rangle_t}$	$\frac{\text{(S-START)} \quad P[\text{pc}^{\text{C.m}}] = \text{start D}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f_1, o \cdot s_1 \rangle \cdot \varphi, z \rangle_t \rightarrow \\ \Vdash_H \langle \langle \text{pc}^{\text{C.m}} + 1, f_1, s_1 \rangle \cdot \varphi, z \rangle_t, \\ \langle \langle \text{1}^{\text{java/lang/Thread.run}}, f_2[0 \mapsto o], \epsilon \rangle, \epsilon \rangle_o}$
$\frac{\text{(S-JOIN)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{join} \\ W'_o = W_o \uplus \langle \langle \text{pc}^{\text{C.m}} + 1, f_1, s_1 \rangle \cdot \varphi_{t_1}, z_{t_1} \rangle_o^o \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f_1, o \cdot s_1 \rangle \cdot \varphi_{t_1}, z_{t_1} \rangle_t, \\ \langle \varphi_{t_2} \cdot (\text{pc}^{\text{java/lang/Thread.run}}, f_1^{t_2}, s_1^{t_2}), z_{t_2} \rangle_o \dagger W_o \rightarrow \\ \Vdash_H \langle \varphi_{t_2} \cdot (\text{pc}^{\text{java/lang/Thread.run}}, f_1^{t_2}, s_1^{t_2}), z_{t_2} \rangle_o \dagger W'_o}$	$\frac{\text{(S-RETURN-RUN)} \quad P[\text{pc}^{\text{java/lang/Thread.run}}] = \text{return}}{P \Vdash_H \langle \langle \text{pc}^{\text{java/lang/Thread.run}}, f, s \rangle, z_1 \rangle_o \dagger \\ \emptyset_o^j \uplus \langle \langle \text{pc}_k, f_k, s_k \rangle \cdot \varphi_k, z_k \rangle_t \rangle_{k \in K}^o \rightarrow \\ \Vdash_H \langle \langle \text{pc}_k, f_k, s_k \rangle \cdot \varphi_k, z_k \rangle_o \rangle_{k \in K} \dagger \emptyset_o^j}$
$\frac{\text{(S-RETURN)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{return} \\ \text{synchronized} \notin \text{mod}(\text{C.m}) \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f_1, v \cdot s_1 \rangle \cdot (\text{pc}^{\text{D.m}}, f_2, \bullet \cdot s_2) \cdot \varphi, z \rangle_t \rightarrow \\ \Vdash_H \langle \text{pc}^{\text{D.m}}, f_2, v \cdot s_2 \rangle \cdot \varphi, z \rangle_t}$	$\frac{\text{(S-RETURN-SYNCH-0)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{return} \quad \text{synchronized} \in \text{mod}(\text{C.m}) \\ H(o).F = 1 \quad H' = H(o).F \mapsto 0 \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f_1, v \cdot s_1 \rangle \cdot (\text{pc}^{\text{D.m}}, f_2, \bullet \cdot s_2) \cdot \varphi, z \uplus \{o\} \rangle_t \rightarrow \\ \Vdash_{H'} \langle \langle \text{pc}^{\text{D.m}}, v \cdot f_2, s_2 \rangle \cdot \varphi, z \rangle_t}$
$\frac{\text{(S-RETURN-SYNCH-N)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{return} \quad \text{synchronized} \in \text{mod}(\text{C.m}) \\ H(o).F = n, n > 1 \quad H' = H(o).F \mapsto n - 1 \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f_1, v \cdot s_1 \rangle \cdot (\text{pc}^{\text{D.m}}, \langle f_2, \bullet \cdot s_2 \rangle \cdot \varphi, z \cup \{o\} \rangle_t \rightarrow \\ \Vdash_{H'} \langle \langle \text{pc}^{\text{D.m}}, f_2, v \cdot s_2 \rangle \cdot \varphi, z \cup \{o\} \rangle_t}$	$\frac{\text{(S-MONITOREXIT-0)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{monitorexit} \\ H(o).F = 1 \\ H' = H(o).F \mapsto 0 \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f, o \cdot s \rangle \cdot \varphi, z \uplus \{o\} \rangle_t \rightarrow \\ \Vdash_{H'} \langle \langle \text{pc}^{\text{C.m}} + 1, f, s \rangle \cdot \varphi, z \rangle_t}$
$\frac{\text{(S-MONITOREXIT-N)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{monitorexit} \\ H(o).F = n, n > 1 \\ H' = H(o).F \mapsto n - 1 \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f, o \cdot s \rangle \cdot \varphi, z \cup \{o\} \rangle_t \rightarrow \\ \Vdash_{H'} \langle \langle \text{pc}^{\text{C.m}} + 1, f, s \rangle \cdot \varphi, z \cup \{o\} \rangle_t}$	$\frac{\text{(S-MONITORENTER-0)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{monitorenter} \\ H = H(o).F = 0 \\ H' = H(o).F \rightarrow 1 \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f, o \cdot s \rangle \cdot \varphi, z \setminus \{o\} \rangle_t \rightarrow \\ \Vdash_{H'} \langle \langle \text{pc}^{\text{C.m}} + 1, f, s \rangle \cdot \varphi, z \uplus \{o\} \rangle_t}$
$\frac{\text{(S-MONITORENTER-N)} \quad \begin{array}{l} P[\text{pc}^{\text{C.m}}] = \text{monitorenter} \\ H = H(o).F = n, n > 0 \\ H' = H(o).F \rightarrow n + 1 \end{array}}{P \Vdash_H \langle \langle \text{pc}^{\text{C.m}}, f, o \cdot s \rangle \cdot \varphi, z \cup \{o\} \rangle_t \rightarrow \\ \Vdash_{H'} \langle \langle \text{pc}^{\text{C.m}} + 1, f, s \rangle \cdot \varphi, z \cup \{o\} \rangle_t}$	

Figure 4.11: Reduction rules: Invocations and synchronizations

following axioms hold:

$$(\nu a)\ell = \ell \quad ((\nu a)\ell')\&\ell = (\nu a)(\ell'\&\ell) \quad ((\nu a)\ell') + \ell = (\nu a)(\ell' + \ell)$$

Additionally, when V ranges over lams that do not contain function invocations, the following axioms hold:

$$V\&V = V \quad V + V = V \quad V\&(\ell' + \ell'') = V\&\ell' + V\&\ell'' \quad (4.1)$$

These axioms permit to rewrite a lam *without function invocations* as a *collection* (operation $+$) *of relations* (elements of a relation are gathered by the operation $\&$). Let \equiv be the least congruence containing the above axioms. (The axioms (4.1) are restricted to terms V that do not contain function invocations because $\mathfrak{m}(\bar{\tau}) \rightarrow \tau'\&((a, b)_t + (b, c)_t) \neq (\mathfrak{m}(\bar{\tau}) \rightarrow \tau'\&(a, b)_t) + (\mathfrak{m}(\bar{\tau}) \rightarrow \tau'\&(b, c)_t)$ because the evaluation of the two lams (see below) may produce terms with different names.)

Operational semantics. Let a *lam context*, noted $L[\]$, be a term derived by the following syntax:

$$L[\] ::= [\] \quad | \quad \ell\&L[\] \quad | \quad \ell + L[\]$$

As usual $L[\ell]$ is the lam where the hole of $L[\]$ is replaced by ℓ . According to the syntax, lam contexts have no ν -binder; that is, the hassle of name captures is avoided. The operational semantics of a program (\mathcal{L}, ℓ) is a transition system where *states* are lams, the *transition relation* is the least one satisfying the rule

$$\frac{\mathfrak{m}(\bar{\tau}) \rightarrow \tau' = (\nu \bar{c})\ell_{\mathfrak{m}} \in \mathcal{L} \quad (\bar{\tau} \rightarrow \tau')\sigma = \bar{\tau}'' \rightarrow \tau''' \quad \bar{c}' \text{ fresh} \quad (\ell_{\mathfrak{m}}\{\bar{c}'/\bar{c}\})\sigma = \ell'_{\mathfrak{m}}}{L[\mathfrak{m}(\bar{\tau}'') \rightarrow \bar{\tau}'''] \rightarrow L[\ell'_{\mathfrak{m}}]}$$

and the initial state is the lam ℓ' such that $\ell \equiv (\nu \bar{c})\ell'$ and ℓ' does not contain any ν -binder, as well as the lam $\ell_{\mathfrak{m}}$. (The class name in the names of lam functions has been dropped, for simplicity.) We write \rightarrow^* for the reflexive and transitive closure of \rightarrow .

By (RED), a lam ℓ is evaluated by successively replacing function invocations with the corresponding lam instances. Name creation is handled by replacing bound names of function bodies with fresh names.

$$\begin{array}{c}
\text{(L-0)} \\
\ell \geq 0 \\
\\
\text{(L-RES)} \\
\frac{\ell \geq \ell'}{(\nu a)\ell \geq (\nu a)\ell'} \\
\\
\text{(L-PLUS)} \\
\frac{\ell_1 \geq \ell'_1 \quad \ell_2 \geq \ell'_2}{\ell_1 + \ell'_1 \geq \ell_2 + \ell'_2} \\
\\
\text{(L-AND)} \\
\frac{\ell_1 \geq \ell'_1 \quad \ell_2 \geq \ell'_2}{\ell_1 \& \ell'_1 \geq \ell_2 \& \ell'_2} \\
\\
\text{(L-INVK)} \\
\frac{\text{BCT}(\mathbf{C.m}) = (\bar{\tau}, t, a) \rightarrow (\nu \bar{a}') \langle \vartheta, \mathcal{T}, \mathcal{R}, K, \bar{\vartheta}, \ell \rangle \quad \bar{b} \text{ fresh}}{\mathbf{C.m}(a[\mathbf{f} : \bar{\tau}], \bar{\tau}') \rightarrow \vartheta^\nabla \geq \ell\{\bar{b}/\bar{a}'\}}
\end{array}$$

Figure 4.12: The later-stage relation

4.9.4 Flattening and circularities

Informally, a lam has a *circularity* when it has a conjunction of dependencies whose transitive closure contains a pair $(a, a)_\surd$ (see Section 4.6). Here the hassle is that lams cannot be always rewritten in a form suitable for defining circularities, namely disjunctions of conjunctions, because the equations (4.1) do not apply to terms containing invocations. To overcome this problem, we define the operation of flattening. The transitive closure of a lam has been defined in Section 4.6.

Definition 9.1. *The flattening of a lam ℓ , noted $(\ell)^b$, is the lam ℓ where every function invocation has been replaced by 0 .*

For example, let $\ell = \mathbf{m}(a, b, c, t) + (a, b)_t \& \mathbf{m}'(b, c, t') \& \mathbf{m}(d, b, c, t)$ (we assume that return types of lam functions are empty). Then $(\ell)^b = 0 + (a, b)_t \& 0 \& 0$.

Definition 9.2. *A lam ℓ has a circularity if $(\ell)^b \equiv \sum_{i \in I} \ell_i$, where every ℓ_i is a conjunction of dependencies, and, for some a and i , $(a, a)_\surd$ is a subterm of ℓ_i^+ .*

A lam program (\mathcal{L}, ℓ) has a circularity if there exists ℓ' such that $\ell \rightarrow^ \ell'$ and ℓ' has a circularity.*

4.9.5 Later stage relation

In the following subject reduction theorem we use a relation, called *later-stage relation* \geq , which is the least congruence with respect to lams that contains the rules presented in Figure 4.12. The notation \geq assumes the presence of a behavioural class table BCT – we should have noted the later-stage relation as \geq_{BCT} , but we prefer to keep the BCT implicit. A relevant property of the later stage relation is the following.

Lemma 4.9.1 (Circularities). *Let $\ell \geq \ell'$. If ℓ' has a circularity then ℓ has also a circularity.*

This result mostly follows from the formal results of [34, 48].

4.9.6 Runtime typing

Definition 9.3 (Γ - H Agreement). *An environment Γ agrees with a heap H , written $\Gamma \approx H$ if*

- $o \in \text{dom}(H)$ implies $o \in \text{dom}(\Gamma)$ and
- $H(o).\text{class} = \mathbf{C}$ and $H(o).f_i \mapsto v_i$, for $i \in 1..n$, imply $\Gamma = o \mapsto ([\mathbf{f}_1^{\text{h}_1} : \chi_1, \dots, \mathbf{f}_i^{\text{h}_i} : \chi_i, \dots, \mathbf{f}_n^{\text{h}_n} : \chi_n], \mathbf{C})$, where, if v_i is an object name, then $v_i = \chi_i$.

Definition 9.4 (S/F/Z-Agreements). *Given Γ , we define the following agreements between the static environments S , F and Z and the corresponding dynamic ones s , f , and z .*

S-agreement.

- $\epsilon \approx_{\Gamma} \epsilon$,
- $\chi \cdot S \approx_{\Gamma} \bullet \cdot s$, if $S \approx_{\Gamma} s$,
- $\chi \cdot S \approx_{\Gamma} v \cdot s$, if $S \approx_{\Gamma} s$ and $\Gamma \vdash v : \chi$.

F-agreement.

