

# Towards a Smart Compilation Manager for Java<sup>\*</sup> (Extended Abstract)

Giovanni Lagorio

DISI - Università di Genova  
Via Dodecaneso, 35, 16146 Genova (Italy)  
lagorio@disi.unige.it

**Abstract.** It is often infeasible to recompile all the sources an application consists of each time a change is made. Yet, a recompilation strategy which does not guarantee the same outcome of an entire recompilation is not useful: why wasting time in debugging a program (a set of `.class` files in the Java case) which might behave *differently* from the program obtained recompiling all the sources from scratch?

We say that a compilation strategy is *sound* if it recompiles, besides the changed sources, all the unchanged sources whose new binary, produced by the overall recompilation, would differ from the existing one (if any) *and* all the sources for which the recompilation would be undefined: indeed, when the entire compilation fails, so should do the partial recompilation.

We say that a compilation strategy is *minimal* if it never recompiles an unchanged source whose new binary would be equal to the existing one. In this paper we present a compilation strategy for a substantial subset of Java which is proved to be *sound* and *minimal*.

## 1 Introduction

When dealing with large applications it is infeasible to recompile all the sources each time a change is made. Of course, separate compilation is the answer to such a problem, but a key point has to be considered: in addition to the modified sources, which files have to be recompiled as well? A recompilation strategy which does not guarantee the same outcome of an entire recompilation is not useful: why wasting time in debugging a program (a set of `.class` files in the Java case) which might behave *differently* from the program obtained recompiling all the sources from scratch? It is known that, using the most common Java compilers (including the standard one), a program which is the result of a successful recompilation may throw linking related exceptions at runtime even if these errors could have been detected at compile time [3]. This is not in contrast

---

<sup>\*</sup> Partially supported by Dynamic Assembly, Reconfiguration Type-checking - EC project IST-2001-33477, APPSEM II - Thematic network IST-2001-38957, and Murst NAPOLI - Network Aware Programming: Oggetti, Linguaggi, Implementazioni.

with the soundness for Java investigated by other papers [6, 7] because they did not take separate compilation into account.

Two contrasting requirements have to be considered: on the one hand re-compilations can be rather expensive (in time), hence they should be avoided when possible. More precisely, they are useless when the recompilation of an unchanged (with respect to the previous compilation) source fragment  $S$ , whose corresponding binary fragment  $B$  is already present, would produce a binary equal to  $B$ . On the other hand, a recompilation strategy which saves time not recompiling a fragment  $S$ , with a corresponding binary fragment  $B$ , whose recompilation would produce a new binary  $B'$  different from  $B$ , could cost a lot of wasted time in debugging an *inconsistent* application, that is, an application that *cannot* be rebuilt by recompiling all the sources.

Albeit some Java IDEs support smart or incremental compilation, to our knowledge there are no publications which explain in detail the inner working of such recompilation strategies. In this paper we chose to analyze Java, a mainstream language, and model its peculiar features, because our final goal is to implement a compilation manager, for a widespread language, whose correctness and minimality (in the sense explained below) can be formally proved.

We say that a compilation strategy is *sound* if it recompiles all the changed sources and the unchanged sources whose new binary, produced by the overall recompilation, would differ from the existing one (if any) *and* all the unchanged sources for which the recompilation would be undefined. This latter requirement is very important: indeed, when the entire recompilation is not defined, so should be the partial recompilation.

Of course, a strategy which recompiles all the sources each time a change is made is trivially sound and, obviously, totally useless in practice. We say that a compilation strategy is *minimal* if it never recompiles an unchanged fragment whose new binary would be equal to the existing one.

In this paper we present a compilation strategy for a substantial subset of Java which is proved to be *sound* and *minimal*.

Section 2 presents the formal framework, Section 3 explains the ideas behind our compilation strategy and Section 4 discusses related and future work.

## 2 Formalization

### 2.1 The Language

In this paper we model a substantial subset of Java at both source (Figure 1) and binary (Figure 2) level. Our model of bytecode is rather abstract: it is basically source code enriched with some annotations (discussed below). However, the same source level expression can be compiled to different binary level expressions, as it happens in Java. For instance, a method invocation  $x.m()$  can be translated to a virtual method invocation or an interface method invocation depending on the static type of  $x$ .

With the exception of arrays and inner-classes we model all the major features of Java: classes (including abstract classes), interfaces, primitive types, access

```

S ::= AM CK class C extends C' implements I1...In { KDSs FDSs MDSs } |
    AM interface I extends I1...In { FDSs MDSs }
AM ::= public | protected | ε | private
CK ::= ε | abstract
KDSs ::= KD1s... KDns
FDSs ::= FD1s... FDns
MDSs ::= MD1s... MDns
KDs ::= AM KH { super(E1s,...,Ens); STMTSs }
FDs ::= AM FINAL FK T f = Es ;
MDs ::= AM MK MH { STMTSs return Es; } | AM abstract MH ;
KH ::= (T1 x1,...,Tn xn) throws ES
MK ::= ε | static | abstract
MH ::= T m(T1 x1,...,Tn xn) throws ES
FINAL ::= ε | final
FK ::= ε | static
T ::= RT | int | bool
RT ::= C | I
ES ::= C1,...,Cn
Es ::= PRIMARYs|ASSIGNs|N|true|false
PRIMARYs ::= null | this | NEWs | x | INVOKEs | super.f | PRIMARYs.f | RT.f
ASSIGNs ::= x = Es | PRIMARYs.f = Es | super.f = Es | RT.f = Es
NEWs ::= new C(E1s,...,Ens)
INVOKEs ::= PRIMARYs.m(E1s,...,Ens) | super.m(E1s,...,Ens) | RT.m(E1s,...,Ens)
STMTSs ::= STMT1s...STMTns
STMTs ::= {STMTSs} | SEs ; | if (Es) STMT1s else STMT2s | while (Es) STMTs |
    try {STMTSs} CATCHESs finally { STMTS1s } | throw Es
SEs ::= ASSIGNs | INVOKEs | NEWs
CATCHESs ::= CATCH1s...CATCHns
CATCHs ::= catch (C x) { STMTSs }

```

*Assumptions:*

- interface names in S are distinct;
- field names in FDS<sup>s</sup> are distinct;
- method/constructor signatures in MDS<sup>s</sup>/KDS<sup>s</sup> are distinct;
- parameter and exception names in both KH and MH are distinct;
- class names in CATCHES<sup>s</sup> are distinct.

