

A Fragment Calculus

– towards a model of Separate Compilation, Linking and Binary Compatibility

Sophia Drossopoulou, Susan Eisenbach and David Wragg

Department of Computing, Imperial College, 180 Queen's Gate, LONDON SW7 2BZ, England

email: `sd, se, dpw@doc.ic.ac.uk`

Abstract

We propose a calculus describing compilation and linking in terms of operations on fragments, i.e. compilation units, without reference to their specific contents. We believe this calculus faithfully reflects the situation within modern programming systems.

Binary compatibility in Java prescribes conditions under which modification of fragments does not necessitate re-compilation of importing fragments. We apply our calculus to formalize binary compatibility, and demonstrate that several interpretations of the language specification are possible, each with different ramifications. We choose a particular interpretation, justify our choice, formulate and prove properties important for language designers and code library developers.

1. Introduction

Separate compilation and linking, although supported by most language implementations, is under-specified in most language descriptions [3]. In the traditional arrangement in languages such as Ada [22, 4] or Modula-2 [23], the compiler checks for type consistency and the linker resolves references and checks the order of compilation. Any units importing modified units have to be re-compiled, and so separate compilation of several units corresponds to the compilation of all units together. Thus, the situation was sufficiently simple for this under-specification not to pose major problems.

However, there exist languages and systems where separate compilation and linking are complex, and justify a formal treatment. For instance in Java [12], because of its intended support for loading and executing remotely produced code whose source code is not necessarily accessible, solutions enforcing consistency through re-compilation are not suitable. Instead, the remit of the linker has been ex-

tended: not only does it have to resolve external references, it also has to ensure that binaries (the compiled units) are structurally correct and that they respect the types of entities they import from other binaries; the order of compilation need not correspond to the import relation.

Certain source code modifications, such as adding a method to a class, are *binary compatible* [8]. The Java language description does not require re-compilation of units importing units modified in binary compatible ways, and claims that successful linking and execution of the altered program is guaranteed. Not only do binary compatible changes not require re-compilation of other units, but such re-compilations *may not be possible*: a binary compatible change to the source code for one class may cause the source code of other classes no longer to be type correct. Separate compilation is *not* equivalent to compilation of all units together.

Binary compatibility has practical importance because of related security issues [5], and implications on library modification policies. It is quite complex – the language specification is sometimes inconsistent, as it considers some changes to be binary compatible, whose combination actually leads to programs which cannot link [7].

Formalizations of such issues tend to suggest calculi describing the underlying source and binary languages, and define modularization and linking in terms of these calculi. We now believe that such approaches have serious disadvantages:

- It is rather cumbersome to establish that a full-fledged calculus (*e.g.* [6, 1, 19, 17]) faithfully reflects the properties of a real language (*e.g.* Java) with respect to linking and separate compilation.
- Such calculi are at an inappropriate level of abstraction. Rather than think in terms of the particular language features, language designers think in terms of “programming in the large” and of properties satisfied by linking and compilation; library developers think in terms of linking capabilities of libraries.

1st phase <pre>class Student { int grade; } class CStudent extends Student { } class Lab { CStudent guy; void f() { guy.grade=100; } }</pre>
2nd phase <pre>class CStudent extends Student { char grade; }</pre>
3rd phase <pre>class Marker { CStudent guy; void g() { guy.grade='A'; } }</pre>

Figure 1. Students — code

In this paper we explore a different avenue: We give an axiomatic definition compilation, linking, well-formedness for source and binary languages, and require some locality properties. We believe that our model distills the essential definitions and properties and reflects the situation in most real programming languages. Also we have taken into account feedback from Sun Java developers [2].

We use this model to formalize what it means for a source code modification to be a binary compatible change. We discovered that several interpretations of the definition in [12] are possible, and discuss their ramifications. We suggest the best interpretation in our view, and we prove properties which allow binary compatible modifications to be applied to interdependent libraries and preserve their linking capabilities. Thus, we clarify the issues around binary compatibility, and we offer a simple and abstract model.