- $\emptyset \approx_{\Gamma} \emptyset$,
- $F[x \mapsto \chi] \approx_{\Gamma} f[x \mapsto v]$, if $F \approx_{\Gamma} f$ and $\Gamma \vdash v : \chi$.

Z-agreement.

- $\epsilon \approx_{\Gamma} \epsilon$,
- $o \cdot Z \approx o \cdot z$ if $Z \approx z$

Definition 9.5 (Reachable instructions set). *Let $\text{reachable}(P, i)$, called the set of addresses reachable from the instruction $i \in \text{dom}(P)$, be the least set satisfying the following equations:*

$$\begin{array}{ll}
 \text{reachable}(P, i) & = \emptyset & \text{if } P[i] = \text{return} \\
 \text{reachable}(P, i) & = \{L\} \cup \text{reachable}(P, L) & \text{if } P[i] = \text{goto } L \\
 \text{reachable}(P, i) & = \{i+1, L\} \cup \text{reachable}(P, i+1) \cup \text{reachable}(P, L) & \text{if } P[i] = \text{if } L \\
 \text{reachable}(P, i) & = \{i+1\} \cup \text{reachable}(P, i+1) & \text{otherwise}
 \end{array}$$

The typing judgement for configurations is

$$P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\Vdash_H \mathfrak{C}) : \ell$$

that is defined by the following two rules:

(T-CONF-AND)

$$\frac{(P, \text{BCT}, \Gamma^{(h)}, F^{(h)}, S^{(h)}, Z^{(h)}, \mathcal{T}^{(h)}, \mathcal{R}^{(h)}, K^{(h)} \vdash \langle \varphi_h, z_h \rangle_{t_h} : \ell_h)_{h \in 1..n}}{\Gamma = \bigcup_{h \in 1..n} \Gamma^{(h)} \quad F = F^{(1)} \dots F^{(n)} \quad S = S^{(1)} \dots S^{(n)} \quad Z = Z^{(1)} \dots Z^{(n)} \\ \mathcal{T} = \bigcup_{h \in 1..n} \mathcal{T}^{(h)} \quad \mathcal{R} = \bigcup_{h \in 1..n} \mathcal{R}^{(h)} \quad K = \bigcup_{h \in 1..n} K^{(h)} \\ \Gamma \approx H}$$

$$P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\|_H \langle (pc^{c_1 \cdot m_1}, f_1, s_1) \cdot \varphi_1, z_1 \rangle_{t_1} \dots \langle (pc^{c_n \cdot m_n}, f_n, s_n) \cdot \varphi_n, z_n \rangle_{t_n} \rangle : \&_{i \in 1..n} \ell_i$$

(T-CONF-SINGLE)

$$\frac{(P, \text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\mathbf{C.m}] : \ell_i; \Psi_i)_{i \in \text{reachable}(P[\mathbf{C.m}], pc^{c.m})} \\ P, \text{BCT}, \Gamma'', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'' \vdash \langle \varphi, (z \setminus Z') \rangle_t : \ell \quad \Gamma'' = \Gamma' \cup \Gamma'' \\ F = F' \cdot F'' \quad S = S' \cdot S'' \quad Z = Z' \cdot Z'' \quad \mathcal{T} = \mathcal{T}' \cup \mathcal{T}'' \quad \mathcal{R} = \mathcal{R}' \cup \mathcal{R}'' \quad K = K' \cup K'' \\ Z \approx z \quad S'_{pc^{c.m}} \approx s \quad F'_{pc^{c.m}} \approx f}{P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash \langle (pc^{c.m}, f, s) \cdot \varphi, z \rangle_t : \sum_{i \in \text{reachable}(P[\mathbf{C.m}], pc^{c.m})} \ell_i + \ell}$$

4.9.7 Subject Reduction

Lemma 4.9.2. *Let $\mathbf{C.m}$ be a method of a JVML_d program. If*

$$\text{BCT}(\mathbf{C.m})(\bar{\tau}, t, a)(\bar{a}') = \langle \vartheta, \mathcal{T}', \mathcal{R}', K', \bar{\vartheta}, \ell \rangle.$$

then

$$\left(\text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K, i \vdash_t P[\mathbf{C.m}] : \ell_i; \Psi_i \right)_{i \in \text{reachable}(P[\mathbf{C.m}], 1^{\mathbf{C.m}})},$$

where

$$\begin{aligned} F_1 &= F_{\top}[0 \mapsto \text{root}(\tau_0), \dots, k \mapsto \text{root}(\tau_k)], \\ \Gamma_1 &= \bigoplus_{i \in 0..k} \text{flat}(\tau_i), \\ S_1 &= \varepsilon, \\ Z_1 &= A, \quad (\text{if } \mathbf{C.m} \text{ is synchronized } A = \text{root}(\tau_0) \cdot a, \text{ otherwise } A = a) \\ \mathcal{T}_1 &= \emptyset, \\ \mathcal{R}_1 &= \emptyset, \\ K_1 &= \emptyset, \\ \ell &= \sum_{i \in \text{reachable}(P[\mathbf{C.m}], 1^{\mathbf{C.m}})} \ell_i, \\ \bigsqcup_{i \in \text{reachable}(P[\mathbf{C.m}], 1^{\mathbf{C.m}})} \Psi_i &= (\vartheta, A, \mathcal{T}', \mathcal{R}', K', \Gamma'), \\ \bar{\vartheta} &= \text{mk_tree}(\Gamma', (\text{root}(\tau_0), \dots, \text{root}(\tau_k))), \text{ and} \\ \bar{a}' &= \text{fn}(\vartheta, \mathcal{T}', \mathcal{R}', K', \bar{\vartheta}) \setminus \text{fn}(\tau_0, \dots, \tau_k, t, a). \end{aligned}$$

Theorem 4.9.3 (Subject Reduction). *If $P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\|_H \mathfrak{C}) : \ell$ and $\Gamma \approx H$ and $\|_H \mathfrak{C} \rightarrow \|_{H'} \mathfrak{C}'$, then there exists ℓ' and $\Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K'$ such that $\Gamma' \approx H'$ and $P, \text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K' \vdash (\|_{H'} \mathfrak{C}') : \ell'$ and $\ell \geq \ell'$.*

Proof. (Sketch) By case analysis on the last reduction rule used.

Case pop. We have $P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\llbracket\!-\! \rrbracket_H \langle (pc^{\text{C.m}}, f, v \cdot s) \cdot \varphi, z \rangle_t) : \ell$ and

$$\frac{\text{(S-POP)} \quad P[pc^{\text{C.m}}] = \text{pop}}{P \llbracket\!-\! \rrbracket_H \langle (pc^{\text{C.m}}, f, v \cdot s) \cdot \varphi, z \rangle_t \rightarrow \llbracket\!-\! \rrbracket_H \langle (pc^{\text{C.m}} + 1, f, s) \cdot \varphi, z \rangle_t}$$

By the configuration typing rules, we get

$$\begin{aligned} & (\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\text{C.m}] : \ell_i)^{i \in \text{reachable}(P[\text{C.m}], pc^{\text{C.m}})} \\ & P, \text{BCT}, \Gamma'', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'' \vdash \langle \varphi, z \rangle_t : \ell_\varphi \quad \Gamma = \Gamma' \cup \Gamma'' \quad F = F' \cdot F'' \\ & S = S' \cdot S'' \quad Z = Z' \cdot Z'' \quad \mathcal{T} = \mathcal{T}' \cup \mathcal{T}'' \quad \mathcal{R} = \mathcal{R}' \cup \mathcal{R}'' \quad K = K' \cup K'' \\ & Z'_{pc^{\text{C.m}}} \approx z \quad S'_{pc^{\text{C.m}}} \approx v \cdot s \quad F'_{pc^{\text{C.m}}} \approx f \end{aligned}$$

From $(\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\text{C.m}] : \ell_i)^{i \in \text{reachable}(P[\text{C.m}], pc^{\text{C.m}})}$ and Definition 9.5 we derive that

$$\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', pc^{\text{C.m}} \vdash_t P[\text{C.m}] : \ell_{pc^{\text{C.m}}}$$

and

$$(\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\text{C.m}] : \ell_i)^{i \in \text{reachable}(P[\text{C.m}], pc^{\text{C.m}} + 1)}.$$

Therefore $\ell = (\ell_{pc^{\text{C.m}}} + (\sum \ell_i)^{i \in \text{reachable}(P[\text{C.m}], pc^{\text{C.m}} + 1)}) + \ell_\varphi$. By the typing rule for **pop**: $S'_{pc^{\text{C.m}}} = \chi \cdot S'_{pc^{\text{C.m}} + 1}$; therefore, since $S'_{pc^{\text{C.m}}} \approx v \cdot s$ then $S'_{pc^{\text{C.m}} + 1} \approx s$, by Definition 9.4.

We can apply again the configuration typing rules and we get that $\ell' = (\sum \ell_i)^{i \in \text{reachable}(P[\text{C.m}], pc^{\text{C.m}} + 1)} + \ell_\varphi$, thus $\ell \geq \ell'$.

Case invokevirtual.

We have $P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\llbracket\!-\! \rrbracket_H \langle (pc^{\text{C.m}}, f, v_n \cdots \cdots v_1 \cdot o \cdot s) \cdot \varphi, z \rangle_t) : \ell$ and then one of three operational rules may have been applied: (S-INVK), (S-INVK-SYNCH-0), or (S-INVK-SYNCH-N). Let us consider the former:

$$\frac{\text{(S-INVK)} \quad \begin{array}{l} P[pc^{\text{C.m}}] = \text{invokevirtual } D.m'(T_1, \dots, T_n) \\ \text{synchronized } \notin \text{mod}(D.m') \end{array}}{P \llbracket\!-\! \rrbracket_H \langle (pc^{\text{C.m}}, f_1, v_n \cdots \cdots v_1 \cdot o \cdot s_1) \cdot \varphi, z \rangle_t \rightarrow \llbracket\!-\! \rrbracket_H \langle (1^{D.m'}, f_2[0 \mapsto o, 1 \mapsto v_1, \dots, n \mapsto v_n], \epsilon) \cdot (pc^{\text{C.m}} + 1, f_1, \bullet \cdot s_1) \cdot \varphi, z \rangle_t}$$

By the configuration typing rules, we get

$$\begin{aligned}
& (\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\mathbf{C.m}] : \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], pc^{\mathbf{C.m}})} \\
& P, \text{BCT}, \Gamma'', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'' \vdash \langle \varphi, z \rangle_t : \ell_\varphi \quad \Gamma = \Gamma' \cup \Gamma'' \\
& F = F' \cdot F'' \quad S = S' \cdot S'' \quad Z = Z' \cdot Z'' \quad \mathcal{T} = \mathcal{T}' \cup \mathcal{T}'' \quad \mathcal{R} = \mathcal{R}' \cup \mathcal{R}'' \\
& K = K' \cup K'' \quad Z'_{pc^{\mathbf{C.m}}} \approx z \quad S'_{pc^{\mathbf{C.m}}} \approx v_n \cdots v_1 \cdot o \cdot s \quad F'_{pc^{\mathbf{C.m}}} \approx f
\end{aligned}$$

From $(\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\mathbf{C.m}] : \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], pc^{\mathbf{C.m}})}$ and Definition 9.5 we derive that

$$\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', pc^{\mathbf{C.m}} \vdash_t P[\mathbf{C.m}] : \ell_{pc^{\mathbf{C.m}}}$$

and

$$(\text{BCT}, \Gamma', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'', i \vdash_t P[\mathbf{C.m}] : \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], pc^{\mathbf{C.m}+1})}.$$

Therefore

$$\ell = \ell_{pc^{\mathbf{C.m}}} + \left(\sum \ell_i \right)^{i \in \text{reachable}(P[\mathbf{C.m}], pc^{\mathbf{C.m}+1})} + \ell_\varphi. \quad (4.2)$$

We assume that $\mathbf{C.m}$ is recursive (this case and that where $pc^{\mathbf{C.m}}$ is inside an iteration are the interesting cases). By (T-INVK) typing rule:

$$\ell_{pc^{\mathbf{C.m}}} = \widehat{\mathcal{T}'_{pc^{\mathbf{C.m}}}} \ \& \ \widehat{\mathcal{R}'_{pc^{\mathbf{C.m}}}} \ \& \ \widehat{Z'_{pc^{\mathbf{C.m}}}}^t \ \& \ \text{D.m}'(\tau_1, \dots, \tau_n, t, [Z'_{pc^{\mathbf{C.m}}}]) \rightarrow \vartheta,$$

where

$$\begin{aligned}
& S'_{pc^{\mathbf{C.m}}} = \chi_n \cdots \chi_1 \cdot S''' \quad \text{typeof}(\Gamma'_{pc^{\mathbf{C.m}}}, \chi_1) = \text{D} \\
& \text{mk_tree}(\Gamma'_{pc^{\mathbf{C.m}}}, (\chi_1, \dots, \chi_n)) = (\tau_1, \dots, \tau_n) \quad \bar{b} = \text{names}(pc^{\mathbf{C.m}}) \\
& \text{BCT}(\text{D.m}')(\tau_1, \dots, \tau_n, t, [Z'_{pc^{\mathbf{C.m}}}])(\bar{b}) = \langle \vartheta, \mathcal{T}''', \mathcal{R}''', K''', (\vartheta_1, \dots, \vartheta_n), \ell_{\text{D.m}'} \rangle \\
& \Gamma'_{pc^{\mathbf{C.m}+1}} = \Gamma'_{pc^{\mathbf{C.m}}}[\text{flat}(\vartheta) \oplus (\bigoplus_{pc^{\mathbf{C.m}} \in 1..n} \text{flat}(\vartheta_{pc^{\mathbf{C.m}}}))] \\
& S'_{pc^{\mathbf{C.m}+1}} = \text{root}(\vartheta) \cdot S''' \quad F'_{pc^{\mathbf{C.m}+1}} = F'_{pc^{\mathbf{C.m}}} \quad Z'_{pc^{\mathbf{C.m}+1}} = Z'_{pc^{\mathbf{C.m}}} \quad K'_{pc^{\mathbf{C.m}+1}} = K'_{pc^{\mathbf{C.m}}} \cup K''' \\
& \mathcal{T}'_{pc^{\mathbf{C.m}+1}}, \mathcal{R}'_{pc^{\mathbf{C.m}+1}} = (\mathcal{T}'_{pc^{\mathbf{C.m}}} \setminus K''') \cup \mathcal{T}''', (\mathcal{R}'_{pc^{\mathbf{C.m}}} \setminus K''') \cup \mathcal{R}'''
\end{aligned}$$

By (T-INVK) typing rule $S'_{pc^{\mathbf{C.m}}} = \chi_n \cdots \chi_1 \cdot S'''$ and $S'_{pc^{\mathbf{C.m}+1}} = \text{root}(\vartheta) \cdot S'''$ so since $S'_{pc^{\mathbf{C.m}}} \approx v_n \cdots v_1 \cdot o \cdot s$ then $S'_{pc^{\mathbf{C.m}+1}} \approx \bullet \cdot s$.