**Fig. 1.** Syntax - Sources

modifiers (including packages, but without the `import` directive), constructors, (instance/static) fields (both in classes and interfaces), (instance/static/abstract) methods, super field accesses and method invocations, exceptions. The treatment of arrays and inner-classes would complicate the model without apparently giving further insights.

Figure 1 gives the syntax of the language. A source fragment  $S$  can be a class declaration or an interface declaration. In the former case it consists of: an access modifier  $AM$ , a class kind  $CK$  (either  $\epsilon$  or `abstract`), the name of the class, the name of the superclass, the list of the implemented interfaces and the declaration of constructors, fields and methods. Analogously, an interface declaration consists of an access modifier, the name of the interface, the name of the superinterfaces and the declaration of fields and methods.

A constructor declaration  $KD^s$  consists of an access modifier  $AM$ , a constructor header  $KH$ , the invocation of a superclass's constructor<sup>1</sup> and a sequence of statements  $STMTS^s$ . A constructor header consists of the sequence of parameters and the exception specification  $ES$ . In this paper we assume for simplicity that any class can be an exception, that is, we do not model the predefined class `Throwable`. So, an exception specification is just a sequence of class names.

A field declaration  $FD^s$  consists of an access modifier  $AM$ , an optional modifier `FINAL`, a field kind  $FK$ , a type  $T$ , the name of the field  $f$  and the initialization expression  $E^s$ .

A method declaration  $MD^s$  can be either concrete or abstract. In the former case it consists of an access modifier  $AM$ , a method kind  $MK$ , a method header  $MH$ , a sequence of statements  $STMTS^s$  and a return expression  $E^s$ . In the latter case it just consists of an access modifier  $AM$ , the keyword `abstract` and a method header. A type  $T$  can be a reference type  $RT$  or a primitive type (`int` or `bool`). We distinguish between class names  $C$  and interface names  $I$  for clarity only, even though they actually range over the same set of names.

An expression  $E^s$  can be: a primary expression, an assignment expression, an integer literal  $N$  or a boolean literal. Some expressions,  $SE^s$ , can be used as statements; they are: assignment  $ASSIGN^s$ , method invocation  $INVOKE^s$  and instance creation  $NEW^s$ .

While Java permits accessing a static member of a class/interface  $RT$  via both the type name  $RT$  or any expression which has static type  $RT$ , here we allow only the former kind of access (because allowing both kinds of access would require additional, uninteresting, typing rules).

Figure 2 gives the syntax of the binary language. As already said, it mostly mimics the source language, except for it is enriched with some annotations enclosed between “ $\ll$ ” and “ $\gg$ ”.

For example, the instance creation expression  $NEW^s$  is translated to  $NEW^b$ , which contains, as annotation, the tuple of types describing the constructor which has been found as most specific at compile time. Analogously, method

---

<sup>1</sup> Invocations of a constructor of the same class (using `this`) are not considered since they are simply syntactic shortcuts (recursive invocations are not allowed - see 8.8.5 of [8]).

```

B ::= AM CK class C extends C' implements I1 ... In { KDSb FDSb MDSb } |
    AM interface I extends I1 ... In { FDSb MDSb }
KDSb ::= KD1b ... KDnb
MDSb ::= MD1b ... MDnb
FDSb ::= FD1b ... FDnb
KDb ::= AM KH { super(E1b, ..., Enb) << T̄ >>c; STMTSb }
FDb ::= AM FINAL FK T f = Eb
MDb ::= AM MK MH { STMTSb return Eb; } | AM abstract MH ;
Eb ::= PRIMARYb | ASSIGNb | N | true | false
PRIMARYb ::= null | this | NEWb | x | INVOKEb
    PRIMARYb. << C.T >>if f | << RT.T >>sf f
ASSIGNb ::= x = Eb | PRIMARYb. << C.T >>if f = Eb | << RT.T >>sf f = Eb
NEWb ::= new C << T̄ >>c (E1b, ..., Enb)
INVOKEb ::= PRIMARYb. << C.m(T̄)T >>vrt m(E1b, ..., Enb) | this. << C.m(T̄)T >>spr m(E1b, ..., Enb) |
    PRIMARYb. << C.m(T̄)T >>stt m(E1b, ..., Enb) | PRIMARYb. << I.m(T̄)T >>int m(E1b, ..., Enb)
STMTSb ::= STMT1b ... STMTnb
STMTb ::= { STMTSb } | SEb ; | if (Eb) STMT1b else STMT2b | while (Eb) STMTb
    try { STMTSb } CATCHESb finally { STMTS1b } | throw Es
SEb ::= ASSIGNb | INVOKEb | NEWb
CATCHESb ::= CATCH1b ... CATCHnb
CATCHb ::= catch (C x) { STMTSb }

```

*Assumptions:*

- interface names in B are distinct;
- field names in FDS<sup>b</sup> are distinct;
- method/constructor signatures in MDS<sup>b</sup>/KDS<sup>b</sup> are distinct;
- class names in CATCHES<sup>b</sup> are distinct.

**Fig. 2.** Syntax - Binaries

invocation expressions are annotated with the signature of the most specific method found at compile time and the static type of the receiver. There are four kinds of (binary) method invocation expressions INVOKE<sup>b</sup>: virtual (instance method invocation), super (invocation via **super**), static and interface.

## 2.2 Type Environments

Type environments  $\Gamma$  are defined in Figure 3. A type assignment  $\gamma$  maps a class/interface name to its type.

The assignment  $C \mapsto [\text{AM}=\text{AM}, \text{CK}=\text{CK}, \text{PARENT}=\text{C}', \text{IS}=\text{I}_1 \dots \text{I}_n, \text{KSS}=\text{KSS}, \text{FSS}=\text{FSS}, \text{MSS}=\text{MSS}]$  has the meaning “the class C has access modifier AM and kind CK, extends C', implements  $\text{I}_1 \dots \text{I}_n$  and has constructor signatures KSS, field signatures FSS and method signatures MSS”. Analogously,  $I \mapsto [\text{AM}=\text{AM}, \text{IS}=\text{I}_1 \dots \text{I}_n, \text{FSS}=\text{FSS}, \text{MSS}=\text{MSS}]$  has the meaning “the interface I has access modifier AM, extends  $\text{I}_1 \dots \text{I}_n$ , and has field signatures FSS and method signatures MSS”.

