

Another Step Towards a Smart Compilation Manager for Java*

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

Categories and Subject Descriptors

D.3.4 [Language Constructs and Features]: Compilers;
D.3.4 [Language Constructs and Features]: Incremental Compilers

General Terms

Languages

Keywords

Java, separate compilation

ABSTRACT

In a recent work we have proposed a compilation strategy (that is, a way to decide which unchanged sources have to be recompiled) for a substantial subset of Java which has been shown to be *sound* and *minimal*. That is, an unchanged source is recompiled if and only if its recompilation produces a different binary or an error. However, that model does not handle two features of Java, namely, compile-time constant fields (`static final` fields initialized by a compile-time constant of a primitive type or `String`) and unreachable code, which turn out to be troublesome for having a sound and minimal compilation strategy. To our best knowledge these two features, probably because of their low-level nature, have been omitted in all models of Java separate compilation written so far. Yet, a compilation strategy for full Java has to deal with them. Thus, in this paper we analyze the implications of handling compile-time constant fields and unreachable code, and extend our previous model in order to handle these two features as well.

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

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC'04 March 14-17, 2004, Nicosia, Cyprus
Copyright 2004 ACM 1-58113-812-1/03/04 ...\$5.00.

1. INTRODUCTION

The intuitive idea behind smart (re)compilation mechanisms is to be able to recompile only a part of a program after having changed some sources, with the implicit requirement that the result of such a recompilation should be equivalent to recompiling all the sources, but cheaper (in time). Note that recompiling only the changed sources is, in most cases, not enough, since these changes could affect the compatibility with other sources. On the other hand, recompiling a source *S*, whose corresponding binary *B* is already present, to obtain the same binary *B* would be a waste of time. In a recent work [8] we have proposed a compilation strategy (that is, a way to decide which unchanged sources have to be recompiled) for a substantial subset of Java which has been shown to be *sound* and *minimal*. That is, an unchanged source is recompiled if and only if its recompilation produces a different binary (or an error). However, that model does not handle two features of Java, namely, compile-time constant fields (`static final` fields initialized by a compile-time constant of a primitive type or `String`) and unreachable code, which turn out to be troublesome for having a sound and minimal compilation strategy. To our best knowledge these two features, probably because of their low-level nature, have been omitted in all models of Java separate compilation written so far. Yet, a compilation strategy for the full Java has to deal with them. Thus, we first discuss the implications of handling compile-time constant fields and unreachable code and then we extend our previous model in order to handle these two features as well. Section 2 discusses, by means of some examples, why the way Java handles compile-time constant fields and unreachable code affects selective recompilation. Section 3 extends the model we have presented in [8]. Finally, Section 4 discusses related and further work.

2. AN INFORMAL PRESENTATION

Because of lack of space we cannot describe in details how the model we extend works, so we recall here the basic idea but refer to the original paper [8] for an extended description. The intuitive idea behind our smart recompilation mechanism is to collect, while compiling a source fragment *S* into a binary *B*, a set of type assumptions which represent the *minimal* requirements on the other fragments which have been used to compile *S* into *B*. So, the result of each successful compilation is a binary, as usual, plus a set of assumptions. After a change to some sources is made, the requirements of all *unchanged* fragments are checked against

the new type environment¹ (the changed sources have to be recompiled anyway, so their assumptions are not checked). All unchanged sources whose type assumptions still hold in the new environment need not to be recompiled, since they would be recompiled to the same binaries. Instead, all other fragments have to be recompiled because either they compile to different binaries or they do not successfully compile anymore.

As said in the introduction, two features (not considered in the original model) heavily affect the smart recompilation mechanism: compile-time constant fields and unreachable code. The former is discussed in Section 2.1 and the latter in Section 2.2.

2.1 Compile-time Constant Fields

Compile-time constant fields are `static final` fields, of a primitive type or `String`, which are initialized by a compile-time constant. For the sake of brevity we will call them *ctc-fields* from now on. Because the value of these fields is a compile-time constant (see 15.28 of [7]), the bytecode generated for an access to a ctc-field is the same that would be generated if the literal corresponding to the value of the field was used in its place.