Therefore we obtain that

$$P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash \langle (pc^{\mathbf{C.m}} + 1, f, \bullet \cdot s) \cdot \varphi, z \rangle_t : \ell_{pc^{\mathbf{C.m}+1}} \quad (4.3)$$

where $\ell_{pc^{\mathbf{C.m}+1}} = \left(\sum \ell_i \right)^{i \in \text{reachable}(P[\mathbf{C.m}], pc^{\mathbf{C.m}+1})} + \ell_\varphi$

By Lemma 4.9.2,

$$\left(\text{BCT}, \Gamma^{\text{D.m}'}, F^{\text{D.m}'}, S^{\text{D.m}'}, Z^{\text{D.m}'}, \mathcal{T}^{\text{D.m}'}, \mathcal{R}^{\text{D.m}'}, K^{\text{D.m}'}, i \vdash_t P[\text{D.m}'] : \ell_i^{\text{D.m}'} ; \Psi_i \right)^{i \in \text{reachable}(P[\text{D.m}'], 1^{\text{D.m}'})},$$

where

$$\begin{aligned}
F_1^{\mathcal{D}.m'} &= F_{\top}^{\mathcal{D}.m'}[0 \mapsto \text{root}(\tau_1), \dots, n \mapsto \text{root}(\tau_n)], \\
\Gamma_1^{\mathcal{D}.m'} &= \bigoplus_{i \in 1..n} \text{flat}(\tau_i), \\
S_1^{\mathcal{D}.m'} &= \varepsilon, \\
Z_1^{\mathcal{D}.m'} &= [Z'_{pc^c.m}], \\
\mathcal{T}_1^{\mathcal{D}.m'} &= \emptyset, \\
\mathcal{R}_1^{\mathcal{D}.m'} &= \emptyset, \\
K_1^{\mathcal{D}.m'} &= \emptyset, \\
\ell^{\mathcal{D}.m'} &= \sum_{i \in \text{reachable}(P[\mathcal{D}.m'], 1^{\mathcal{D}.m'})} \ell_i^{\mathcal{D}.m'}, \\
\prod_{i \in \text{reachable}(P[\mathcal{D}.m'], 1^{\mathcal{D}.m'})} \Psi_i &= (\vartheta, [Z'_{pc^c.m}], \mathcal{T}^{\mathcal{D}.m'}, \mathcal{R}^{\mathcal{D}.m'}, K^{\mathcal{D}.m'}, \Gamma^{\mathcal{D}.m'}), \\
\bar{\vartheta} &= \text{mk_tree}(\Gamma^{\mathcal{D}.m'}, (\text{root}(\tau_1), \dots, \text{root}(\tau_n))), \text{ and} \\
\bar{b} &= \text{fn}(\vartheta, \mathcal{T}^{\mathcal{D}.m'}, \mathcal{R}^{\mathcal{D}.m'}, K^{\mathcal{D}.m'}, \bar{\vartheta}) \setminus \text{fn}(\tau_1, \dots, \tau_n, t, a).
\end{aligned}$$

Thus

$$\begin{aligned}
P, \text{BCT}, \Gamma^{\mathcal{D}.m'} \cup \Gamma, F^{\mathcal{D}.m'} \cdot F, S^{\mathcal{D}.m'} \cdot S, Z^{\mathcal{D}.m'} \cdot Z, \mathcal{T}^{\mathcal{D}.m'} \cup \mathcal{T}^{\mathcal{D}.m'}, \mathcal{R}^{\mathcal{D}.m'} \cup \mathcal{R}^{\mathcal{D}.m'}, K^{\mathcal{D}.m'} \cup K \vdash \\
\langle \langle 1^{\mathcal{D}.m'}, f'[0 \mapsto o, 1 \mapsto v_1, \dots, n \mapsto v_n, \text{this} \mapsto o], \epsilon \rangle \\
\cdot (pc^c.m + 1, f, s) \cdot \varphi, z \rangle_t : \ell^{\mathcal{D}.m'} + \ell_{pc^c.m+1}
\end{aligned}$$

and $\ell \geq \ell^{\mathcal{D}.m'} + \ell_{pc^c.m+1}$

The last step is to prove that $\Gamma^{\mathcal{D}.m'} \cup \Gamma \approx H$, which follows directly from Definition 9.3 and the fact that $\Gamma \approx H$.

The other two cases of the operational semantics are proven in a similar way, except that we may have to check the agreement of the modified environment z , but this is straightforward from the definition of agreement and the typing rule for configuration.

Case start. We have $P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\Vdash_H \langle (pc^c.m, f_1, o \cdot s_1) \cdot \varphi, z \rangle_t) : \ell$ and then

$$\begin{array}{c}
(\text{s-START}) \\
\frac{P[pc^c.m] = \text{start D}}{P \Vdash_H \langle (pc^c.m, f_1, o \cdot s_1) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle (pc^c.m + 1, f_1, s_1) \cdot \varphi, z \rangle_t, \langle (1^{\mathcal{D}.run}, f_2[0 \mapsto o], \epsilon), \epsilon \rangle}
\end{array}$$

By the configuration typing rules, we obtain

$$\begin{aligned}
&(\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\mathbf{C}.m] : \ell_i)^{i \in \text{reachable}(P[\mathbf{C}.m], pc^c.m)} \\
&P, \text{BCT}, \Gamma'', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'' \vdash \langle \varphi, z \rangle_t : \ell_\varphi \quad \Gamma = \Gamma' \cup \Gamma'' \\
&F = F' \cdot F'' \quad S = S' \cdot S'' \quad Z = Z' \cdot Z'' \quad \mathcal{T} = \mathcal{T}' \cup \mathcal{T}'' \quad K = K' \cup K'' \\
&Z'_{pc^c.m} \approx z \quad S'_{pc^c.m} \approx o \cdot s \quad F'_{pc^c.m} \approx f
\end{aligned}$$

From $(\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\mathbf{C.m}] : \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], pc^{\mathbf{C.m}})}$ and Definition 9.5 we derive that

$$\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', pc^{\mathbf{C.m}} \vdash_t P[\mathbf{C.m}] : \ell_{pc^{\mathbf{C.m}}}$$

and

$$(\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\mathbf{C.m}] : \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], pc^{\mathbf{C.m}+1})}.$$

Therefore

$$\ell = \ell_{pc^{\mathbf{C.m}}} + \left(\sum \ell_i \right)^{i \in \text{reachable}(P[\mathbf{C.m}], pc^{\mathbf{C.m}+1})} + \ell_\varphi. \quad (4.4)$$

By (T-START) typing rule:

$$\ell_{pc^{\mathbf{C.m}}} = \widehat{\mathcal{T}'_{pc^{\mathbf{C.m}}}} \ \& \ \widehat{\mathcal{R}'_{pc^{\mathbf{C.m}}}} \ \& \ \widehat{Z'_{pc^{\mathbf{C.m}}}}^t$$

where

$$\begin{aligned} S'_{pc^{\mathbf{C.m}}} &= o \cdot S'_{pc^{\mathbf{C.m}+1}} \\ \text{typeof}(\Gamma'_{pc^{\mathbf{C.m}}}, o) &= \mathbf{D} \quad \text{mk_tree}(\Gamma'_{pc^{\mathbf{C.m}}}, o) = \tau \quad \bar{b} = \text{names}(pc^{\mathbf{C.m}}) \\ \text{BCT}(\mathbf{D.run})(\tau, o, \text{lock}_o)(\bar{b}) &= \langle \text{void}, \mathcal{T}''', \mathcal{R}''', K''', \vartheta, \ell \rangle \\ \Gamma'_{pc^{\mathbf{C.m}+1}} &= \Gamma'_{pc^{\mathbf{C.m}}} \oplus \text{flat}(\vartheta) \quad F'_{pc^{\mathbf{C.m}}} = F'_{pc^{\mathbf{C.m}+1}} \quad Z'_{pc^{\mathbf{C.m}}} = Z'_{pc^{\mathbf{C.m}+1}} \quad K'_{pc^{\mathbf{C.m}+1}} = K'_{pc^{\mathbf{C.m}}} \\ \mathcal{T}_{pc^{\mathbf{C.m}+1}}, \mathcal{R}_{pc^{\mathbf{C.m}+1}} &= \begin{cases} \mathcal{T}_{pc^{\mathbf{C.m}}} \cup \{o\} \cup \mathcal{T}''', \mathcal{R}_{pc^{\mathbf{C.m}}} \cup \mathcal{R}''' & \text{if } i \text{ is not recursive} \\ & \text{and } o \notin \mathcal{T}_{pc^{\mathbf{C.m}}} \cup \mathcal{R}_{pc^{\mathbf{C.m}}} \\ \mathcal{T}_{pc^{\mathbf{C.m}}}, \mathcal{R}_{pc^{\mathbf{C.m}}} \cup \mathcal{R}''' \cup \mathcal{T}''' \cup \{o\} & \text{otherwise} \end{cases} \end{aligned}$$

Let $S'_{pc^{\mathbf{C.m}}} \approx o \cdot s_1$; then $S'_{pc^{\mathbf{C.m}+1}} \approx s_1$. Therefore we obtain that

$$P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash \langle (pc^{\mathbf{C.m}} + 1, f_1, s_1) \cdot \varphi, z \rangle_t : \ell_{pc^{\mathbf{C.m}+1}} \quad (4.5)$$

where $\ell_{pc^{\mathbf{C.m}+1}} = \left(\sum \ell_i \right)^{i \in \text{reachable}(P[\mathbf{C.m}], pc^{\mathbf{C.m}+1})} + \ell_\varphi$. By Lemma 4.9.2,

$$\left(\text{BCT}, \Gamma^{\mathbf{D.run}}, F^{\mathbf{D.run}}, S^{\mathbf{D.run}}, Z^{\mathbf{D.run}}, T^{\mathbf{D.run}}, K^{\mathbf{D.run}}, i \vdash_t P[\mathbf{D.run}] : \ell_i^{\mathbf{D.run}} ; \Psi_i \right)^{i \in \text{reachable}(P[\mathbf{D.run}], 1^{\mathbf{D.run}})},$$

where

$$\begin{aligned} F_1^{\mathbf{D.run}} &= F_{\top}^{\mathbf{D.run}}[0 \mapsto \text{root}(\tau)], \\ \Gamma_1^{\mathbf{D.run}} &= \text{flat}(\tau), \\ S_1^{\mathbf{D.run}} &= \varepsilon, \\ Z_1^{\mathbf{D.run}} &= [Z'_{pc^{\mathbf{C.m}}}], \\ \mathcal{T}_1^{\mathbf{D.run}} &= \emptyset, \\ \mathcal{R}_1^{\mathbf{D.run}} &= \emptyset, \\ K_1^{\mathbf{D.run}} &= \emptyset, \\ \ell^{\mathbf{D.run}} &= \sum_{i \in \text{reachable}(P[\mathbf{D.run}], 1^{\mathbf{D.run}})} \ell_i^{\mathbf{D.run}}, \\ \bigsqcup_{i \in \text{reachable}(P[\mathbf{D.run}], 1^{\mathbf{D.run}})} \Psi_i &= (\vartheta, [Z'_{pc^{\mathbf{C.m}}}], \mathcal{T}^{\mathbf{D.run}}, \mathcal{R}^{\mathbf{D.run}}, K^{\mathbf{D.run}}, \Gamma^{\mathbf{D.run}}), \\ \vartheta &= \text{mk_tree}(\Gamma^{\mathbf{D.run}}, o), \\ \bar{b} &= \text{fn}(\text{void}, \mathcal{T}^{\mathbf{D.run}}, \mathcal{R}^{\mathbf{D.run}}, K^{\mathbf{D.run}}, \vartheta) \setminus \text{fn}(\tau, o, \text{lock}_o). \end{aligned}$$

Thus

$$P, \text{BCT}, \Gamma^{\text{D.run}} \cup \Gamma, F^{\text{D.run}} \cdot F, S^{\text{D.run}} \cdot S, Z^{\text{D.run}} \cdot Z, T^{\text{D.run}} \cup \mathcal{T}, \mathcal{R}, K^{\text{D.run}} \cup K \vdash \\ (\Vdash_H \langle (pc^{\text{C.m}} + 1, f_1, s_1) \cdot \varphi, z \rangle_t, \\ \langle (1^{\text{D.run}}, f_2[0 \mapsto o], \epsilon), \epsilon \rangle_o) : \ell^{\text{D.run}} \& \ell_{pc^{\text{C.m}}+1}.$$

and $\ell \geq \ell^{\text{D.run}} \& \ell_{pc^{\text{C.m}}+1}$

The last step is to prove that $\Gamma^{\text{D.run}} \cup \Gamma \approx H$, which follows directly from Definition 9.3 and the fact that $\Gamma \approx H$.