$\Gamma ::= \gamma_1 \dots \gamma_n$	
$\gamma ::= \mathbf{C} \mapsto [\text{AM}=\text{AM}, \text{CK}=\text{CK}, \text{PARENT}=\mathbf{C}', \text{IS}=\mathbf{I}_1 \dots \mathbf{I}_n, \text{KSS}=\text{KSS}, \text{FSS}=\text{FSS}, \text{MSS}=\text{MSS}] \mid$	
$\mathbf{I} \mapsto [\text{AM}=\text{AM}, \text{IS}=\mathbf{I}_1 \dots \mathbf{I}_n, \text{FSS}=\text{FSS}, \text{MSS}=\text{MSS}]$	
$\bar{\mathbf{T}} ::= \mathbf{T}_1 \dots \mathbf{T}_n$	
$\mathbf{T}^\perp ::= \mathbf{T} \mid \perp$	$\bar{\mathbf{T}}^\perp ::= \mathbf{T}_1^\perp \dots \mathbf{T}_n^\perp$
$\text{KS} ::= \text{AM } \bar{\mathbf{T}} \text{ throws ES}$	$\text{KSS} ::= \text{KS}_1 \dots \text{KS}_n$
$\text{FS} ::= \text{AM FK T f}$	$\text{FSS} ::= \text{FS}_1 \dots \text{FS}_n$
$\text{MS} ::= \text{AM MK T m}(\bar{\mathbf{T}}) \text{ throws ES}$	$\text{MSS} ::= \text{MS}_1 \dots \text{MS}_n$
$\lambda ::= \mathbf{T} \leq \mathbf{T}' \mid$	
$\text{RT} <_1 \text{RT}' \mid$	
$\text{RT} \not\prec \exists \mathbf{T} \mid$	
$\text{RT} \not\prec \exists_{\mathbf{C}} \text{CK}^* \mathbf{C} \mid$	
$\text{RT} \not\prec \exists_{\mathbf{I}} \mathbf{I} \mid$	
$\text{RT} \not\prec \text{Cns}(\mathbf{C}, \bar{\mathbf{T}}^\perp) = [\text{PAR}=\bar{\mathbf{T}}, \text{ES}=\text{ES}] \mid$	
$\text{RT} \not\prec \text{Fld}(\text{RT}', \mathbf{f}) = [\text{FINAL}=\text{FINAL}^*, \text{FK}=\text{FK}, \text{T}=\mathbf{T}] \mid$	
$\text{RT} \not\prec \text{Mth}(\text{RT}', \mathbf{m}, \bar{\mathbf{T}}^\perp) = [\text{MK}=\text{MK}^*, \text{RET}=\mathbf{T}, \text{PAR}=\bar{\mathbf{T}}, \text{ES}=\text{ES}]$	
$\text{CK}^* ::= \text{CK} \mid \_$	$\text{FINAL}^* ::= \text{FINAL} \mid \_$
$\text{MK}^* ::= \text{MK} \mid \text{not-static}$	

**Fig. 3.** Type environments and Type assumptions

Type assumptions  $\lambda$ , also defined in Figure 3, describe fine-grained requirements; they are:

- $\mathbf{T} \leq \mathbf{T}'$  with the meaning “ $\mathbf{T}$  is a subtype of  $\mathbf{T}'$ ”;
- $\text{RT} <_1 \text{RT}'$  with the meaning “ $\text{RT}$  *directly* extends  $\text{RT}'$ ”;
- $\text{RT} \not\prec \exists \mathbf{T}$  with the meaning “type  $\mathbf{T}$  exists and is accessible from code contained in type  $\text{RT}$ ”<sup>2</sup>;
- $\text{RT} \not\prec \exists_{\mathbf{C}} \text{CK}^* \mathbf{C}$  with the meaning “class  $\mathbf{C}$ , with kind  $\text{CK}^*$ , exists and is accessible from code contained in  $\text{RT}$ ”.  $\text{CK}^* = \_$  means “any kind”, that is, we do not care whether the class is abstract or not;
- $\text{RT} \not\prec \exists_{\mathbf{I}} \mathbf{I}$  with the meaning “interface  $\mathbf{I}$  exists and is accessible from code contained in  $\text{RT}$ ”;
- $\text{RT} \not\prec \text{Cns}(\mathbf{C}, \bar{\mathbf{T}}^\perp) = [\text{PAR}=\bar{\mathbf{T}}, \text{ES}=\text{ES}]$  with the meaning “the most specific constructor for class  $\mathbf{C}$  and parameter types  $\bar{\mathbf{T}}^\perp$ , invoked from code contained in  $\text{RT}$ , has parameter types  $\bar{\mathbf{T}}$  and can throw exceptions which are compatible with the exception specification  $\text{ES}$ ”<sup>3</sup>;
- $\text{RT} \not\prec \text{Fld}(\text{RT}', \mathbf{f}) = [\text{FINAL}=\text{FINAL}^*, \text{FK}=\text{FK}, \text{T}=\mathbf{T}]$  with the meaning “if code contained in  $\text{RT}$  looks up a field named  $\mathbf{f}$  in type  $\text{RT}'$ , then it finds a field with

<sup>2</sup> If you think the symbol “ $\not\prec$ ” as an eye, then you can interpret any assumption of the form “ $\text{RT} \not\prec \dots$ ” as: “ $\text{RT}$  *sees*  $\dots$ ” and this is supposed to help ☺.

<sup>3</sup> That is, any exception  $\mathbf{C}$  such that a  $\mathbf{C}$ ’s superclass is contained in  $\text{ES}$  - see 11.2 of [8].

- a final modifier  $\text{FINAL}^*$ , kind  $\text{FK}$  and type  $\text{T}$ ".  $\text{FINAL}^* = \_$  means we do not care whether the field is final or not;
- $\text{RT} \not\prec \text{Mth}(\text{RT}', \text{m}, \bar{\text{T}}^\perp) = [\text{MK}=\text{MK}^*, \text{RET}=\text{T}, \text{PAR}=\bar{\text{T}}, \text{ES}=\text{ES}]$  with the meaning “the most specific method, invoked from code contained in  $\text{RT}$ , for a method named  $\text{m}$ , with parameter types  $\bar{\text{T}}^\perp$  on a receiver with static type  $\text{RT}'$  is a method which has kind  $\text{MK}^*$ , return type  $\text{T}$ , parameter types  $\bar{\text{T}}$  and can throw exceptions which are compatible with the exception specification  $\text{ES}$ ”<sup>3</sup>.