Consider, for instance, the following declarations (assuming that each class is declared in a separate file).

```
class A {
    static final int CONST_A = 1 ;
}
class B {
    static final int CONST_B = 1 ;
}
class FirstClient {
    int ma() {
        return A.CONST_A ;
    }
    int mb() {
        return B.CONST_B ;
    }
}
class SecondClient {
    int m() {
        return A.CONST_A+B.CONST_B ;
    }
}
```

These sources are compiled to the following binaries²:

```
class A {
    static final int CONST_A = 1 ;
}
class B {
    static final int CONST_B =1 ;
}
class FirstClient {
    int ma() { return 1 ; }
    int mb() { return 1 ; }
}
```

¹A type environment is a map from class/interface names to their types.

²As in [8], we use a very abstract view of the bytecode. In a real `.class` file the initializer expressions for all ctc-fields are not present (see 13.4.8 of [7]) and their values are stored as attributes of type `ConstantValue` (see 4.7.2 of [9]).

```
class SecondClient {
    int m() { return 2 ; }
}
```

As shown in the example, in the binaries every symbolic reference to ctc-fields has disappeared. Indeed, there is no need to calculate at run-time a value which is a compile-time constant. This means that every time the value of a ctc-field is changed and the class where it has been declared is recompiled, its clients might need to be recompiled as well, in order to have a sound compilation strategy. Indeed, the bytecode corresponding to the method `ma` declared in class `FirstClient` changes each time the value of `A.CONST_A` changes. So, we need to keep track of the dependencies between the definition of ctc-fields and their uses. For this reason, we introduce a new kind of type assumption. However, the naive approach of adding assumptions of the form “field name=value” is sound but not minimal. As the example shows, client classes `FirstClient` and `SecondClient` use both constants, but the former needs both `A.CONST_A` and `B.CONST_B` to be equal to 1 to be recompiled to the same bytecode, while the latter just needs their sum to be equal to 2. Thus, if we change, say, `CONST_A` to 0 and `CONST_B` to 2, then the first client has to be recompiled, while the second has not. Indeed, after the change, the expression `A.CONST_A+B.CONST_B` has still the same value it had before, so class `SecondClient` would be recompiled to the same existing binary. Hence, to achieve minimality we use type assumptions of the form “expression=value” where the expression is built on values and references to fields (which must be constant). Moreover, we need assumptions of the form “`f = ⊥`” which express the fact that a field `f` is *not* a ctc-field. This is needed, in this case, because every access to `f` is compiled to a field-access instruction in the bytecode but, if `f` became a ctc-field, then the compile-time value of `f` would be directly used instead.

Note that compiling a single source expression may lead to several assumptions; for instance, the source expression `A.CONST_A+aMethod(A.CONST_A>=B.CONST_B)`, is compiled³ to `1+aMethod(true)` and generates two type assumptions:

1. `A.CONST_A=1` and
2. `A.CONST_A>=B.CONST_B=true`.

These assumptions model the fact that this source expression is compiled to the same bytecode whenever `A.CONST_A` is equal to 1 (because this is the value of the left operand of the sum) and `B.CONST_B` is any integer less or equal to 1 (that is, the value `A.CONST_A` is constrained to be by the first assumption). In some cases these assumptions could be simplified; for instance, the second one can be simplified, as just noted, to `1>=B.CONST_B=true`. In the general case, anyway, this would require an analysis of assumptions deduced from different expressions and we think it is not worth the effort.

2.2 Unreachable Code

The way Java handles unreachable code is peculiar: what is a warning in most languages is a compile-time error in Java, see 14.20 of [7]. The relation between this choice and selective recompilation lies in the fact that a change in some source may make unreachable a piece of code, contained

³Ignoring method call annotations, which are of no interest in the context of this example.

in an unchanged source, which was reachable before. So, a source which was perfectly legal is not correct anymore. In this case, a sound compilation strategy must trigger the compilation on that source, in order to raise the error (unreachable code) that a global recompilation would raise.