Case return. We have

$$P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\Vdash_H \langle (pc^{\text{C.m}}, f_1, v \cdot s_1) \cdot (pc^{\text{D.m}}, f_2, \bullet \cdot s_2) \cdot \varphi, z \rangle_t) : \ell$$

The following rules may have been applied: (S-RETURN), (S-RETURN-RUN), (S-RETURN-SYNCH-0), and (S-RETURN-SYNCH-N). Let us consider the first one (the other ones are similar).

$$\frac{\text{(S-RETURN)} \quad P[pc^{\text{C.m}}] = \text{return} \quad \text{synchronized} \notin \text{mod}(\text{C.m})}{P \Vdash_H \langle (pc^{\text{C.m}}, f_1, v \cdot s_1) \cdot (pc^{\text{D.m}}, f_2, \bullet \cdot s_2) \cdot \varphi, z \rangle_t \rightarrow \Vdash_H \langle pc^{\text{D.m}}, f_2, v \cdot s_2 \rangle_t \cdot \varphi, z \rangle_t}$$

By the configuration typing rules, we have

$$\begin{aligned} & (\text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K', i \vdash_t P[\text{C.m}] : \ell_i)^{i \in \text{reachable}(P[\text{C.m}], pc^{\text{C.m}})} \\ & P, \text{BCT}, \Gamma'', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'' \vdash \langle (pc^{\text{D.m}}, f_2, \bullet \cdot s_2) \cdot \varphi, z \rangle_t : \ell_\varphi \quad \Gamma = \Gamma' \cup \Gamma'' \\ & F = F' \cdot F'' \quad S = S' \cdot S'' \quad Z = Z' \cdot Z'' \quad \mathcal{T} = \mathcal{T}' \cup \mathcal{T}'' \quad \mathcal{R} = \mathcal{R}' \cup \mathcal{R}'' \quad K = K' \cup K'' \\ & Z'_{pc^{\text{C.m}}} \approx z \quad S'_{pc^{\text{C.m}}} \approx v \cdot s_1 \quad F'_{pc^{\text{C.m}}} \approx f_1 \end{aligned}$$

Therefore

$$\ell = \left(\left(\sum \ell_i \right)^{i \in \text{reachable}(P[\text{C.m}], pc^{\text{C.m}})} + \left(\sum \ell_i \right)^{i \in \text{reachable}(P[\text{D.m}], pc^{\text{D.m}})} + \ell_\varphi \right) \quad (4.6)$$

Moreover, $S''_{pc^{\text{D.m}'}} \approx \bullet \cdot s_2$ and $S''_{pc^{\text{D.m}'}}$ has been set while typing the invocation of C.m as $S''_{pc^{\text{D.m}'}} = \text{root}(\vartheta) \cdot S'''$, where $\text{root}(\theta)$ is the value returned by the method invocation, then $S''_{pc^{\text{D.m}'}} \approx v \cdot s_2$.

Therefore we obtain

$$\ell' = \left(\sum \ell_i \right)^{i \in \text{reachable}(P[\text{D.m}], pc^{\text{D.m}})} + \ell_\varphi \quad (4.7)$$

then $\ell \geq \ell'$.

Case new. We have

$$P, \text{BCT}, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\Vdash_H \langle (pc^{\text{C.m}}, f, s) \cdot \varphi, z \rangle_t) : \ell$$

and

$$\frac{\text{(S-NEW)} \quad P[\text{pc}^{\mathbf{C.m}}] = \mathbf{new D} \quad o \notin \text{dom}(H) \quad H' = H[o \mapsto (\rho^D, D)]}{P \Vdash_H \langle (\text{pc}^{\mathbf{C.m}}, f, s) \cdot \varphi, z \rangle_t \rightarrow \Vdash_{H'} \langle (\text{pc}^{\mathbf{C.m}} + 1, f, o \cdot s) \cdot \varphi, z \rangle_t}$$

By the configuration typing rules, we get

$$\begin{aligned} & (\text{BCT}, \Gamma'', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'', i \vdash_t P[\mathbf{C.m}] : \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], \text{pc}^{\mathbf{C.m}})} \\ & P, \text{BCT}, \Gamma''', F''', S''', Z''', \mathcal{T}''', \mathcal{R}''', K''' \vdash \langle \varphi, z \setminus Z'' \rangle_t : \ell_\varphi \\ & \Gamma = \Gamma'' \cup \Gamma''' \quad F = F'' \cdot F''' \quad S = S'' \cdot S''' \\ & Z = Z'' \cdot Z''' \quad \mathcal{T} = \mathcal{T}'' \cup \mathcal{T}''' \quad K = K'' \cup K''' \\ & Z \approx z \quad S''_{\text{pc}^{\mathbf{C.m}}} \approx s \quad F''_{\text{pc}^{\mathbf{C.m}}} \approx f \end{aligned}$$

Therefore

$$\ell = ((\sum \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], \text{pc}^{\mathbf{C.m}})} + \ell_\varphi) \quad (4.8)$$

From $(\text{BCT}, \Gamma'', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'', i \vdash_t P[\mathbf{C.m}] : \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], \text{pc}^{\mathbf{C.m}})}$ and Definition 9.5 we derive that

$$\text{BCT}, \Gamma'', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'', \text{pc}^{\mathbf{C.m}} \vdash_t P[\mathbf{C.m}] : \ell_{\text{pc}^{\mathbf{C.m}}}$$

and

$$(\text{BCT}, \Gamma'', F'', S'', Z'', \mathcal{T}'', \mathcal{R}'', K'', i \vdash_t P[\mathbf{C.m}] : \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], \text{pc}^{\mathbf{C.m}+1})}.$$

Therefore $\ell = (\ell_{\text{pc}^{\mathbf{C.m}}} + (\sum \ell_i)^{i \in \text{reachable}(P[\mathbf{C.m}], \text{pc}^{\mathbf{C.m}+1})}) + \ell_\varphi$.

By these hypotheses, it is easy to show the existence of $\Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K'$ and ℓ' such that $\Gamma' \approx H'$ and $P, \text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K' \vdash (\Vdash_{H'} \langle (\text{pc}^{\mathbf{C.m}} + 1, f, o \cdot s) \cdot \varphi, z \rangle_t) : \ell'$. In particular, letting $\text{fields}(\mathbf{D}) = \bar{\mathbf{f}}$, $\Gamma' = \Gamma[o \mapsto (\bar{\mathbf{f}} : \bar{\mathbf{T}}, \mathbf{D})]$ and $S'_{\text{pc}^{\mathbf{C.m}+1}} = o \cdot S''_{\text{pc}^{\mathbf{C.m}}}$.

Here there is a critical point of the proof of subject reduction because it is not evident that $\ell \geq \ell'$. Actually, by **new** typing rule $a = \text{names}(\text{pc}^{\mathbf{C.m}})$, $\Gamma''_{\text{pc}^{\mathbf{C.m}+1}} = \Gamma''_{\text{pc}^{\mathbf{C.m}}}[a \mapsto (\bar{\mathbf{f}} : \bar{\mathbf{T}}, \mathbf{D})]$, and $S''_{\text{pc}^{\mathbf{C.m}+1}} = a \cdot S''_{\text{pc}^{\mathbf{C.m}}}$ and we cannot assume that $a = o$. In particular, since $a = \text{names}(\text{pc}^{\mathbf{C.m}})$, it may be a name that is already “in use” – this is the case when $\mathbf{C.m}$ is either recursive or $\text{pc}^{\mathbf{C.m}}$ is inside an iteration.

There are two cases: (i) \mathbf{D} is not a subclass of **Thread** and (ii) \mathbf{D} is a subclass of **Thread**. In case (i) it easy to verify that $\ell \geq \ell'\{a/o\} \geq \ell'$. In case (ii) we may have $\ell' = (b, c)_a \ \& \ (c, b)_o$. When ℓ' has such a value, we have $\ell'\{a/o\} \not\geq \ell'$ because the name o may appear as index of a dependency pair. In such situation, $(b, c)_a \ \& \ (c, b)_o$ is a circular dependency, while $\ell'\{a/o\} = (b, c)_a \ \& \ (c, b)_a$ is not – see Sections 4.6 and 4.9.5.

To solve this criticality, whenever a new thread is created – either because a `start` is executed or because it is returned by a method invocation, *cf.* the sets \mathcal{T} and \mathcal{R} – we verify whether we are inside a recursion or an iteration, which are the two cases when $names(pc^{c.m})$ may return a name that already exists. In these two cases, the name a is added to the set \mathcal{R} instead of \mathcal{T} . As a consequence, the lam ℓ' will contain terms $RUN((o[\overline{f} : \tau], D))$ while, correspondingly, ℓ will contain terms $RUN((o[\overline{f} : \tau], D))\{a/o\}$. However, since

$$RUN((o[\overline{f} : \tau], D)) = D.run((o[\overline{f} : \tau], D), o, lock_o) \ \& \ (\nu o') RUN((o'[\overline{f} : \tau], D)) ,$$

replacing o in $RUN((o[\overline{f} : \tau], D))$ with any other object name does not change the circularity predicate (because the function is unfolded in a lam that has an invocation of the same function on a fresh name in parallel). Therefore, it is possible to derive $\ell\{a/t\} \succcurlyeq \ell$. \square

4.9.8 Deadlocks and circularities

To conclude the proof of correctness we need to bridge the gap between the notion of deadlock in JVM and the notion of circularity in a lam program.

Definition 9.6. *A JVM configuration $\Vdash_H \langle \varphi_1, z_1 \rangle_{t_1} \cdots \langle \varphi_n, z_n \rangle_{t_n}$ is deadlocked if there are $i_1, \dots, i_k \in 1..n$ such that, for every $j \in \{i_1, \dots, i_k\}$ one of the following holds*

- $\varphi_j = (pc_j, f_j, s_j) \cdot \varphi'_j$ and $P[pc_j] = \mathit{monitorenter}$ and $s_j = a \cdot s'_j$ and $a \in z_h$ with $h \in \{i_1, \dots, i_k\} \setminus j$;
- $\varphi_j = (pc_j, f_j, s_j) \cdot \varphi'_j$ and $P[pc_j] = \mathit{join}$ and $s_j = t_h \cdot s'_j$ with $h \in \{i_1, \dots, i_k\}$.

The following statement is a straightforward consequence of the definitions.

Proposition 4.9.4. *Let $\Vdash_H \mathfrak{C}$ be deadlocked and let $P, BCT, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\Vdash_H \mathfrak{C}) : \ell$. Then ℓ has a circularity.*

Theorem 4.9.5. *Let P be a $JVML_d$ program and $\Gamma \approx H, F \approx_\Gamma f_\perp[0 \mapsto \mathit{main}]$, $S \approx_\Gamma \epsilon$, and $Z \approx_\Gamma \epsilon$. If*

1. $P, BCT, \Gamma, F, S, Z, \mathcal{T}, \mathcal{R}, K \vdash (\Vdash_H \langle (1^{c.main}, f_\perp[0 \mapsto \mathit{main}], \epsilon), \epsilon \rangle_{main}) : \ell$
2. and $\Vdash_H \langle (1^{c.main}, f_\perp[0 \mapsto \mathit{main}], \epsilon), \epsilon \rangle_{main} \rightarrow^* \Vdash_{H'} \mathfrak{C}'$
3. and $\Vdash_{H'} \mathfrak{C}'$ is deadlocked

then ℓ has a circularity.