These assumptions can be thought as “the minimal pieces of information” needed to compile a certain source to a certain bytecode, as we will detail in the sequel. The rules defining the judgment  $\Gamma \vdash \lambda$  are omitted for lack of space. A forthcoming extended work, containing all the rules, will be available at <http://www.disi.unige.it/person/LagorioG/publications.html> shortly. Also, the rules for a small subset of Java can be found in [2].

Next subsection shows how and when these assumptions are used in the process of compilation, while Section 3 explain how to exploit type assumptions in order to obtain a recompilation strategy which is both sound and minimal.

### 2.3 Compilation

Compilation of expressions is expressed by the following judgment:

$$\text{RT}; \text{II}; \text{ES}; \Gamma \vdash \text{E}^s \rightsquigarrow \text{E}^b : \text{T}$$

with the meaning “expression  $\text{E}^s$  has type  $\text{T}$  and compiles to binary expression  $\text{E}^b$  when contained in type  $\text{RT}$ , in a local environment  $\text{II}$ , in a context where exceptions  $\text{ES}$  can be thrown and in a type environment  $\Gamma$ ”. Type  $\text{RT}$  is needed to model the access control; for instance,  $\text{RT}$ ’s `private` methods can be invoked only by expressions inside  $\text{RT}$ . The local environment  $\text{II}$  maps parameter names and `this` to their respective types (`this` is undefined when typing expressions contained in static contexts). Figure 4 show some selected rules defining this judgment.

Compilation of statements is expressed by the following judgment:

$$\text{RT}; \text{II}; \text{ES}; \Gamma \vdash \text{STMT}^s \rightsquigarrow \text{STMT}^b$$

with the meaning “statement  $\text{STMT}^s$  is compiled to  $\text{STMT}^b$  when contained in type  $\text{RT}$ , in a local environment  $\text{II}$ , in a context where exceptions  $\text{ES}$  can be thrown and in a type environment  $\Gamma$ ”. The rules defining this judgment are omitted because of lack of space.

Before describing the compilation of fragments, we need to introduce the notion of *compilation environment*. A compilation environment  $ce$  maps fragment names to the corresponding fragment.

$$ce : \text{RT} \rightarrow \text{S} \cup \text{B}$$

Note that  $ce$  models a compilation environment from a compiler’s point of view; that is, for each fragment name  $ce$  returns either a source or a binary

$\begin{array}{l} \forall i \in 1..n \text{ RT}; II; \text{ES}; \Gamma \vdash E_i^s \rightsquigarrow E_i^b : T_i \\ \Gamma \vdash \text{RT} \not\prec \text{Cns}(\text{C}, T_1 \dots T_n) = [\text{PAR}=\bar{\text{T}}, \text{ES}=\text{ES}] \\ \Gamma \vdash \text{RT} \not\prec \exists c \in \text{C} \end{array}$	
$\text{RT}; II; \text{ES}; \Gamma \vdash \text{new C}(E_1^s, \dots, E_n^s) \rightsquigarrow \text{new C} \ll \bar{\text{T}} \gg_c (E_1^b, \dots, E_n^b) : \text{C}$	
$\frac{\begin{array}{l} \Gamma \vdash \text{C} <_1 \text{C}' \\ \Gamma \vdash \text{C} \not\prec \text{Fld}(\text{C}', \text{f}) = [\text{FINAL}=\_, \text{FK}=\epsilon, \text{T}=\text{T}] \end{array}}{\text{C}; II; \text{ES}; \Gamma \vdash \text{super.f} \rightsquigarrow \text{this} \ll \text{C}'.\text{T} \gg_{\text{if}} : \text{T}} \quad \text{this} \in \text{Def}(II)$	
$\frac{\begin{array}{l} \text{RT}; II; \text{ES}; \Gamma \vdash E_1^s \rightsquigarrow E_1^b : \text{C} \\ \Gamma \vdash \text{RT} \not\prec \text{Fld}(\text{C}, \text{f}) = [\text{FINAL}=\epsilon, \text{FK}=\epsilon, \text{T}=\text{T}] \\ \text{RT}; II; \text{ES}; \Gamma \vdash E_2^s \rightsquigarrow E_2^b : \text{T}_2 \\ \Gamma \vdash \text{T}_2 \leq \text{T} \end{array}}{\text{RT}; II; \text{ES}; \Gamma \vdash E_1^s.\text{f} = E_2^s \rightsquigarrow E_1^b \ll \text{C}.\text{T} \gg_{\text{if}} = E_2^b : \text{T}}$	$\frac{\Gamma \vdash \text{RT} \not\prec \text{Fld}(\text{RT}', \text{f}) = [\text{FINAL}=\_, \text{FK}=\text{static}, \text{T}=\text{T}]}{\text{RT}; II; \text{ES}; \Gamma \vdash \text{RT}'.\text{f} \rightsquigarrow \ll \text{RT}'.\text{T} \gg_{\text{sf}} : \text{T}}$
$\frac{\begin{array}{l} \text{RT}; II; \text{ES}; \Gamma \vdash E^s \rightsquigarrow E^b : \text{RT}' \\ \Gamma \vdash \text{RT} \not\prec \exists c \text{ - RT}' \\ \Gamma \vdash \text{RT} \not\prec \text{Mth}(\text{RT}', \text{m}, T_1 \dots T_n) = [\text{MK}=\text{not-static}, \text{RET}=\text{T}, \text{PAR}=\bar{\text{T}}, \text{ES}=\text{ES}] \\ \forall i \in 1..n \text{ RT}; II; \text{ES}; \Gamma \vdash E_i^s \rightsquigarrow E_i^b : T_i \end{array}}{\text{RT}; II; \text{ES}; \Gamma \vdash E^s.\text{m}(E_1^s, \dots, E_n^s) \rightsquigarrow E^b \ll \text{RT}'.\text{m}(\bar{\text{T}}) \gg_{\text{vrt}} (E_1^b, \dots, E_n^b) : \text{T}}$	
$\frac{\begin{array}{l} \text{RT}; II; \text{ES}; \Gamma \vdash E^s \rightsquigarrow E^b : \text{RT}' \\ \Gamma \vdash \text{RT} \not\prec \exists_1 \text{RT}' \\ \Gamma \vdash \text{RT} \not\prec \text{Mth}(\text{RT}', \text{m}, T_1 \dots T_n) = [\text{MK}=\text{abstract}, \text{RET}=\text{T}, \text{PAR}=\bar{\text{T}}, \text{ES}=\text{ES}] \\ \forall i \in 1..n \text{ RT}; II; \text{ES}; \Gamma \vdash E_i^s \rightsquigarrow E_i^b : T_i \end{array}}{\text{RT}; II; \text{ES}; \Gamma \vdash E^s.\text{m}(E_1^s, \dots, E_n^s) \rightsquigarrow E^b \ll \text{RT}'.\text{m}(\bar{\text{T}}) \gg_{\text{int}} (E_1^b, \dots, E_n^b) : \text{T}}$	