Two ways to make unreachable a piece of code, which was reachable, without changing its source are: changing the value of a `ctc`-field and changing the exception specification of a method. Examples of both cases are shown below, starting with the simpler one: changing the value of a `ctc`-field (making it a member of the infamous “inconstant constants⁴” club ☹). For instance, consider:

```
class ThirdClient {
  void m() {
    int x = 0 ;
    while (A.CONST_A==B.CONST_B)
      ++x ;
  }
}
```

This client⁵ class can be compiled only when both `ctc`-fields `A.CONST_A` and `B.CONST_B` have the same value; when they do not, the increment of `x` can never be reached.

The following example shows that changing an exception specification can do the trick too.

```
class E extends Exception { ... }
class Foo {
  Foo m() throws E { ... }
  Foo m2() { ... }
}
class Client {
  ...
  try {
    new Foo().m().m2() ;
  } catch (E e) {
    ++x ;
  }
  ...
}
```

If `E` is removed from the exception specification of method `m`, then the increment of `x` inside the `catch` becomes unreachable. As for `ctc`-fields, we introduce a new assumption to model the requirements which guarantee a piece of code to be reachable. Again, a naive assumption like “method `m` throws `E`” is not enough to obtain a minimal, although sound, compilation strategy: indeed, if we move `E` from the exception specification of `m` to the exception specification of `m2`, then the code inside the `catch` remains reachable. The nesting of `try` statements complicates the matter further; consider, for instance, the following code:

```
try {
  m1() ;
  try {
    m2() ;
  } catch (E1 e1) { }
} catch (E2 e2) { ++x ; }
```

⁴See 13.4.8 of [7].

⁵of classes `A` and `B` which have been defined in Section 2.1.

In this example the increment of `x` is reachable if and only if:

- method `m1` has an exception specification which includes a subclass of `E2` (which can, of course, be `E2` itself) *or*
- method `m2` has an exception specification which includes a subclass of `E2` which is not a subclass of `E1` (otherwise such an exception would be captured by the `catch` clause of the inner `try`).

Using $\varepsilon(m)$ to indicate the exception specification of the method `m` we can express the above assumptions as shown below:

$$(\varepsilon(m1) \cup_{\text{Exc}} (\varepsilon(m2) \setminus_{\text{Exc}} \{E1\})) \supseteq_{\text{Exc}} \{E2\}$$

The operators \cup_{Exc} , \setminus_{Exc} and \supseteq_{Exc} are discussed in the next section.

3. EXTENSIONS TO THE MODEL

Albeit the model we have presented in [8] handles a substantial subset of Java, we need to extend it in order to model `ctc`-fields and unreachable code. In this paper we only describe the extension and briefly recall the original definitions where needed to understand the new parts. Figure 1 summarizes the changes.

- The syntax of source expressions E^s and binary expressions E^b has been enriched in order to include arithmetic expressions in the language; this change is necessary to show how constant expressions are compiled.
- The keyword `break` has been added to the syntax of source statements $STMT^s$ and binary statements $STMT^b$ because, albeit this is not strictly necessary, it makes the typing rules for the statement `while` a bit more interesting when dealing with unreachable code.
- As before, type environments Γ map a class/interface name to its type. These types contain, among other information, the signatures of the declared fields, which have been enriched to describe also `ctc`-fields. A field signature FS for a `ctc`-field consists of: the access modifier AM , the kind⁶, the type T , the name f and the compile-time value ν . This change implies that discarding bodies from sources and binaries, as described in [8], is no longer enough for extracting a type environment from a compilation environment (a collection of sources and binaries); now the initialization expressions for the `ctc`-fields declared in the source fragments must be evaluated as well.
- Exception specifications ES , as in the previous model, are sets of class names (for simplicity we allow any class to be used as an exception). In this new model we introduce exception expressions EE , which describe a set of exceptions in an abstract way. The exception expressions $\varepsilon(C(\bar{T}))$ and $\varepsilon(RT.m(\bar{T}))$ represent, respectively, the exception specification of the constructor for class `C` with parameter types \bar{T} and of the method named `m`, with parameter types \bar{T} , declared in the reference type `RT`. These abstract sets of exceptions, and

⁶The kind is always `static` for `ctc`-fields, but the kind is kept in the signature for uniformity.