The paper is organized as follows: In section 2 we introduce Java binary compatibility. In section 3 we describe a generic model of compilation and linking, in terms of a calculus of fragments. In section 4 we extend this model to describe updating and compiling into fragments. In sections 5 and 6 we define link compatibility, formulate and prove its properties. Finally, in section 7 we draw conclusions and outline further work.

2. Binary compatibility in Java

The motivation for the concept of binary compatibility in Java is the intention to support large scale re-use of software available on the Internet [13], and in particular, to avoid the *fragile base class problem*, found in most C++ implementations, where a field (data member or instance variable) access is compiled into an offset from the beginning of the

object, fixed at compile-time. If new fields are added and the class is re-compiled, then offsets may change, and object code that previously compiled using the original definition of the class may not execute safely together with the object code of the modified class. Similar problems may arise with virtual function calls.

Development environments usually attempt to compensate by automatically re-compiling all units importing the modified units; however, this strategy would be too restrictive in some cases. For instance, if one developed a local program P, which imported a library L1, the source for L1 was not available, L1 imported library L2, and L2 was modified, then re-compilation of L1 would not be possible. Any further development of P would therefore be impossible.

In contrast, Java promises that if the modification to L2 were binary compatible, then the binaries of the modified L2, the original L1 and the current P can be linked without error. This is possible, because Java binaries carry more type information than object code usually does.

The example in figure 1 demonstrates some of the issues connected with binary compatibility. It consists of three phases. In the first phase we create the classes Student, CStudent, and Lab. The class CStudent inherits the instance variable grade of type int. In class Lab the field guy, of class CStudent, is assigned grade 100. This program is well-formed and compiles producing binary files Student.class, CStudent.class and Lab.class. In the second phase we add the field grade of type char to class CStudent, and re-compile CStudent, producing CStudent'.class. In the third phase we define a new class, Marker. In the body of its method g(), we assign the grade 'A' to guy. The class Marker is type correct, and thus it can be compiled to produce the file Marker.class.

The two changes, *i.e.* the addition of field grade in class CStudent, and the creation of class Marker, are binary compatible changes. So, the corresponding binaries, *i.e.* Student.class, CStudent'.class, Lab.class and Marker.class, can safely be linked together.

The sources are *not* type correct any more. An attempt to re-compile the class Lab would flag a type error for the assignment guy.grade=100, since the expression guy.grade now refers to the field in class CStudent which is of type char. Also, the compiled form of the expression guy.grade in the binary Lab.class refers to an integer, whereas the compiled form of the same expression in the binary Marker.class refers to a character. The two compiled forms exist at the same time, and refer to different fields of a CStudent object; *cf.* figure 3, where guy[Student].grade represents the first and guy[CStudent].grade represents the second access.

Similar situations can arise for method calls.

3. Fragments

As in [3], we consider *fragments* as the basic units participating in compilation and linking. They represent parts of programs or libraries, and they need *not* be self-contained. The exact nature of fragments is language dependent: In Java fragments would be classes and interfaces, in Ada fragments would be packages, in Modula-2 fragments would be modules, *etc.*

For the current discussion we are not interested in the contents of the fragments. However, as we are interested in compilation and linking we distinguish \mathcal{S} fragments from \mathcal{B} fragments, where:

- \mathcal{S} is the *source language*,
- \mathcal{B} is the *binary language* containing all necessary information for execution and for compilation of importing fragments.

\mathcal{S} may stand for Java, Pascal, Ada, *etc.* \mathcal{B} may stand for the Java class files, the Modula-2 `.o` and `.sym` files, *etc.* In our previous work applied to Java[7, 6], \mathcal{S} was represented by $\text{Java}_{se}\text{-program} \times \text{environment}$ pairs, \mathcal{B} was represented by $\text{Java}_{se}\text{-program} \times \text{environment}$ pairs. Possible \mathcal{S} fragments for the students example are shown in figure 2, and possible (high level) \mathcal{B} fragments are shown in figure 3. The difference between the fragments in figure 2 and 3 is, that field accesses in the latter are enriched with information necessary for execution.