Fig. 4. Selected expression typing rules

fragment, but not both. In fact, even if both are present, only one is considered by the compiler, usually the most up-to-date according to the file's attributes. As a consequence, a type environment  $\Gamma$  can be extracted from a compilation environment  $ce$  by disregarding the code while retaining signatures and information about type hierarchy. Let us assume the function  $extractEnv : ce \rightarrow \Gamma$  does this job.

Assume we have a *consistent* compilation environment  $ce$ , that is, if we compile all the sources in  $ce$  we obtain, for the sources which have been re-compiled, the same binaries already present in  $ce$ . Assume, then, to change a bunch of sources, say  $\text{RT}_1, \dots, \text{RT}_n$ , obtaining  $ce_{\text{new}}$  (and a corresponding  $\Gamma_{\text{new}} = extractEnv(ce_{\text{new}})$ ). What do we have to do in order to obtain a new *consistent* compilation environment?

First of all, we have to check that the environment  $\Gamma_{\text{new}}$  is well-formed, for instance, it must not contain a cycle in the type hierarchy. We formally capture this notion with the judgment  $\vdash \Gamma_{\text{new}} \diamond$ . One could argue that we do not want



to check the well-formedness of the *whole*  $\Gamma_{\text{new}}$  after having changed, say, a *tiny detail* in one source. Right, we do not want to; yet, we need to *unless* we know something more (see Section 3).

After having checked the environment  $\Gamma_{\text{new}}$ , we need to decide which fragments to (re)compile. Of course we have to recompile those which have been changed, that is,  $\text{RT}_1 \dots \text{RT}_n$  but what else? Are you tempted to answer “`make clean; make all` will do”? A lot of real-world programmers would do just that. When dealing with a small/medium application and a quite powerful computer you can do that. And, it works. It works *fine* actually, but what if you *can not* do that? Next section addresses this crucial point. For now, just assume we have to compile  $\text{RT}_1, \dots, \text{RT}_m$  (obviously  $m \geq n$ ).

When we have a well-formed  $\Gamma_{\text{new}}$  and do know which fragments to recompile we can go for it: in our model “running the compilation” on each fragment  $\text{RT}_i$  amounts to prove the following judgments:

$$\text{RT}_i; \Gamma_{\text{new}} \vdash ce_{\text{new}}(\text{RT}_i) \rightsquigarrow B_i$$

Figure 5 shows some selected rules defining this judgment. In defining the compilation of a set of classes we assume to compile them one by one, that is, we assume that no global optimizations take place. This reflects the fact that in languages with dynamic linking like Java, the concept of “program” is only significant at runtime so it is safer to leave cross class optimizations to virtual machines like HotSpot [9].

### 3 A Smart Strategy

When dealing with an updated compilation context  $ce_{\text{new}}$  (in respect to a previous  $ce_{\text{old}}$ ) there are two steps to perform: checking whether the corresponding new type environment  $\Gamma_{\text{new}}$  is well-formed and, when it is, decide which (unchanged) source fragments have to be recompiled besides the changed ones. Some information gathered during a previous compilation can be used to speedup both these steps. We first show how we can check only the “updated part” of an environment  $\Gamma_{\text{new}}$  when a previous well-formed environment  $\Gamma_{\text{old}}$  is known. Then, we show how the type assumptions used to compile a fragment can be used later to decide whether it has to be recompiled.

We define  $\text{leaves}_\Gamma(\text{RT}) = \{\text{RT}' \mid \Gamma \vdash \text{RT}' \leq \text{RT} \wedge \forall \text{RT}'' \Gamma \vdash \text{RT}'' \leq \text{RT}' \implies \text{RT}'' = \text{RT}'\}$  and the judgment  $\Gamma \vdash \text{okOvr RT}$  with the meaning “RT correctly extends its parent types (up to Object) in  $\Gamma$ ”. That is, RT’s hierarchy is acyclic and the Java rules on method overriding/hiding are respected. The rules defining such a judgment are omitted for lack of space.

**Definition 1.** A type environment  $\Gamma_{\text{new}}$  is well-formed w.r.t. another type environment  $\Gamma_{\text{old}}$  iff the following conditions hold:

$$[\text{add}] \text{RT} \in \text{Def}(\Gamma_{\text{new}}) \setminus \text{Def}(\Gamma_{\text{old}}) \implies \begin{cases} \Gamma_{\text{new}} \vdash \text{okOvr RT} \\ \text{used}_{\Gamma_{\text{new}}}(\text{RT}) \subseteq \text{Def}(\Gamma_{\text{new}}) \end{cases}$$

These metarules assume  $\Gamma$  to be well-formed, that is,  $\vdash \Gamma \diamond$ .