Proof. The proof is an immediate consequence of the previous results. By 1 and 2 and Theorem 4.9.3 we have the existence of ℓ' and $\Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K'$ such that $\Gamma' \approx H'$ and $P, \text{BCT}, \Gamma', F', S', Z', \mathcal{T}', \mathcal{R}', K' \vdash (\Vdash_{H'} \mathfrak{C}') : \ell'$ and $\ell \succcurlyeq \ell'$. By 3 and Proposition 4.9.4, we have that ℓ' has a circularity. By $\ell \succcurlyeq \ell'$ and Lemma 4.9.1, we have that ℓ has a circularity as well. \square

4.10 Related work

Several techniques have been designed for detecting deadlocks of Java programs. We start by briefly discussing a set of techniques addressing the analysis of Java bytecode. We then discuss separately deadlock analysis techniques using static-time techniques and those requiring either model generation or runtime executions.

4.10.1 Static approaches

Analyses using static time techniques are grouped according to the approach used.

Type Systems

There have been few works based on type systems applied specifically to the deadlock detection in Java. Some of the approaches discussed below do not target this problem in particular, however their techniques, to some extent, tribute to solve some of the problems faced in the development of our solution.

A work that has been key for the development of our solution is [49]. The target of the analysis in this work was not precisely the deadlock verification, instead, this work focused on the verification of well-formedness of the Java bytecode with respect to the locking primitives. The approach taken in JaDA for the design of the behavioral type system and for tackling complex Java features like exception handling has been, to some extent, taken from this work.

Another work that has a lot in common with ours is the deadlock analyzer for ABS programs [35]. In fact, our solution can be seen as the natural extension of this work to the Java language. This work and ours share the same approach based on behavioral type systems, moreover the abstract representations (*lams*) and the subsequent analysis in both works are quite similar. The main differences are introduced according to the necessity of

tackling a diverse concurrency model and also in relation with some of the intrinsic features of the `Java` language.

On the other hand, one of the most complete works regarding deadlock analysis in `Java` with type systems is [19], which defines a type system that derives the order of lock acquisitions in *SafeJava* programs (a subset of `Java with annotations` written as part of the code). Well-typed programs will be verified deadlock free. An extension of this type system for JVMML has been defined in [54].

Our technique differs from this approach in two key aspects. First of all, we aim to a fully automatic tool allowing to analyze existing programs without user intervention. In addition, our types are behavioral. Therefore, deadlock-freedom is not a property of the type system, instead the deadlock detection is performed over the behaviors obtained by the typing process.

Another example of a type system based approach is the work from [25] for detecting race conditions. Like our proposal the type system annotates the access to shared variables and detect read/write contradictions between different threads that may occur concurrently. The downside of this approach is the little support for object-oriented concepts like complex data-types and inheritance.

Data Flow Analysis

A standard approach for the detection of circular dependencies in concurrent programs is based on the construction of an execution flow graph and then search for cycles within this graph. Due to aliasing – two process may access the same shared resource by using a different names –, this construction quickly becomes complex because elements with different names in the graph may refer to the same object. Tools like `JLint` use data-flow analysis to deal with aliasing (`JLint` can be downloaded at <http://jlint.sourceforge.net/>) [12].

Our tool uses an special names generation function for coping with aliases; we also rely on data flow analysis for determining more precisely type of method bodies and method signatures.

Another work using data flow analysis for coping with aliasing is [59]. This technique defines a formal set of rules that control the lock order in their data flow analysis, in similar way to the type system proposed in this thesis. Like our technique this work presents a sound strategy for detecting re-entrant locks, however unlike ours the re-entrance is only detected when the lock expressions involve local variables but not fields. Also, unlike our tool, this technique does not detect circularities with a common prefix – so false positives –, thus leading to a higher number of false positive outcomes.

The work of Nayik [51] defines a set of necessary conditions for the existence of deadlocks, which is then verified by data flow analysis. This is the theory behind the **Chord** tool discussed in Section 5.4. We notice that this tool is reported as one of the most effective ones by [26,51].

Another technique based on this approach is the work of [27]. This technique uses a points-to analysis in combination with a *may happen in parallel* analysis to detect circular dependencies in the code. This technique does not target the Java language directly, instead it targets actor-based concurrency models for which there are some implementations of third party concurrency libraries for Java. Like this technique our solution, ensures that the points of the program involved in dependency cycles, can indeed happen concurrently. This analysis ensures the removal of a considerable number of false positives.

4.10.2 Non-static approaches

Non static techniques usually analyze programs that either have finite models or that have a finite set of relevant inputs. The advantage of these techniques is that it is possible reduce the case explosion and obtain a very precise analysis.

Model Checking

A model checking approach composed of two steps is presented by [47]. A first step generates a so called *trace program* that records the critical concurrent operations and discards non critical parts. In a second step this program is analyzed using an off-the-shelf model checker. Like this technique, our approach tries to get rid of non critical parts of the program. In our case, we use lams, thus allowing to perform the analysis in a reduced but key part of the program, without constraining the program model to be finite.

Monitoring Analysis

Monitoring techniques allow detecting potential deadlocks at runtime. The main idea behind this approach is to analyze only real scenarios. The analysis of circularities in our approach has some similarities to the one presented in [17]. In particular, like this approach, in our one, the locked objects are tagged with the label identifying the threads acquiring the locks. This allows to deal with re-entrance problems in a very clever and simple way. On the other hand the two approaches differ substantially when parallel codes must be detected. Our technique detects the parallel states in the type system,

which is done at static time and, therefore, some over-approximation is introduced to cope with the non-determinism of the scheduler. In [17], being the analysis at runtime, it is possible to tag each segment of the program that is reached by the execution flow, in a manner that these tags specify the exact order of lock acquisitions. The technique in [17] is also interesting because it uses an hybrid strategy for detecting potential deadlocks that might occur because of different scheduler choices (than the current one). The theory of [17] has been prototyped in the `GoodLock` tool used in our benchmark results comparison in Section 5.4. This technique has been refined in `Sherlock`, a tool using symbolic executions [26].

4.11 Conclusions

Deadlocks have proven to be one of the most common flaws in modern concurrent software. Only in the *Oracle Bug Database*, the issue tracking system of the Java Development Kit (JDK), there are currently more than 40 unresolved deadlock problems.

In this chapter we have presented the theoretical tools for the analysis of deadlocks in Java programs. Our approach is based on a behavioral type system for JVM (the Java bytecode) that abstracts object dependencies in concurrent programs. These dependencies are then analyzed by means of a decision algorithm that detects circularities.

There are two features in the Java language that entangle the most the analysis process: the recursive invocations and the creation of fresh objects (that are not part of the original arguments), and in particular, the creation of new threads. The combination of these two features may lead to models with potentially infinite dependencies. The strength of the analysis here proposed is the ability to reduce these models to finite ones.

Existing techniques for the analysis of deadlocks in Java programs are either unsound [14, 51] or target only a subset of the language [54]. In our solution we put special attention to both of these aspects.

We have successfully managed to analyze most of the complex features of Java like: recursive data types, static members, inheritance, arrays structures and exception handling. We have also provided a formal proof of the correctness of the type system and the ulterior analysis.

The correctness of this approach is proven in two parts. The first part focuses in the correctness of the type system. This proof is, as usual, done by demonstrating that the subject reduction theorem holds for the static semantics of the type system. The second part of the proof focuses on demonstrating that the type (and therefore the dependencies) extracted at static

time is preserved during the runtime execution of the program.

Still, for some `Java` programs this typing process may not succeed, static deadlock analysis is after all, undecidable. Our type system may fail, in particular, in presence of data race problems, that do not allow identifying the exact behavior of the program.

There are two major elements of the `Java` concurrency features that remained out of the scope of this analysis: the synchronization mechanisms based on the `wait`, `notify` and `notifyAll` methods; and the components of the `java.util.concurrent` package.

The first of these features has been discussed in this chapter. We noted that the analysis of circularities involving these methods can be dual to the analysis targeting the `synchronize` blocks constructions treated in our solution. We remark that the modularity of our solution could easily allow to extend the analysis process by adding the new verifications like the one proposed to this aim.

Chapter 5

The JaDA tool

Summary

In this chapter we describe the JaDA Tool. This is the prototype implementation of the theory discussed in Chapter 4. We describe in depth the algorithms for the type system inference and the behavioral types analysis. We also review the existing works in the literature addressing these or similar problems. We assess the performance of JaDA against some of these, as well as, against a commercial grade tool. These assessments involve a set of well known and ad-hoc Java and Scala programs. The chapter ends with some final remarks about this solution.

5.1 Introduction

The proposed solution follows the general approach described in Chapter 1. Figure 5.1 particularizes this approach for the case of deadlock analysis in Java programs.

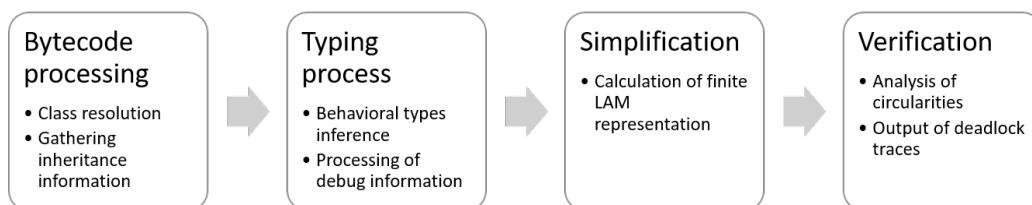


Figure 5.1: Behavioral types based approach for Java deadlock analysis

For the typing process, we have defined an inference system that associates abstract behaviors to Java bytecode instructions. This system consists

of a number of rules which are identical for most of the instructions except for those that have effects in the synchronization process. The rules that need to be treated in a special way are invocations, locks acquisitions and releases, object manipulation and control flow operations. The abstract behaviors extracted from Java bytecode, called *lams* [34, 48], are infinite state models that define dependencies between threads. The authors of [34, 48] demonstrated that the deadlock analysis in *lams* is decidable and proposed an algorithm to this purpose. In this work we extend both the *lams* model and the deadlock detection algorithm in order to cope with some of the Java features.

The inference of the behavioral types and this algorithm constitute the core of a prototype called JaDA – the *Java Deadlock Analyzer* –. The current release of JaDA handles all the main features of Java including threads and synchronizations, constructors, arrays, exceptions, static members, interfaces, inheritance, recursive data types. Few features are not yet covered in the current JaDA version, i.e., `wait-notify-notifyAll` operations, and constructs involving the Java `volatile` modifier.

Chapter contents

The contents of this Chapter correspond to the description of JaDA: the prototype implementation of the theory presented in Chapter 4. Sections 5.2 and 5.3 present the core algorithms behind this tool. The former corresponds to the automatic inference of the behavioral types out of bytecode instructions, while the latter corresponds to the analysis of the circularities within these behaviors. In Section 5.4 we compare the results of JaDA against the results of some of the existing tools, including a commercial grade one. Section 5.5 describes the main features of the tool and its configuration possibilities. Finally, this chapter concludes in Section 5.6 with the main remarks about the JaDA tool, and the possible future lines of work based upon this solution.

5.2 Type inference

The theory described in Chapter 4 relies on program annotations (see rules for method definitions in Figure 4.5 and the function `BCT`). However, the JaDA tools does not require any of these annotation to be made by the programmer. It is able to *infer* types out of the code. The implementation of the type inference mechanism and the further analysis of the resulting behavioral use two iterative processes.

The inference of lams and method types from JVMML programs is an iterative process that, given a set of method types (which include method effects, such as the structure and the updates of the arguments and the returned object, when applicable, the new threads spawned and the synchronized threads), at each step, computes the types of method instructions and the method types. The overall process reiterates again if method types are modified, till an invariant is reached.

In order to infer the right types we use a standard solution that relies on subtypes and relaxes the equality constraint of types of successor instructions in Figures 4.4 and 4.5. We briefly outline this solution.

Record types ϑ are equipped with a binary relation $<$: that is the reflexive and transitive closure of the following pairs:

- $\vartheta <: \top$, and
- $(a[\dots, \mathbf{f}^h : \vartheta, \dots], \mathcal{C}) <: (a[\dots, \mathbf{f}^{h'} : \vartheta', \dots], \mathcal{C}')$, if $\mathcal{C} \subseteq \mathcal{C}'$ and, for all fields, $\mathbf{h} \leq \mathbf{h}'$ and $\vartheta <: \vartheta'$.

(we also have $\perp <: \vartheta$, where \perp has been introduced to deal with inheritance, see Section 4.8). We notice that the constraint $\mathcal{C} \subseteq \mathcal{C}'$ is the identity in JVMML_d, but it is subset relation when inheritance comes into the scene. Clearly, $<$: defines an upper-semi lattice on record types, which allows us to use the operation \sqcup for computing least upper-bounds in our algorithm. The relation $<$: is extended to functions Γ_i , F_i and S_i as follows. Let $\Gamma_i <: \Gamma_i'$ whenever Γ_i and Γ_i' have the same domain and, for every χ , $mk_tree(\Gamma_i, \chi) <: mk_tree(\Gamma_i', \chi)$. Similarly for $F_i <: F_i'$, with the assumption that we have an implicit environment Γ_i for computing the record types of $F_i(x)$ and $F_i'(x)$. As regards stacks, $S_i <: S_i'$ whenever $S_i = \chi_1 \cdots \chi_n$, $S_i' = \chi'_1 \cdots \chi'_n$ (they have the same length) and, for every j , $\Gamma_i(\chi_j) <: \Gamma_i(\chi'_j)$. Therefore, in the typing rules of Figures 4.4 and 4.5, the equality constraints between Γ_i , F_i , and S_i and Γ_{i+1} , F_{i+1} , and S_{i+1} are relaxed in favor of $<$:. For instance $\Gamma_i = \Gamma_{i+1}$ becomes $\Gamma_i <: \Gamma_{i+1}$. We also relax identities $T_i = T_{i+1}$ into $T_i \subseteq T_{i+1}$.