E^s	::=	...	$E_1^s + E_2^s$	$E_1^s - E_2^s$...
E^b	::=	...	$E_1^b + E_2^b$	$E_1^b - E_2^b$...
$STMT^s$::=	...	break ;			
$STMT^b$::=	...	break ;			
Γ	::=	$\gamma_1 \dots \gamma_n$				
γ	::=	$C \mapsto$	$[\dots, FSS = FS_1 \dots FS_n, \dots]$			
		$I \mapsto$	$[\dots, FSS = FS_1 \dots FS_n, \dots]$			
β	::=	true false				
ι	::=	0 1 -1 2 -2 ...				
ν	::=	β ι				
FS	::=	...	AM static T f = ν			
ES	::=	$\{C_1, \dots, C_n\}$				
EE	::=	$\varepsilon(C(\bar{T}))$ $\varepsilon(RT.m(\bar{T}))$ ES				
		$EE_1 \cup_{Exc} EE_2$ $EE \setminus_{Exc} ES$				
λ	::=	...	$RT \triangleleft E^s = \nu$ $RT.f = \perp$ $EE \supseteq_{Exc} ES$			

Figure 1: Extensions to the previous model

also the actual sets of exceptions ES , can be combined together using the operators “union” \cup_{Exc} and “minus” \setminus_{Exc} which, in first approximation, can be thought as the corresponding set operations. However, they have to deal with the fact that any exception C actually stands for “ C or any subclass”. Because of lack of space we do not give here the formal semantics of these operators. An in-depth discussion of these issues can be found in [4].

- Three new kinds of type assumptions λ have been added. The first two are needed to deal with constant expressions, the last one with code reachability. The assumption $RT \triangleleft E^s = \nu$ has the meaning: “if the code contained in RT evaluates E^s it gets the value ν ”; $RT.f = \perp$ has the meaning: “the field f declared in RT is not a ctc-field”; finally, $EE \supseteq_{Exc} ES$ has the meaning: “the (abstract) set of exceptions EE contains the (actual) set ES ”.

Compilation of expressions is expressed by the following judgment:

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

with the meaning “expression E^s has type T , throws exceptions EE and compiles to binary expression E^b when contained in type RT , in a local environment Π , in a context where exceptions ES can be thrown and in a type environment Γ ”. Type RT is needed to model the access control; for instance, **private** methods of RT can be invoked only by expressions inside RT . The local environment Π maps parameter names and **this** to their respective types (**this** is undefined when typing expressions contained in static contexts). The only difference with respect to the previous model is the presence of EE , which is needed for expressing the requirements on code reachability (see below). Figure 2 shows some selected rules defining this judgment. The first two rules model the compilation of a compile-time constant expression of type, respectively, **int** and **bool**. As explained before, these expressions are compiled directly to their corresponding value and, obviously, their evaluation throw no

exceptions (so EE is the empty set). The third rule models the compilation of an expression which accesses the static field f of the type RT' . In this case the field is not a ctc-field (because of the premise $\Gamma \vdash RT'.f = \perp$) so the field-access expression is compiled to a field-access binary expression (which is annotated with the type of f). In the last rule two expressions, which are not both constant, are combined using an operator Θ : the result is a non constant expression that can throw any exception its operands can throw.

Compilation of statements is expressed by the following judgment:

$$RT; \Pi; ES; \Gamma \vdash STMT^s \rightsquigarrow STMT^b \text{ throws: } EE \text{ ccn: } \beta$$