Anyway, we do *not* want  $\vdash \Gamma \diamond$  to be a premise of any of them (see the text for full details).

$$\begin{array}{c}
\frac{\text{C}; \Gamma \vdash \text{KDS}^s \rightsquigarrow \text{KDS}^b \quad \text{C}; \Gamma \vdash \text{FDS}^s \rightsquigarrow \text{FDS}^b \quad \text{C}; \Gamma \vdash \text{MDS}^s \rightsquigarrow \text{MDS}^b}{\Gamma \vdash \text{AM CK class C extends C' implements } I_1, \dots, I_m \{ \text{KDS}^s \text{ FDS}^s \text{ MDS}^s \} \rightsquigarrow \text{AM CK class C extends C' implements } I_1, \dots, I_m \{ \text{KDS}^b \text{ FDS}^b \text{ MDS}^b \}} \quad \text{AM} \in \{\epsilon, \text{public}\} \\
\\
\frac{\text{I}; \Gamma \vdash \text{FDS}^s \rightsquigarrow \text{FDS}^b \quad \text{I}; \Gamma \vdash \text{MDS}^s \rightsquigarrow \text{MDS}^b}{\Gamma \vdash \text{AM interface I extends } I_1 \dots I_m \{ \text{FDS}^s \text{ MDS}^s \} \rightsquigarrow \text{AM interface I extends } I_1 \dots I_m \{ \text{FDS}^b \text{ MDS}^b \}} \quad \begin{array}{l} \text{MDS}^s = \text{MD}_1^s \dots \text{MD}_n^s \\ \forall i \in 1..n \text{ MD}_i^s = \text{public abstract} \dots \\ \text{FDS}^s = \text{FD}_1^s \dots \text{FD}_m^s \\ \forall i \in 1..m \text{ FD}_i^s = \text{public static final} \dots \\ \text{AM} \in \{\epsilon, \text{public}\} \end{array} \\
\\
\frac{\begin{array}{l} \Gamma \vdash \text{C} <_1 \text{C}' \\ \forall i \in 1..n \text{ C}; \text{H}; \text{ES}; \Gamma \vdash \text{E}_i^s \rightsquigarrow \text{E}_i^b : \text{T}_i \\ \text{C}; \text{H}; \text{ES}; \Gamma \vdash \text{STMTS}^s \rightsquigarrow \text{STMTS}^b \\ \Gamma \vdash \text{C} \not\prec \text{Cns}(\text{C}', \text{T}_1 \dots \text{T}_n) = [\text{PAR}=\bar{\text{T}}, \text{ES}=\text{ES}] \end{array}}{\text{C}; \Gamma \vdash \text{AM } (\text{T}_1 \text{ x}_1, \dots, \text{T}_n \text{ x}_n) \text{ throws ES } \{ \text{super}(\text{E}_1^s, \dots, \text{E}_n^s); \text{STMTS}^s \} \rightsquigarrow \text{AM } (\text{T}_1 \text{ x}_1, \dots, \text{T}_n \text{ x}_n) \text{ throws ES } \{ \text{super}(\text{E}_1^b, \dots, \text{E}_n^b) \ll \bar{\text{T}} \gg_c; \text{STMTS}^b \}} \quad \text{H} = \{ x_1 \mapsto \text{T}_1, \dots, x_n \mapsto \text{T}_n, \text{this} \mapsto \text{C} \} \\
\\
\frac{\begin{array}{l} \text{RT}; \text{H}; \emptyset; \Gamma \vdash \text{E}^s \rightsquigarrow \text{E}^b : \text{T}' \\ \Gamma \vdash \text{T}' \leq \text{T} \end{array}}{\text{RT}; \Gamma \vdash \text{AM FINAL FK T f} = \text{E}^s ; \rightsquigarrow \text{AM FINAL FK T f} = \text{E}^b ;} \quad \text{H} = \text{This}(\text{FK}, \text{C}) \\
\\
\frac{\begin{array}{l} \text{C}; \text{H}; \text{ES}; \Gamma \vdash \text{STMTS}^s \rightsquigarrow \text{STMTS}^b \\ \text{C}; \text{H}; \text{ES}; \Gamma \vdash \text{E}^s \rightsquigarrow \text{E}^b : \text{T}' \\ \Gamma \vdash \text{T}' \leq \text{T} \end{array}}{\text{C}; \Gamma \vdash \text{AM MK T m}(\text{T}_1 \text{ x}_1, \dots, \text{T}_n \text{ x}_n) \text{ throws ES } \{ \text{STMTS}^s \text{ return } \text{E}^s; \} \rightsquigarrow \text{AM MK T m}(\text{T}_1 \text{ x}_1, \dots, \text{T}_n \text{ x}_n) \text{ throws ES } \{ \text{STMTS}^b \text{ return } \text{E}^b; \}} \quad \begin{array}{l} \text{H} = \{ x_1 \mapsto \text{T}_1, \dots, x_n \mapsto \text{T}_n \} \cup \text{This}(\text{MK}, \text{C}) \\ \text{MK} \neq \text{abstract} \end{array} \\
\\
\text{This}(\_, \text{I}) = \emptyset \\
\text{This}(\text{static}, \text{C}) = \emptyset; \\
\text{This}(\epsilon, \text{C}) = \{ \text{this} \mapsto \text{C} \}
\end{array}$$

**Fig. 5.** Selected compilation rules

$$\begin{array}{l}
[rmv] \text{ } Def(\Gamma_{old}) \setminus Def(\Gamma_{new}) \neq \emptyset \implies \forall RT \in Def(\Gamma_{new}) \text{ used}_{\Gamma_{new}}(RT) \subseteq Def(\Gamma_{new}) \\
[cng] \text{ } RT \in Def(\Gamma_{old}) \cap Def(\Gamma_{new}), \Gamma_{old}(RT) \neq \Gamma_{new}(RT) \implies \\
\quad \left\{ \begin{array}{l} \forall RT' \in \text{leaves}_{\Gamma_{new}}(RT) \Gamma_{new} \vdash \text{okOvr } RT' \\ \text{used}_{\Gamma_{new}}(RT) \subseteq Def(\Gamma_{new}) \end{array} \right.
\end{array}$$

where “used by RT in  $\Gamma$ ”,  $\text{used}_{\Gamma}(RT)$ , means all types directly referenced by RT.

**Theorem 1.** *If  $\vdash \Gamma_{old} \diamond$  holds, and  $\Gamma_{new}$  is well-formed w.r.t.  $\Gamma_{old}$  then,  $\vdash \Gamma_{new} \diamond$  holds.*

*Proof* Two requirements have to be met:

- $\Gamma_{new}$  must be closed, that is,  $\text{used}_{\Gamma_{new}}(RT) \subseteq Def(\Gamma_{new})$  for all  $RT \in Def(\Gamma_{new})$ ;
- overriding rules must be satisfied, that is, for any  $RT \in Def(\Gamma_{new})$  the judgment  $\Gamma_{new} \vdash \text{okOvr } RT$  must be valid.