The process typing the instructions uses a data-flow algorithm whose fix point is reached when all constraints are satisfied. The pseudo-code of the algorithm is the following (a visual diagram of the algorithm is shown in Figure 5.2):

```

1  TypeInference(Bct, P, Env_0) -> <Env, L>
2      Env = array[P.size]
3      L = array[P.size]
4      Env[0] = Env_0
5      pending = {0}
6      while( !pending.isEmpty() )
7          index = pending.pop()
8          current = P[index]
9          L[index] = current.type.lam(Bct,Env[index])
10         next_env = current.type.env(Bct,Env[index])
11         if( !current.isReturn() )
12             if( Merge(next_env, current.next, Env) )
13                 pending.push(current.next)
14             if( current.isIf() )
15                 if( Merge(next_env, current.jump, Env) )
16                     pending.push(current.jump)

```

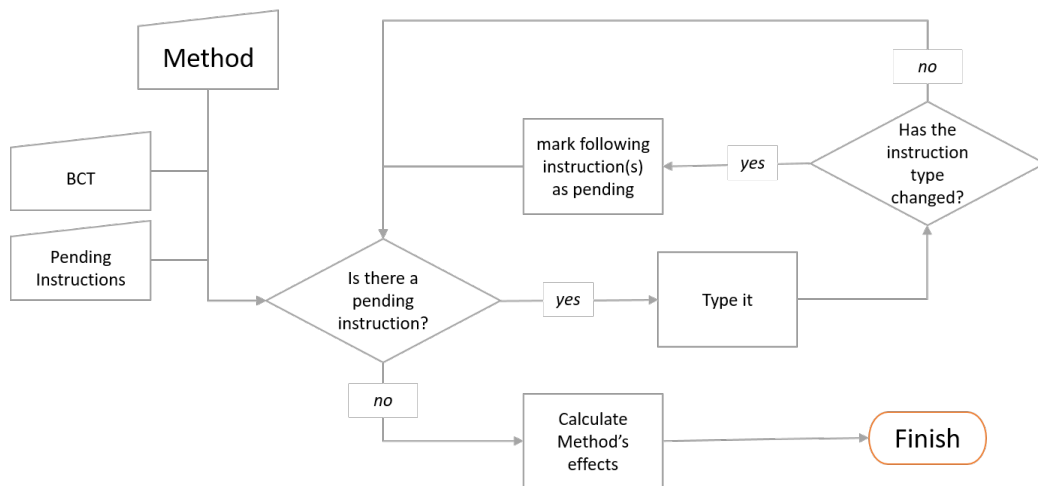


Figure 5.2: Type inference of method's bodies in JaDA.

`TypeInference` starts with the typing environment of the first instruction (see the rules for method definitions in Section 4.5.1) and initializes the set `pending` to `{0}` (only instruction 0 must be typed (see line 5)). The typing environment includes Γ , F , S , Z , T , R , and K of Section 4.5. The type of every instruction is saved in `L` (line 9). After typing the current instruction we get the typing environment for the next instruction; this environment has to match the current typing environment for the next instruction (line 12). If the two typing environments do not match, they are merged using the $<$: relation and the corresponding instruction has to be typed again (line 13). Notice that in a first pass all instructions will be typed and their

typing environment computed, this is because the initial environment for every instruction (except the first one) is `null`. If the current instruction is an *if jump* then the same process is repeated now comparing with the environment at instruction *jump*. The `Merge` function checks if the existing environment is `null`, if so the environment for the current instruction is updated and the method returns true because the instruction needs to be typed again. If an environment for the next instruction already exists then the existing environment is merged with the one just computed one and the function `Merge` returns true again. `Merge` returns false if the existing typing environment remains unchanged.

The process `TypeInference` terminates because, at every iteration, the environment is augmented or remains unchanged (because of the merge operation). In a method body object names are finite (see paragraph *The function names(·)* in Section 4.6). Our prototype does not use recursive records (recursive types are banned in $JVML_d$ and are represented by finite structure in $JVML$, see Section 4.8). Additionally the upper-semi lattice of $\langle \cdot \rangle$ has a finite height. These premises give necessary conditions for the termination of the above process.

The `TypeInference` process assumes the existence of a BCT. Actually, the computation of the BCT is performed by another iterative process that uses `TypeInference`. The following pseudo code highlights the algorithm (a visual representation of the algorithm is shown in Figure 5.3):

```

1  BCT(ClassList) -> <Bct>
2    foreach class in ClassList
3      foreach m in class.Methods
4        Bct[m] = Empty(m)
5
6    changes = true
7    while(changes)
8      changes = false
9      foreach class in ClassList
10       foreach m in class.Methods
11         <Env, L> = TypeInference(Bct, m.P, m.Env_0)
12         changes = changes OR Update(Bct, m, Effects(L))

```

BCT starts by initializing the effects of every method to be empty (line 4). Then every method body is typed with the current `Bct` (line 11) and the typing environments, lams and the method's effects are re-computed from its type (12). This process is repeated until the `Bct` reaches a stable state (line 7). The arguments for the termination of `Bct` are similar to those of `TypeInference`.




Figure 5.3: Type inference of methods' behaviors in JaDA.

5.3 Analysis of circularities

The pseudo code of the algorithm for computing the abstract model of a lam function in Section 4.6 is highlighted here (a visual representation of the algorithm is shown in Figure 5.4):

```

1  ExpandAndCleanCCT(Bct)
2  foreach (method in Bct)
3    State[method] = {}
4    do
5      changed = false
6      foreach( method in Bct )
7        newState = Replace(State, method.type.lam)
8        newState = Normalize(newState)
9        foreach( set in newState )
10         set = Transitive_Closure(set)
11         set = Clean(set)
12         changed = changed OR IsEqual(State[method],newState)
13         State[method] = newState
14     while( changed )
  
```

`ExpandAndCleanCCT` computes the set of states of every method in the BCT. For each method, the function computes its state (line 6) by instantiating the method invocations in its lam with the models computed in the previous iteration (line 7; at the beginning the model is empty, see line 3). Then the state is simplified as a disjunction of conjunctions (the `Normalize` invocation at line 8). That is the model is a set of elements that are sets of dependency pairs. Every set is then transitively closed (line 10) and then simplified by

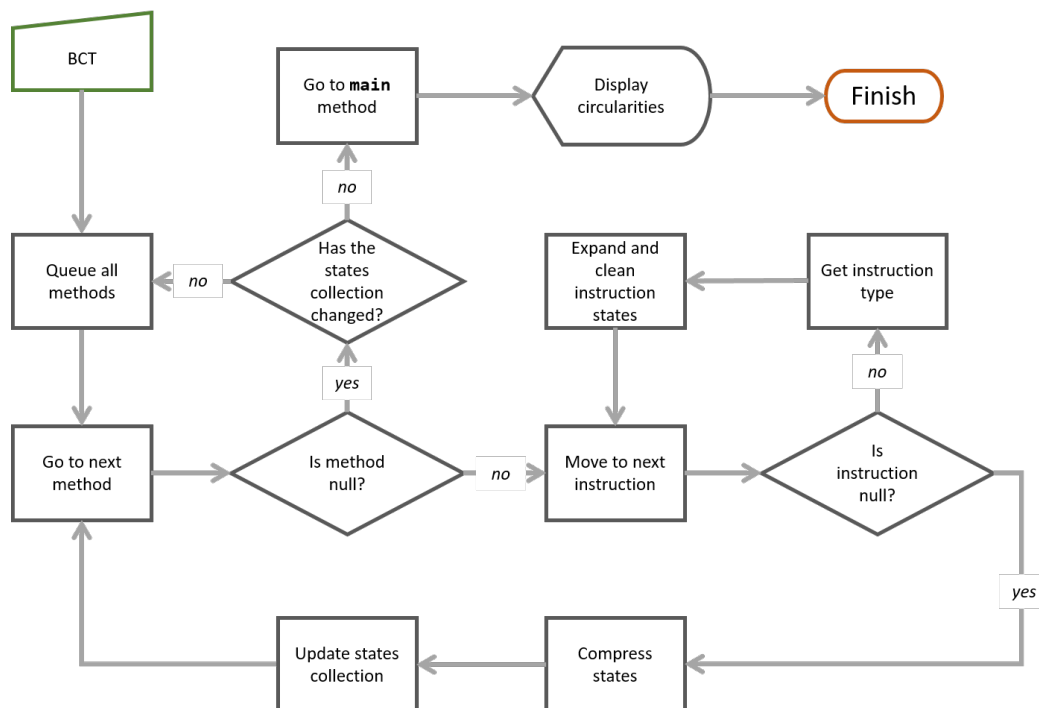


Figure 5.4: JaDA Analysis of behavioral types

means of the `Clean` operation in line 10 . After all method states have been recomputed, we check if the new model is different than the old one (line 12). This process goes on until every method has reach an stable set of final states.

JaDA returns a deadlock if it finds a circularity in the lam of the `main` method. Since the method invocations are stored in the structure implementing the dependency pairs, it is also able to return a trace corresponding to the deadlock.

5.4 Related tools and assessments

In this section we compare JaDA with other deadlock analysis tools for Java. For this comparison we have chosen those that, to the bests of our knowledge, are more powerful tools so far. In Table 5.1, the related tools have been classified according to the type of analysis they perform (see Chapter 4.10 for a discussion about analysis techniques for deadlock detection). We have chosen `Chord` for static analysis [51], `GoodLock` for hybrid analysis [17] and `Sherlock` for dynamic analysis [26]. We have also considered a commercial

static analysis tool, `ThreadSafe`¹ [15].

It is important to remark that we considered in our comparison the `Sherlock` tool in spite of its dynamic nature. The main reason for this is that `Sherlock`'s results provided a number of *verified* deadlocks for each one of these programs. Therefore, we expect the results of static analyzers to output numbers at least greater than the ones given by `Sherlock`. We notice that this is not the case either for `Chord` or `ThreadSafe` (see example `RayTracer`).

Table 5.1: Comparison with different deadlock detection tools. The inner cells show the number of deadlocks detected by each tool. The output labeled “*” are imprecise: see the text.

benchmarks	Static		Hybrid	Dynamic	Commercial
	JaDA	Chord	GoodLock	Sherlock	ThreadSafe
Sor	1	1	7	1	4
Hedc	*	24	23	20	1
Vector	*	3	14	4	0
RayTracer	*	1	8	2	0
MolDyn	*	3	6	1	0
MonteCarlo	*	2	23	2	0
Xalan	*	42	210	9	1
BuildNetwork	3	0			0
Philosophers2	1	0			1
PhilosophersN	3	0			0
StaticFields	1	1			1
ThreadArrays	1	1			1
ThreadArraysWJoins	1	1			0
ScalaSimpleDeadlock	1				
ScalaPhilosophersN	3				

The source of all benchmarks in Table 5.1 is available either at [26, 51] or in the `JaDA-deadlocks` repository². Out of the four chosen tools, we were able to install and effectively test only two of them: `Chord` and `ThreadSafe`; the results corresponding to `GoodLock` and `Sherlock` come from [26]. We also had problems in testing `Chord` with some of the examples in the benchmarks, perhaps due to some miss-configurations, that we were not able to solve even after getting in contact with the author because `Chord` has been discontinued.

The first block belongs to a group of well known Java programs used as benchmarks for several Java analysis tools. In its current state `JaDA` only detects 1 deadlock (and its full trace) in all of the seven analyzed programs

¹<http://www.contemplateld.com/threadsafe>

²<https://github.com/abelunibo/Java-Deadlocks>

from this group. After a manual review we have noticed that the current release of JaDA does not correctly manage arrays of threads. The tool, though, is able to verify the necessary conditions for deadlock existence, but it is not able to produce the correct traces nor individuate the exact number of deadlocks involved. Enhancing the support for this feature is among the top priorities in the development of JaDA. We have, on the other hand, isolated patterns involving arrays of threads, examples `ThreadArrays` and `ThreadArraysWJoins`, and in both case, for this shorter examples the tool outputs the expected results. Moreover, in the case of arrays of threads with `join` operations, we remark that `Chord` and `ThreadSafe` both fail even for a simplified example of this pattern.

In addition, in the case of the `Hedc` program, the concurrency patterns rely on `wait-notify-notifyAll`, which are not supported by JaDA.

The other programs in the second block corresponds to examples designed to test our tool against complex deadlock scenarios like the `Network` program. We notice that both `Chord` and `ThreadSafe` fail to detect those kinds of deadlocks. In addition we report two examples of `Scala` programs [53], these programs have been compiled with the version 2.11 of the `Scala` compiler. We remark that, to the best of our knowledge at the moment of writing this, there were no static deadlock analysis tools for such language. While the results in Table 5.1 are far from ideal, they are encouraging. We hope to improve more convincing ones as soon as JaDA overcomes its current limitations.