with the meaning “statement $STMT^s$ is compiled to $STMT^b$, throws exceptions EE and can complete normally β when contained in type RT , in a local environment Π , in a context where exceptions ES can be thrown and in a type environment Γ ”. Figure 3 shows some selected rules defining this judgment. The difference with respect to the previous model is the presence of EE and the flag β , both required to model unreachable code. The intuitive idea is that a reachable statement can complete normally (so the flag β is true) when, at run-time, the flow of execution may continue beyond such a statement. For instance, a sequence of statements can complete normally if and only if all statements it consists of can complete normally (see the first two rules of Figure 3), the evaluation of a statement expression SE^b always complete normally (third rule), while a statement **break** cannot ever complete normally (fourth rule). These notions are described in detail in 14.20 of [7]. Figure 3 contains an apparent asymmetry in how the statements **if** and **while** are handled. That is, **while (false) {...}** is not allowed, because its body would be unreachable, but the similar cases **if (false) {...} else {...}** and **if (true) {...} else {...}** are ok (even though in both cases one of the two branches is indeed unreachable). The typing rules reflect the language specification which allows this use of statement **if** as a way to express conditional compilation. The last two rules in Figure 3 deal with throwing and catching exceptions. A statement **throw**, whose expression E^s has type C , can throw all the exceptions the evaluation of E^s can throw, plus, of course, C itself. A statement **throw** never completes normally. The last rule is the trickiest of all: a statement **try** can throw the exceptions its body can throw EE_0 , minus all the caught ones $\{C_1, \dots, C_n\}$, plus the ones that can be thrown by statement inside its **finally** clause, EE_{n+1} . Moreover, the statement **try** can complete normally only when at least its body $STMT_0^s$ or one of the catch clauses $STMT_i^s$ can complete normally *and* the **finally** clause can complete normally.

4. RELATED AND FURTHER WORK

This paper describes an extension of the model presented in [8], which, in turn, is based on ideas recently discussed with Ancona [2, 3] and has been inspired by Dmitriev’s paper [6]. The solution presented here is similar to attribute recompilation, according to the classification given in [1] - here attributes correspond to assumptions.

The initial goal was to provide a compilation strategy able to recompile the minimum number of source fragments still guaranteeing the soundness of the approach (that is, guaranteeing that the result of a smart recompilation is equivalent to recompiling all the sources). The idea of having some

$\frac{\Gamma \vdash \text{RT} \not\sim E^s = \iota}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash E^s \sim \iota : \text{int throws: } \emptyset}$	
$\frac{\Gamma \vdash \text{RT} \not\sim E^s = \beta}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash E^s \sim \beta : \text{bool throws: } \emptyset}$	
$\frac{\Gamma \vdash \text{RT}' . f = \perp \quad \Gamma \vdash \text{RT} \not\sim \text{Fld}(\text{RT}', f) = [\text{FINAL}=_, \text{FK}=\text{static}, \text{T}=\text{T}]}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{RT}' . f \sim \ll \text{RT}' . \text{T} \gg_{\text{sf}} f : \text{T throws: } \emptyset}$	
$\frac{\text{RT}; \Pi; \text{ES}; \Gamma \vdash E_1^s \sim E_1^b : \text{int throws: } \text{EE}_1 \quad \text{RT}; \Pi; \text{ES}; \Gamma \vdash E_2^s \sim E_2^b : \text{int throws: } \text{EE}_2}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash E_1^s \Theta E_2^s \sim E_1^b \Theta E_2^b : \text{int throws: } \text{EE}_1 \cup_{\text{Exc}} \text{EE}_2} \quad \Theta \in \{+, -, \dots\}$	$E_1^b \text{ or } E_2^b \text{ is not a value}$