The former requirement is met by hypothesis when some type defined in  $\Gamma_{old}$  has been removed from  $\Gamma_{new}$ , see *[rmv]* in Definition 1. When no types have been removed, the unchanged types cannot, trivially, refer to undefined types because, by hypothesis,  $\Gamma_{old}$  is closed. By *[add]* and *[cng]* of Definition 1 new and updated types refer only to types defined in  $\Gamma_{new}$ .

The latter requirement can be proved by case analysis. Consider  $Def(\Gamma_{new})$  as the union of three disjoint sets:  $U$ ,  $C$  and  $N$ . These sets contain, respectively, the unchanged, changed and new types in  $\Gamma_{new}$  with respect to  $\Gamma_{old}$ . Formally:

$$\begin{aligned}
U &= \{RT \mid RT \in Def(\Gamma_{new}) \cap Def(\Gamma_{old}), \Gamma_{new}(RT) = \Gamma_{old}(RT)\} \\
C &= \{RT \mid RT \in Def(\Gamma_{new}) \cap Def(\Gamma_{old}), \Gamma_{new}(RT) \neq \Gamma_{old}(RT)\} \\
N &= Def(\Gamma_{new}) \setminus Def(\Gamma_{old})
\end{aligned}$$

Let us consider  $N \cup C$  first. For any new type RT, contained in  $N$ , the judgment  $\Gamma_{new} \vdash \text{okOvr } RT$  is valid by the hypothesis, see *[add]* of Definition 1. For any changed type RT, contained in  $C$ , the judgment  $\Gamma_{new} \vdash \text{okOvr } RT$  is valid because of *[cng]* of Definition 1.

It remains to prove that the judgment holds for the unchanged types, contained in  $U$ . Since they are unchanged the *direct* supertypes of any type in  $U$  are the same in  $\Gamma_{old}$  and  $\Gamma_{new}$ . Furthermore, each direct supertype of RT must be contained in  $Def(\Gamma_{old})$  and in  $Def(\Gamma_{new})$  because of, respectively, the fact that  $\vdash \Gamma_{old} \diamond$  holds and the fact that  $\Gamma_{new}$  is closed (which we have proved before). Hence, these direct supertypes must be contained in  $Def(\Gamma_{old}) \cap Def(\Gamma_{new})$  which, by definition, is equal to:  $U \cup C$ . If a direct supertype is in  $U$ , then the same reasoning can be applied; so, for any  $RT \in U$ , only two cases are possible:

- *all* supertypes of RT are in  $U$ ; then, the hierarchy of RT has not changed and by the hypothesis  $\vdash \Gamma_{old} \diamond$  is valid, so the judgment  $\Gamma_{new} \vdash \text{okOvr } RT$  is valid too;
- there exists a supertype  $RT'$  of RT which is in  $C$ , whose subtypes till RT are in  $U$ . Then, by *[cng]* there exists a type  $RT''$  such that  $RT'' \leq RT$  and  $\Gamma_{new} \vdash \text{okOvr } RT''$  holds. So,  $\Gamma_{new} \vdash \text{okOvr } RT$  must be valid too.  $\square$

In summary, as long as we keep trace of a previous well-formed environment  $\Gamma_{\text{old}}$  we can check for the well-formedness of any new environment  $\Gamma_{\text{new}}$  by examining the changes w.r.t.  $\Gamma_{\text{old}}$ . The most expensive check must be performed when some classes are removed; this is acceptable because classes are added/changed more often than removed in the usual software development cycle. The requirement of having a previous well-formed environment available may seem restrictive, but it is not, since one can always use the empty environment (which is trivially well-formed) as  $\Gamma_{\text{old}}$  when starting to use our strategy from scratch.

Assume we have a compilation environment  $ce_{\text{old}}$ , a corresponding well-formed type environment  $\Gamma_{\text{old}} = \text{extractEnv}(ce_{\text{old}})$ , and we can prove:

$$\text{RT}; \Gamma_{\text{old}} \vdash ce_{\text{old}}(\text{RT}) \rightsquigarrow B$$

The proof tree for this judgment can be proved to be unique, and contains a set of assumptions  $A = \lambda_1, \dots, \lambda_n$ . We call it the *requirements* for RT in  $\Gamma_{\text{old}}$  and write

$$\text{Reqs}(\text{RT}, \Gamma_{\text{old}}) = A$$

Assume to change  $ce_{\text{old}}$  leaving the fragment for RT untouched. That is, assume to have another compilation environment  $ce_{\text{new}}$ , with a corresponding well-formed environment  $\Gamma_{\text{new}}$ , such that  $ce_{\text{new}}(\text{RT}) = ce_{\text{old}}(\text{RT})$ .

Does RT compile in  $ce_{\text{new}}$ ? If it does, can we say something about the corresponding binary?

The following theorem states that, if  $\Gamma_{\text{new}}$  still satisfy the requirements of RT in  $\Gamma_{\text{old}}$ , then we do know that RT compiles in  $\Gamma_{\text{new}}$  and it compiles to the *same* binary.

**Theorem 2.** *If  $ce_{\text{old}}$  and  $ce_{\text{new}}$  are two compilation environments which share the same source for RT, that is,  $ce_{\text{old}}(\text{RT}) = ce_{\text{new}}(\text{RT})$ , the corresponding type environments  $\Gamma_{\text{old}} = \text{extractEnv}(ce_{\text{old}})$  and  $\Gamma_{\text{new}} = \text{extractEnv}(ce_{\text{new}})$  are well-formed, and the judgment  $\text{RT}; \Gamma_{\text{old}} \vdash ce_{\text{old}}(\text{RT}) \rightsquigarrow B$  can be proved, then:*

$$\text{RT}; \Gamma_{\text{new}} \vdash ce_{\text{new}}(\text{RT}) \rightsquigarrow B \iff \Gamma_{\text{new}} \vdash \text{Reqs}(\text{RT}, \Gamma_{\text{old}})$$