5.5 JaDA Deliverable

JaDA has been designed to run on bytecode generated by the `Java` compiler³ and it assumes that the bytecode has been already checked by the `Java` Bytecode Verifier (therefore it does not contain either syntactic or semantic errors). JaDA also requires that every dependency is matched by a corresponding bytecode. Although the bytecode is not executed, JaDA computes every necessary information to solve key issues for the analysis, such as the informations about inheritance. The loading of the existing types is done dynamically in a sand-boxed class loader⁴ to avoid security risks. The full set of dependencies can be specified in JaDA through a *classpath*-like configuration (see property `class-path` in Section 5.5.1).

JaDA is available in three forms: a demo website [31], a command line tool and an Eclipse plug-in (see Figure 5.5). All of them share the same

³We have tested JaDA against the 1.6, 1.7 and 1.8 versions of the `Java` compiler, and against the 1.8 version of the Eclipse Java Compiler (ECJ)

⁴<https://docs.oracle.com/javase/7/docs/api/java/lang/ClassLoader.html>

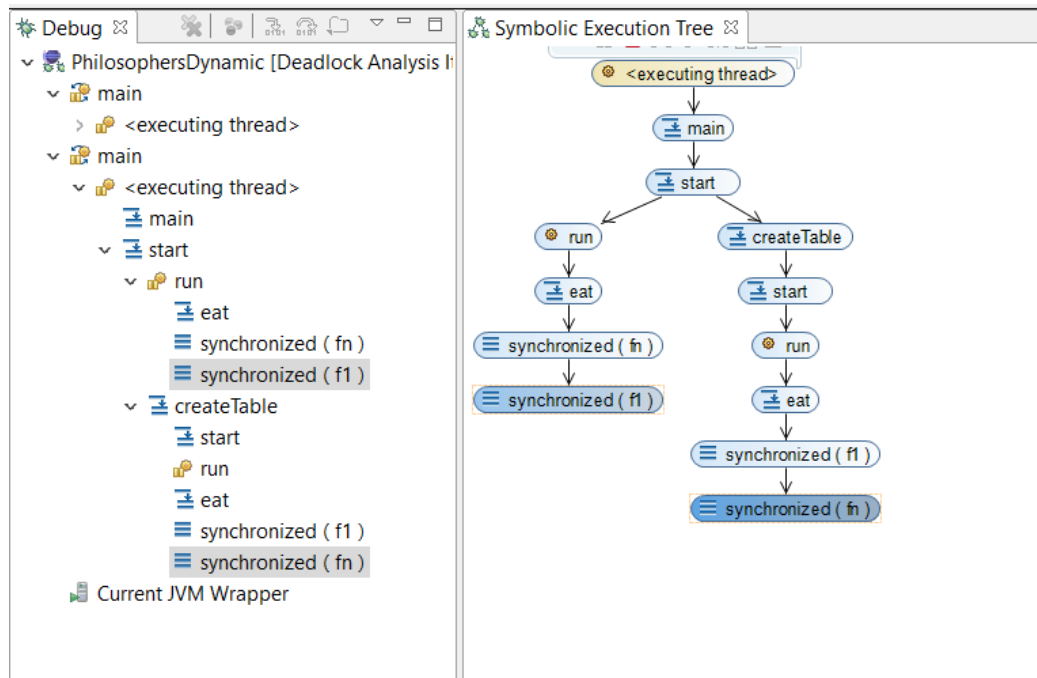


Figure 5.5: Visual representation of a deadlock trace

core: the technique discussed in this work. At the moment of writing, the demo website only allows analyzing single-file programs and to use a subset of the options previously described. The command line tool and the Eclipse plug-in are available through direct requests.

Deadlock trace report

An important result of this solution is the ability to produce the deadlock traces. To this aim the traces of every lock chain have to be maintained at every step of the process. This is allowed by a careful design and some extra work during the type inference process. Since the analysis here presented is static, the resulting traces does not include variable values nor any extra information on the state of the program. It does, however, produce the exact paths involved in the potential deadlock. The Eclipse plug-in, for example, outputs the execution graph causing the deadlock with links to the source code that originates it (see Figure 5.5).

Figure 5.5 shows a deadlock trace representation in the JaDA Eclipse plug-in. This plug-in is the richest delivery form of this prototype, the deadlock trace includes direct links to the source code from where the deadlock is generated. We notice that this is only possible when debugging information

has been compiled into the bytecode.

5.5.1 Customization

The main goal of JaDA is to provide a fully automatic tool for the analysis of deadlocks. We note that even though the user interaction (e.g. through code annotations) may allow gaining in precision, this is a cumbersome task and having a tool completely dependent in this sense may lead to a tool only applicable in non-realistic scenarios.

In order to provide some flexibility, JaDA supports a set of settings to customize the analysis.

<target>: this setting specifies the target file or folder to analyze. It is mandatory. The type of files admitted are: Javaclass files (“`.class`”), Java jar files (“`.jar`”) and compressed zip files (“`.zip`”). In the case of folders, the content of the folder is analyzed recursively.

verbose[=*<value>*]: the value ranges from 1 to 5, the default and more expressive value is 5.

class-path *<classpath>*: Standard Java classpath description. If the target contains dependencies other than those in the standard library, they must be specified via this option.

target-method *<methodName>*: fully qualified target method (should be a void method without arguments). It compels JaDA to analyze the specified method. If this option is not set, the analysis chooses the first `main` method found.

analysis-extent[=*<value>*]: Indicates the extent of the analysis. Possible values are (default value is `classpath`): `full` –analyzes every dependency including the system and classpath-included libraries–, `classpath` –analyzes every library in the classpath–, `custom` –analyzes the classes specified through the property `additional-targets`– and `self` –does not analyze any class but the target one–.

additional-targets *<classes>*: if `analysis-extent` is set to `custom` this property must contain a comma separated list of the fully qualified names of a subset of classes in the `classpath` to include in the analysis. Such a feature is useful for avoiding type known libraries.

custom-types *<file>*: a setting file to specify predefined behavioral types.

`static-constructors[=<value>]`: indicates when the static constructors should be processed, the possibilities are `before-all` and `non-deterministically`. The default option is `before-all`.

Third-party tools involved

The analysis process in JaDA starts with the parsing of the bytecode. This is a cumbersome task because of the length and verbosity of the JVMIL syntax. JaDA relies on the ASM framework [21] for the bytecode extraction and manipulation. (Other third party tools have been also designed for manipulating and analyzing the bytecode: the page <https://java-source.net/open-source/bytecode-libraries> contains a list of existing tools for this purpose. ASM provides a wide set of tools for interacting with the bytecode, including code generation and code analysis. It is also light-weight, open source, very well documented and up to date with the latests versions of Java.

The Eclipse plug-in integration has been inspired in the Symbolic Execution Debugger tool [39].

5.6 Conclusions

In this chapter we have presented the JaDA Tool. It is a static deadlock analysis tool that targets JVMIL, the Java bytecode. Therefore, it supports the analysis compiled Java programs, as well as, programs written in other languages that are also compiled in JVMIL, like Scala. In fact, although there are several tools that target Java, this is to the best of our knowledge the first static deadlock analyzer applied to Scala programs.

The technique underlying JaDA uses a behavioral type system that abstracts the main features of the programs with respect to the concurrent operations. These types are then analyzed in order to verify the presence of circular dependencies that can potentially lead to deadlock states. The solution here described has been successfully applied to another programming language: ABS [35]

JaDA is designed to run in a fully automatic fashion, meaning that the inference of the program type and the subsequent analysis could be done unassisted. Nevertheless, user intervention is possible and may enhance the precision of the analysis, for example in presence of native methods. The possibilities to configure JaDA are listed in Section 5.5. In particular, one of the main advantage of this approach is the possibility to provide, manually,

the behavioral types of methods that cannot be typed automatically, like native methods or methods relying on reflection.

Even though the tool is still under development, we have been able to assess it by analyzing a set of `Java` and `Scala` programs. This contribution also reports a comparison between `JaDA`'s results and those of existing deadlock analysis tools, among which is a commercial grade one. The assessments have proved that our tool excels in the analysis of recursive programs that a priori may lead to infinite state models, like the case of the example in Section 4.3.3.

Chapter 6

General conclusions

Summary

We conclude this work with the main remarks about the techniques that has been so far described. We also discuss future lines of research derived from the solutions here presented.

6.1 Conclusions

This thesis describes an approach based on the use of behavioral types for the static analysis of computer programs. We have successfully applied this technique in two different problems, in two different programming languages.

Behavioral type systems constitute a powerful way to reason abstractly in the program effects. The main characteristic of this approach is the modularity, these type systems are built compositionally from the level of program instructions to the level of more general programming units like methods or classes. This poses several advantages, for example, the possibility to focus during the development phase on subsets of the programming language and extend this gradually afterwards to support all the language functionalities. Moreover, this modularity is extended also to the overall system. In this case, the properties targeted by the analysis are not properties of the type system, like it could be, for example, the type safety in standard type systems. Instead, behavioral type systems produce an abstraction of the program that remove spurious information for the analysis. This abstraction can be analyzed afterwards, in a second stage. Having the analysis decoupled from the type system allows to combine several analyzers, but more importantly, eases the demonstration of the correctness of the whole solution.

The demonstration of the correctness of these techniques is usually done

in two steps. A first step, ensures that: (i) the behavior obtained for the program at static time is preserved during the program execution; and (ii) the execution does not introduce at runtime information that was not considered in the program abstraction. This proof is usually done by means of a standard subject reduction theorem demonstrated by induction on the program runtime evolution. In a second part, one can demonstrate the correctness of the analyzer, or even delegate such analysis to other certificated third party tools.

The main contribution of this work is the application of this technique to two problems: the analysis of virtual machines usage in dynamic Cloud Computing scenarios, and the analysis of deadlocks in Java. In both cases we have respectively developed prototype implementations: The **SRA** tool and The **JaDA** tool, that allow to interact with the theory here proposed.

Both tools can be considered novelties in their particular field, at least to some extent. For the case of **SRA**, this is the first automatic analyzer that considers Cloud related resource usage. The language targeted by this analysis is an object-oriented concurrent language that allows the allocation and deallocations of objects representing virtual machines. Moreover, these virtual machines are in turn, active objects, that may run asynchronous operations that also create and release other virtual machine objects. This combination of settings has not been, to the best of our knowledge, considered by any of the existing similar tools. For the case of the **JaDA** tool, this is one of the few analyzers capable of targeting a very wide set of the **Java** language features. It is though, the first of these that proposes a sound approach without impacting in the precision of the analysis. It is also the first analyzer capable of detecting deadlocks in **Scala**, a high end programming language that is also compiled to **Java** bytecode.

6.2 Future work

The results obtained by these two analyzers have been compared with other scientific or commercial tools. In both of the problems tackled, our tools proved competitive, state-of-the-art level results. However, this experience shows that there is still room for improvement, which leaves us with some concrete future lines of work:

SRA tool enhancements

Removing technical restrictions. The restrictions described in Section 2.4 are technical simplifications and do not constitute particular limitations to

the theory behind. However, these can be alleviated by retaining more expressive notations for the effect of a method, *i.e.* by considering the methods' possible releases as a set of sets instead of a simple set. Such a notation is more suited for modeling nondeterministic behaviors and it might be made even more expressive by tagging the methods' side effects with conditions specifying when each effect is yielded. Clearly, the management of these domains becomes more complex and the trade-off between simplicity and expressiveness of the behavioral types must be carefully evaluated.

Integrating new automatic solvers. The existing cost equations solvers tend to specialize in particular scenarios. A good alternative to improve the SRA tool could be to integrate other (more powerful) solvers, such as those based on theorem provers or in multivariate cost analysis. This extension could definitely enrich the domain of programming patterns that can be covered by this analyzer.

Targeting higher-end languages. A big limitation of the SRA tool is the targeted language. This reduces the possibility to analyze *real life* programs. Targeting high end programming languages like Java or C (as the C4B analyzer) would give the SRA tool a higher practical value. Another, good step in this direction is to target a programming language with a formal model as ABS [46], of which `vm1` is a very basic sub-calculus.

In particular, the latter of these lines of work may be of special interest. Behavioral types are abstract specifications of programs that highlight resource usages. One of the main applications of these abstract specifications is to bridge the gap between programs and performance constraints in Service Level Agreement (SLA, in short) documents [32]. As discussed in [32], a language like ABS is intended to be a language able to interact, for example, with the Amazon Elastic Cloud Computing features library. Therefore, this language might be used to model Cloud services, where a primary issue is the compatibility with the SLA document. The technique conveyed here would allow (i) extracting the abstract specifications by means of the type (inference) system, (ii) computing the cost by means of the solver, and (iii) verifying whether the computed costs complies with the SLA or not. This solution demonstrates that these steps are all feasible and correct.

JaDA tool enhancements

Full Java support. A logical first step in this direction is to extend JaDA to cover the remaining Java features with respect to the concurrency patterns.

As described so far, there are some limitations on the programs that can be analyzed by this tool. The main ones are: the analysis of the high level concurrency support introduced with the `java.util.concurrent` package of the Java standard library; the analysis of the synchronization mechanisms provided by the methods `wait`, `notify` and `notifyAll` of the class `Object`; and the analysis of the instructions involving the Java `volatile` keyword.