Figure 2: Selected expression typing rules

$\text{RT}; \Pi; \text{ES}; \Gamma \vdash \{ \} \sim \{ \} \text{ throws: } \emptyset \text{ ccn: true}$	
$\frac{i \in 1..(n-1) \text{ RT}; \Pi; \text{ES}; \Gamma \vdash \text{STMT}_i^s \sim \text{STMT}_i^b \text{ throws: } \text{EE}_i \text{ ccn: true} \quad \text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{STMT}_n^s \sim \text{STMT}_n^b \text{ throws: } \text{EE}_n \text{ ccn: } \beta}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash \{ \text{STMT}_1^s \dots \text{STMT}_n^s \} \sim \{ \text{STMT}_1^b \dots \text{STMT}_n^b \} \text{ throws: } \text{EE}_1 \cup_{\text{Exc}} \dots \cup_{\text{Exc}} \text{EE}_n \text{ ccn: } \beta}$	
$\frac{\text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{SE}^s \sim \text{SE}^b : \text{T throws: } \text{EE}}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{SE}^s ; \sim \text{SE}^b ; \text{ throws: } \text{EE ccn: true}}$	
$\text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{break} ; \sim \text{break} ; \text{ throws: } \emptyset \text{ ccn: false}$	
$\frac{\text{RT}; \Pi; \text{ES}; \Gamma \vdash E^s \sim E^b : \text{bool throws: } \text{EE}_0 \quad i \in 1..2 \text{ RT}; \Pi; \text{ES}; \Gamma \vdash \text{STMT}_i^s \sim \text{STMT}_i^b \text{ throws: } \text{EE}_i \text{ ccn: } \beta_i}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{if } (E^s) \text{ STMT}_1^s \text{ else STMT}_2^s \sim \text{if } (E^b) \text{ STMT}_1^b \text{ else STMT}_2^b \text{ throws: } \text{EE}_0 \cup_{\text{Exc}} \text{EE}_1 \cup_{\text{Exc}} \text{EE}_2 \text{ ccn: } \beta_1 \vee \beta_2}$	
$\frac{\text{RT}; \Pi; \text{ES}; \Gamma \vdash E^s \sim E^b : \text{bool throws: } \text{EE}_0 \quad \text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{STMT}^s \sim \text{STMT}^b \text{ throws: } \text{EE}_1 \text{ ccn: } -}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{while } (E^s) \text{ STMT}^s \sim \text{while } (E^b) \text{ STMT}^b \text{ throws: } \text{EE}_0 \cup_{\text{Exc}} \text{EE}_1 \text{ ccn: true}} \quad E^b \text{ is not a value}$	$E^b \text{ is not a value}$
$\frac{\Gamma \vdash \text{RT} \not\sim E^s = \text{true} \quad \text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{STMT}^s \sim \text{STMT}^b \text{ throws: } \text{EE ccn: } -}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{while } (E^s) \text{ STMT}^s \sim \text{while } (\text{true}) \text{ STMT}^b \text{ throws: } \text{EE ccn: } \beta} \quad \beta = \text{ there is a break for this while inside STMT}^s$	$\beta = \text{ there is a break for this while inside STMT}^s$
$\frac{\text{RT}; \Pi; \text{ES}; \Gamma \vdash E^s \sim E^b : \text{C throws: } \text{EE}}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{throw } E^s ; \sim \text{throw } E^b ; \text{ throws: } \text{EE} \cup_{\text{Exc}} \{ \text{C} \} \text{ ccn: false}}$	
$\frac{\text{RT}; \Pi; \text{ES} \cup_{\text{Exc}} \{ \text{C}_1, \dots, \text{C}_n \}; \Gamma \vdash \text{STMT}_0^s \sim \text{STMT}_0^b \text{ throws: } \text{EE}_0 \text{ ccn: } \beta_0 \quad i \in 1..n \text{ RT}; \Pi[x_i \mapsto \text{C}_i]; \text{ES}; \Gamma \vdash \text{STMT}_i^s \sim \text{STMT}_i^b \text{ throws: } \text{EE}_i \text{ ccn: } \beta_i \quad \text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{STMT}_{n+1}^s \sim \text{STMT}_{n+1}^b \text{ throws: } \text{EE}_{n+1} \text{ ccn: } \beta_{n+1} \quad i \in 1..n \text{ } \Gamma \vdash (\text{EE}_0 \setminus_{\text{Exc}} \{ \text{C}_1, \dots, \text{C}_{i-1} \}) \supseteq_{\text{Exc}} \{ \text{C}_i \}}{\text{RT}; \Pi; \text{ES}; \Gamma \vdash \text{try } \{ \text{STMT}_0^s \} \text{ catch } (\text{C}_1 \text{ } x_1) \{ \text{STMT}_1^s \} \dots \text{ catch } (\text{C}_n \text{ } x_n) \{ \text{STMT}_n^s \} \text{ finally } \{ \text{STMT}_{n+1}^s \} \sim \text{try } \{ \text{STMT}_0^b \} \text{ catch } (\text{C}_1 \text{ } x_1) \{ \text{STMT}_1^b \} \dots \text{ catch } (\text{C}_n \text{ } x_n) \{ \text{STMT}_n^b \} \text{ finally } \{ \text{STMT}_{n+1}^b \} \text{ throws: } (\text{EE}_0 \setminus_{\text{Exc}} \{ \text{C}_1, \dots, \text{C}_n \}) \cup_{\text{Exc}} \text{EE}_{n+1} \text{ ccn: } (\beta_0 \vee \dots \vee \beta_n) \wedge \beta_{n+1}}$	