*Proof (sketch)* If  $\Gamma_{\text{new}} \vdash \text{Reqs}(\text{RT}, \Gamma_{\text{old}})$ , that is,  $\Gamma_{\text{new}} \vdash \lambda$  holds for all  $\lambda \in \text{Reqs}(\text{RT}, \Gamma_{\text{old}})$ , the proof for  $\text{RT}; \Gamma_{\text{old}} \vdash ce_{\text{old}}(\text{RT}) \rightsquigarrow B$  is a proof for  $\text{RT}; \Gamma_{\text{new}} \vdash ce_{\text{new}}(\text{RT}) \rightsquigarrow B$  too. On the other hand, if  $\Gamma_{\text{new}} \not\vdash \text{Reqs}(\text{RT}, \Gamma_{\text{old}})$ , then there exists at least one  $\lambda \in \text{Reqs}(\text{RT}, \Gamma_{\text{old}})$  such that  $\Gamma_{\text{new}} \not\vdash \lambda$ . Since the proof tree for a compilation judgment is unique, either there is no proof tree for compiling RT in  $\Gamma_{\text{new}}$  or the compilation produces another binary  $B'$  which differs from B.  $\square$

When the compilation for a fragment RT is undefined, or the binary produced by the compilation differs from the existing one, a sound strategy requires the compilation of RT.

To sum up, our compilation strategy is quite simple: *a fragment has to be recompiled if and only if the new environment does not entail its requirements.* This strategy is both *sound* and *minimal*.

So far so good: if we keep trace of the requirements for a fragment when we compile it we can apply a sound and minimal strategy. However, from a practical point of view there is another point to ponder: the cost of checking whether the requirements of a fragment are entailed by a type environment. If this checking costed more than compiling the source fragment, then all the reasoning so far would be useless. Luckily, this is not the case. In typechecking a source fragment a compiler must necessarily perform all the steps which are necessary to check the validity of the entailment. In addition, a compiler must, of course, parse the source and generate the code. So, checking whether the entailment holds is definitely faster than recompiling, and we expect this to be much faster. Of course, the global cost of using a smart strategy is not easy to determine because it depends on the particular compiler and the compilation context. That is, on the one hand there is the additional cost of checking the entailment, on the other hand there is the saving in not compiling the fragments whose requirements are entailed by the new environment. In a sense, the time used to find that a fragment’s requirements are *not* entailed may appear “wasted”, because that fragment has to be recompiled. However, if the compilation manager and the compiler are tightly integrated, the compiler’s typechecking step can use the results of the previous entailment checking step as a sort of cache and skip many checks when recompiling the fragment. Using this trick, the “wasted” time is extremely small: it just consists in finding that a single requirement  $\lambda$  does not hold.

We can model an *extended* compilation environment as follows:

$$ce_+ : \langle \Gamma, RT \mapsto \mathbf{S} \cup \mathbf{B} \cup (\mathbf{B} \times \Lambda) \rangle$$

That is, a pair consisting of the previous (well-formed) environment  $\Gamma$  and a function which, for each fragment name, returns either: the source only, the binary only (when the source is not available) or a pair consisting of the binary and the requirements (the idea is that they are the result of a previous compilation).

## 4 Related and Further Work

This paper, together with [3, 4], can be considered a step towards a better support for separate compilation of Java-like languages.

The solution presented here, which extends the ideas in [2], is similar to attribute recompilation, according to the classification given in [1] - here attributes correspond to assumptions.

An inspiring source of our work has certainly been Dmitriev’s paper [5], which describes a make technology, based on smart dependency checking, that aims to keep a project consistent while reducing the number of files to be recompiled. A freely downloadable tool, *Javamake*, is based on such a paper and implements the selective recompilation upon any Java compiler. Unfortunately, as pointed out by the author himself, there is no proof of the correctness of the approach, which is not based on theoretical foundations. So, it might happen that *Javamake* fails to force the recompilation of some classes which is actually needed for ensuring

the consistency of the project. Conversely, *Javamake* cannot avoid a considerable amount of unnecessary recompilations. Hence, it is neither sound nor minimal.

The final goal of the work presented in this paper is implementing a smart compilation manager for the whole Java language, à la *Javamake*, but based on a formal model on which the correctness can be actually proved. To achieve this result, there are some subtle and Java-peculiar features which must be addressed.

- *Unreachable code* is not just a bad idea: it is forbidden (see 14.20 of [8]) - this is the most challenging issue to be tackled.
- *Final methods*: an invocation of a final instance method is compiled to different bytecode w.r.t. an invocation of a non-final instance method (non-virtual invocation vs virtual one, see 15.12.3 of [8]). This issue should not pose major problems.
- *Accessing a final field initialized by a constant expression* is a constant expression (see 15.28 of [8]) and so should be compiled directly to the corresponding value. This issue should not pose problems too, as long as an assumption keeps track that such a constant is really an access to a static final field.

*Acknowledgements* We warmly thank Elena Zucca, Davide Ancona and Sophia Drossopoulou for their useful suggestions and feedback.

## References

1. Rolf Adams, Walter Tichy, and Annette Weinert. The cost of selective recompilation and environment processing. *ACM Transactions on Software Engineering and Methodology*, 3(1):3–28, January 1994.
2. D. Ancona and G. Lagorio. Stronger Typings for Separate Compilation of Java-like Languages. Technical report, DISI, March 2003.
3. D. Ancona, G. Lagorio, and E. Zucca. A formal framework for Java separate compilation. In B. Magnusson, editor, *ECOOP 2002 - Object-Oriented Programming*, number 2374 in Lecture Notes in Computer Science, pages 609–635. Springer, 2002.
4. D. Ancona, G. Lagorio, and E. Zucca. True separate compilation of Java classes. In *ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP'02)*, pages 189–200. ACM Press, 2002.
5. M. Dmitriev. Language-specific make technology for the Java programming language. *ACM SIGPLAN Notices*, 37(11):373–385, 2002.
6. S. Drossopoulou and S. Eisenbach. Describing the semantics of Java and proving type soundness. In J. Alves-Foss, editor, *Formal Syntax and Semantics of Java*, number 1523 in Lecture Notes in Computer Science, pages 41–82. Springer, 1999.
7. S. Drossopoulou, T. Valkevych, and S. Eisenbach. Java type soundness revisited. Technical report, Dept. of Computing - Imperial College of Science, Technology and Medicine, September 2000.
8. J. Gosling, B. Joy, G. Steele, and G. Bracha. *The Java™ Language Specification, Second Edition*. Addison-Wesley, 2000.
9. SUN Microsystems. The Java HotSpot Virtual Machine, 2001. Technical White Paper.