Enhance the data race problems detection. In addition to deadlocks, another issue closely related to concurrent programming is the data race problem. The behavioral type here described currently keeps track of the access to shared objects from different threads, see *Record Types* in Section 4.4. At this moment, when such scenario is detected, the typing process does not succeed. The JaDA tool reports an inconclusive deadlock to ensure the soundness of the output. However, a much more informative output would report the details of the data race condition.

Applying this approach in other high end languages. One of the main contributions of this work has been, without doubts, the application of this technique to a high-end language like Java. In this case, the tool targets specifically the Java bytecode, JVMIL, which allows analyzing also another sophisticated language like Scala. Like Java, many common programming languages compile to intermediate stack based semantics bytecode. The technique described in this work for the type inference in JVMIL might be, with more or less effort, adapted to match some of these languages. A good first choice in this sense would be the IL language which is the compilation target of languages in the .Net platform like C#, VB, F#. We remark that thanks to the modularity of this approach, the analysis here described can be re-used without any modifications.

Bibliography

- [1] Wolfgang Ahrendt, Bernhard Beckert, Daniel Bruns, Richard Bubel, Christoph Gladisch, Sarah Grebing, Reiner Hähnle, Martin Hentschel, Mihai Herda, Vladimir Klebanov, Wojciech Mostowski, Christoph Scheben, Peter H. Schmitt, and Mattias Ulbrich. The key platform for verification and analysis of java programs. In *Verified Software: Theories, Tools and Experiments - 6th International Conference, VSTTE 2014, Vienna, Austria, July 17-18, 2014, Revised Selected Papers*, pages 55–71, 2014.
- [2] Elvira Albert, Puri Arenas, Jesús Correas, Samir Genaim, Miguel Gómez-Zamalloa, Germán Puebla, and Guillermo Román-Díez. Object-sensitive cost analysis for concurrent objects. *Softw. Test. Verif. Reliab.*, 25(3):218–271, May 2015.
- [3] 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.
- [4] Elvira Albert, Puri Arenas, Samir Genaim, Germán Puebla, and Damiano Zanardini. Cost analysis of java bytecode. In *Programming Languages and Systems, 16th European Symposium on Programming, ESOP 2007, Held as Part of the Joint European Conferences on Theory and Practics of Software, ETAPS 2007, Braga, Portugal, March 24 - April 1, 2007, Proceedings*, pages 157–172, 2007.
- [5] Elvira Albert, Puri Arenas, Samir Genaim, German Puebla, and Damiano Zanardini. COSTA: design and implementation of a cost and termination analyzer for java bytecode. In *Formal Methods for Components and Objects, 6th International Symposium, FMCO 2007, Amsterdam, The Netherlands, October 24-26, 2007, Revised Lectures*, pages 113–132, 2007.

- [6] 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.
- [7] Elvira Albert, Jesús Correas, and Guillermo Román-Díez. Peak cost analysis of distributed systems. In *Proceedings of SAS 2014*, volume 8723 of *Lecture Notes in Computer Science*, pages 18–33. Springer, 2014.
- [8] Elvira Albert, Jesús Correas, and Guillermo Román-Díez. Non-Cumulative Resource Analysis. In *Proceedings of TACAS 2015*, Lecture Notes in Computer Science, pages 85–100. Springer, 2015.
- [9] Elvira Albert, Samir Genaim, and Miguel Gómez-Zamalloa. Parametric inference of memory requirements for garbage collected languages. *SIGPLAN Not.*, 45(8):121–130, 2010.
- [10] Diego Esteban Alonso-Blas and Samir Genaim. On the limits of the classical approach to cost analysis. In *Proceedings of SAS 2012*, volume 7460 of *Lecture Notes in Computer Science*, pages 405–421. Springer, 2012.
- [11] Davide Ancona, Viviana Bono, Mario Bravetti, Joana Campos, Giuseppe Castagna, Pierre-Malo Deniérou, Simon J. Gay, Nils Gesbert, Elena Giachino, Raymond Hu, Einar Broch Johnsen, Francisco Martins, Viviana Mascardi, Fabrizio Montesi, Rumyana Neykova, Nicholas Ng, Luca Padovani, Vasco T. Vasconcelos, and Nobuko Yoshida. Behavioral types in programming languages. *Foundations and Trends in Programming Languages*, 3(2-3):95–230, 2016.
- [12] Cyrille Artho and Armin Biere. Applying static analysis to large-scale, multi-threaded java programs. In *13th Australian Software Engineering Conference (ASWEC 2001)*, pages 68–75, 2001.
- [13] Robert Atkey. *Amortised Resource Analysis with Separation Logic*. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010.
- [14] Robert Atkey and Donald Sannella. Threadsafe: Static analysis for java concurrency. *ECEASST*, 72, 2015.
- [15] Robert Atkey and Donald Sannella. Threadsafe: Static analysis for java concurrency. *ECEASST*, 72, 2015.
- [16] 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.
- [17] Saddek Bensalem and Klaus Havelund. Dynamic deadlock analysis of multi-threaded programs. In *in Hardware and Software Verification and Testing*, volume 3875 of *Lecture Notes in Computer Science*, pages 208–223. Springer, 2005.
- [18] Leonard Berman. The complexity of logical theories. *Theoretical Computer Science*, 11(1):71 – 77, 1980.
- [19] Chandrasekhar Boyapati, Robert Lee, and Martin C. Rinard. Ownership types for safe programming: preventing data races and deadlocks. In *Proceedings of OOPSLA 2002*, pages 211–230. ACM, 2002.
- [20] 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.
- [21] Eric Bruneton. Asm 4.0 a java bytecode engineering library. <http://download.forge.objectweb.org/asm/asm4-guide.pdf>. Last accessed: 2016-12-03.
- [22] Quentin Carbonneaux, Jan Hoffmann, and Zhong Shao. Compositional certified resource bounds. In *Proceedings of the 36th PLDI Conference*, pages 467–478. ACM, 2015.
- [23] Wei-ngan Chin, Huu Hai Nguyen, Shengchao Qin, and Martin Rinard. Memory usage verification for oo programs. In *Proceedings of SAS 2005*, volume 3672 of *Lecture Notes in Computer Science*, pages 70–86. Springer, 2005.
- [24] Haskell B Curry, Robert Feys, and William Craig. Combinatory logic. I, 1958. North-Holland, Amsterdam.
- [25] Mohamed A El-Zawawy and Hamada A Nayel. Type systems based data race detector. *Computer and Information Science*, 5(4):53, 2012.
- [26] Mahdi Eslamimehr and Jens Palsberg. Sherlock: scalable deadlock detection for concurrent programs. In *Proceedings of the 22nd International Symposium on Foundations of Software Engineering (FSE-22)*, pages 353–365. ACM, 2014.

- [27] Antonio Flores-Montoya, Elvira Albert, and Samir Genaim. May-happen-in-parallel based deadlock analysis for concurrent objects. In *Formal Techniques for Distributed Systems - Joint IFIP WG 6.1 International Conference, FMOODS/FORTE 2013, Held as Part of the 8th International Federated Conference on Distributed Computing Techniques, DisCoTec 2013, Florence, Italy, June 3-5, 2013. Proceedings*, pages 273–288, 2013.
- [28] Antonio Flores Montoya and Reiner 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.
- [29] Stephen N. Freund and John C. Mitchell. A type system for the java bytecode language and verifier. *J. Autom. Reasoning*, 30(3-4):271–321, 2003.
- [30] Abel Garcia and Cosimo Laneve. Static analyzer of resource usage upper bounds, 2015.
- [31] Abel Garcia and Cosimo Laneve. JaDA – the Java Deadlock Analyzer. Available at JaDA.cs.unibo.it from October 28, 2016.
- [32] Elena Giachino, Stijn de Gouw, Cosimo Laneve, and Behrooz Nobakht. Statically and dynamically verifiable sla metrics, 2016. To appear on *Lecture Notes in Computer Science*.
- [33] Elena Giachino, Einar Broch Johnsen, Cosimo Laneve, and Ka I Pun. Time complexity of concurrent programs - - A technique based on behavioural types -. In *Formal Aspects of Component Software - 12th International Conference, FACS 2015, Niterói, Brazil, October 14-16, 2015, Revised Selected Papers*, pages 199–216, 2015.
- [34] Elena Giachino, Naoki Kobayashi, and Cosimo Laneve. Deadlock analysis of unbounded process networks. In *Proceedings of 25th International Conference on Concurrency Theory CONCUR 2014*, volume 8704 of *Lecture Notes in Computer Science*, pages 63–77. Springer, 2014.
- [35] Elena Giachino, Cosimo Laneve, and Michael Lienhardt. A framework for deadlock detection in ABS. *Software and Systems Modeling*, 2015.
- [36] James Gosling, William N. Joy, and Guy L. Steele Jr. *The Java Language Specification*. Addison-Wesley, 1996.

- [37] 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.
- [38] Arie Gurfinkel and Sagar Chaki. Combining predicate and numeric abstraction for software model checking. *STTT*, 12(6):409–427, 2010.
- [39] Martin Hentschel, Richard Bubel, and Reiner Hähnle. Symbolic execution debugger (SED). In *Runtime Verification - 5th International Conference, RV 2014, Toronto, ON, Canada, September 22-25, 2014. Proceedings*, pages 255–262, 2014.
- [40] Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. Multivariate amortized resource analysis. In *Proceedings of POPL 2011*, pages 357–370. ACM, 2011.
- [41] Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. Multivariate amortized resource analysis. *ACM Trans. Program. Lang. Syst.*, 34(3):14, 2012.
- [42] Jan Hoffmann and Zhong Shao. Automatic static cost analysis for parallel programs, 2015. [Online; accessed 11-February-2015].
- [43] Martin Hofmann and Steffen Jost. Static prediction of heap space usage for first-order functional programs. In *Proceedings of POPL 2003*, pages 185–197. ACM, 2003.
- [44] Martin Hofmann and Steffen Jost. Type-based amortised heap-space analysis. In *Proceedings of ESOP 2006*, volume 3924 of *Lecture Notes in Computer Science*, pages 22–37. Springer, 2006.
- [45] Martin Hofmann and Dulma Rodriguez. Efficient type-checking for amortised heap-space analysis. In *Proceedings of CSL 2009*, volume 5771 of *Lecture Notes in Computer Science*, pages 317–331. Springer, 2009.
- [46] Einar Broch Johnsen, Reiner Hähnle, Jan Schäfer, Rudolf Schlatte, and Martin Steffen. ABS: A core language for abstract behavioral specification. In *Proceedings of FMCO 2010*, volume 6957 of *Lecture Notes in Computer Science*, pages 142–164. Springer, 2011.
- [47] Pallavi Joshi, Mayur Naik, Koushik Sen, and David Gay. An effective dynamic analysis for detecting generalized deadlocks. In *Proceedings of the 18th International Symposium on Foundations of Software Engineering*, pages 327–336. ACM, 2010.

- [48] Naoki Kobayashi and Cosimo Laneve. Deadlock analysis of unbounded process networks. Available at the journal website, To Appear.
- [49] Cosimo Laneve. A type system for JVM threads. *Theoretical Computer Science*, 290(1):741 – 778, 2003.
- [50] Sajee Mathew. Overview of amazon web services. at d0.awsstatic.com/whitepapers/aws-overview.pdf, November 2014.
- [51] Mayur Naik, Chang-Seo Park, Koushik Sen, and David Gay. Effective static deadlock detection. In *31st International Conference on Software Engineering (ICSE 2009)*, pages 386–396. ACM, 2009.
- [52] Jorge Navas, Mario Méndez-Lojo, and Manuel V Hermenegildo. Customizable resource usage analysis for java bytecode. Submitted, 2008.
- [53] Martin Odersky and al. An Overview of the Scala Programming Language. Technical Report IC/2004/64, EPFL, Lausanne, Switzerland, 2004.
- [54] Pratibha Permandla, Michael Roberson, and Chandrasekhar Boyapati. A type system for preventing data races and deadlocks in the Java Virtual Machine Language: 1. In *Proceedings of the 2007 Conference on Languages, Compilers, and Tools for Embedded Systems (LCTES'07)*, page 10. ACM, 2007.
- [55] Raymie Stata and Martin Abadi. A type system for Java bytecode sub-routines. *ACM Transactions on Programming Languages and Systems*, 21(1):90–137, January 1999.
- [56] Robert Endre Tarjan. Amortized computational complexity. *SIAM Journal on Algebraic Discrete Methods*, 6(2):306–318, 1985.
- [57] Alan Mathison Turing. On computable numbers, with an application to the entscheidungsproblem. *J. of Math*, 58(345-363):5, 1936.
- [58] Ben Wegbreit. Mechanical program analysis. *Communications of the ACM*, 18(9):528–539, 1975.
- [59] Amy Williams, William Thies, and Michael D. Ernst. Static deadlock detection for java libraries. In *19th European Conference on Object-Oriented Programming (ECOOP 2005)*, volume 3586 of *Lecture Notes in Computer Science*, pages 602–629. Springer, 2005.