Figure 3: Selected statement typing rules

type assumptions which describe *exactly* what is needed for a certain source fragment to be compiled to a specific binary fragment is very similar to have a type system with principal typings [10]. These ideas have been formalized for a Java-like language by Ancona and Zucca [5].

The intuition behind our work is to shorten the recompilation time by recompiling less fragments, even though, from a practical point of view, this might not be the case: the time required to decide which files have to be recompiled has to be taken in account as well. Evaluating that time is a difficult task which heavily depends on how much the compilation manager and the compiler are integrated. In a production system they have to be tightly integrated, but in a prototype implementation they would probably consist of two different programs: the compilation manager and a standard Java compiler slightly modified (in order to produce the type assumptions in addition to the standard bytecode). Unfortunately, in order to deal with ctc-fields and unreachable code we had to add some new type assumptions which are rather complex.

We think that a contribution of this paper is showing clearly how some Java features were not designed with compositional compilation in mind. While the way ctc-fields are handled could be explained for efficiency reasons, the rationale in banning some⁷ unreachable code is not clear: a warning would be enough to help the programmer in spotting a potential problem in the code, whereas an error is just annoying when the programmer knows what she/he is doing. Albeit we do not like how these two features have been designed, our goal is to analyze Java as an example of a real-world mainstream language, so we have to bear with them.

Because much of the complexity of the new type assumptions stems from avoiding unnecessary recompilations, it could make sense to trade ease of implementation with minimality if this can also help in gaining a speed-up in checking whether a fragment has to be recompiled. Of course, we do *not* mean to give up on minimality altogether, because that would mean to compile all the sources each time and this contradicts the whole idea of smart recompilation. A compromise could consist in “relaxing” some of the assumptions so that they continue to work for the most common cases and “fail” (so an unnecessary recompilation may be triggered) on some, hopefully rare, cases.

The next step is to implement the ideas presented here, experimenting with “less precise” type assumptions and comparing the various strategies in a quantitative way as it has been done with a number of different mechanisms in [1].

Acknowledgements We warmly thank Elena Zucca and Davide Ancona for their useful suggestions and feedback.

5. 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 and G. Lagorio. Stronger Typings for Separate Compilation of Java-like Languages (Extended Abstract). In *5th Intl. Workshop on Formal Techniques for Java Programs 2003*, July 2003.
- [4] D. Ancona, G. Lagorio, and E. Zucca. A core calculus for Java exceptions. In *ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA 2001)*, SIGPLAN Notices. ACM Press, October 2001.
- [5] D. Ancona and E. Zucca. Principal typings for Java-like languages. In *ACM Symp. on Principles of Programming Languages 2004*, January 2004.
- [6] M. Dmitriev. Language-specific make technology for the Java programming language. *ACM SIGPLAN Notices*, 37(11):373–385, 2002.
- [7] J. Gosling, B. Joy, G. Steele, and G. Bracha. *The Java™ Language Specification, Second Edition*. Addison-Wesley, 2000.
- [8] G. Lagorio. Towards a Smart Compilation Manager for Java. In Blundo and Laneve, editors, *Italian Conf. on Theoretical Computer Science 2003*, number 2841 in Lecture Notes in Computer Science, pages 302–315. Springer, October 2003.
- [9] T. Lindholm and F. Yellin. *The Java Virtual Machine Specification*. The Java Series. Addison-Wesley, Second edition, 1999.
- [10] J.B. Wells. The essence of principal typings. In *International Colloquium on Automata, Languages and Programming 2002*, number 2380 in Lecture Notes in Computer Science, pages 913–925. Springer, 2002.

⁷It is of course undecidable to determine whether a statement is reachable in the general case.