A *fragment system* distills the basic concepts necessary for the description of compilation and linking:

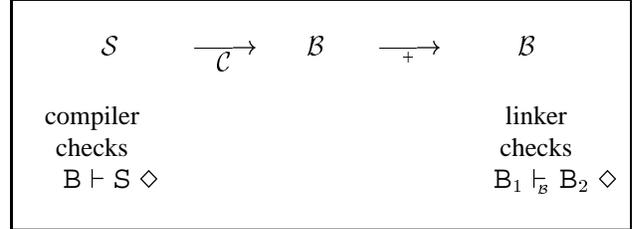
Definition 1 A tuple $(\mathcal{S}, \mathcal{B}, \epsilon, \text{Ids}, \mathcal{C}, _+ _ , \mathcal{D}, _ \vdash _ \diamond, _ \vdash_{\mathcal{B}} _ \diamond)$ is a fragment system, iff

- $\mathcal{S}, \mathcal{B}, \text{Ids}$ are sets, $\mathcal{S} \cap \mathcal{B} = \{\epsilon\}$
- \mathcal{C} is a mapping, $\mathcal{C} : \mathcal{S} \times \mathcal{B} \longrightarrow \mathcal{B}$
- $_+ _$ is a commutative, associative operator, $_+ _ : \mathcal{S} \times \mathcal{S} \longrightarrow \mathcal{S} \cup \mathcal{B} \times \mathcal{B} \longrightarrow \mathcal{B}$, with ϵ as the identity element
- \mathcal{D} a mapping, $\mathcal{D} : \mathcal{S} \cup \mathcal{B} \longrightarrow \text{Ids}$
- $_ \vdash _ \diamond$ is a relation in $\mathcal{B} \times \mathcal{S}$, and $_ \vdash_{\mathcal{B}} _ \diamond$ is a relation in $\mathcal{B} \times \mathcal{B}$

and the axioms 1, 2, 3, 4 (given later) are satisfied. The elements of $\mathcal{S} \cup \mathcal{B}$ are called fragments.

The sets \mathcal{S}, \mathcal{B} and Ids stand, respectively, for the source language, the binary language and a set of identifiers. The ϵ fragment denotes either an empty \mathcal{S} fragment or an empty

\mathcal{B} fragment. \mathcal{S} fragments are compiled into \mathcal{B} fragments by the total function \mathcal{C}^1 using environment information from imported \mathcal{B} fragments. Several fragments may be put together to form larger fragments, using the linking operator $_+ _$. The total function \mathcal{D} extracts the identifiers of all entities declared at the outer level of fragments. For \mathcal{S} fragment \mathcal{S} , \mathcal{B} fragment \mathcal{B} , the predicate $\mathcal{B} \vdash \mathcal{S} \diamond$ represents the checks performed on \mathcal{S} by the compiler in the environment of type information from \mathcal{B} , whereas the predicate $\mathcal{B}_1 \vdash_{\mathcal{B}} \mathcal{B}_2 \diamond$ represents the checks performed by the linker.



\mathcal{S} fragments will be named $\mathcal{S}, \mathcal{S}', \mathcal{S}_1, \text{etc.}$ \mathcal{B} fragments will be named $\mathcal{B}, \mathcal{B}', \mathcal{B}_1, \text{etc.}$ and fragments which may belong to either \mathcal{S} or \mathcal{B} will be named $\mathcal{F}, \mathcal{F}' \text{etc.}$ In the remainder of this section we discuss the operations and predicates $\mathcal{C}, _+ _, \mathcal{D}, _ \vdash _ \diamond, _ \vdash_{\mathcal{B}} _ \diamond$ in more detail, and formulate requirements on these in terms of axioms.

3.1. Compilation and well-formedness

Compilation (\mathcal{C}) of \mathcal{S} code produces \mathcal{B} code using environment information from \mathcal{B} code. Thus, \mathcal{B} is expected to contain two different kinds of information: the first is code for the execution of the particular fragment, the second is environment information for the compilation of importing fragments. In many language implementations this is stored in different formats and in different files – e.g. the `.o` and the `.sym` files of some Modula-2 implementations. However, since compilation produces both kinds of information, no generality is lost by not distinguishing them, and by expecting \mathcal{B} to contain execution *and* environment information. We expect $\mathcal{C}\{\mathcal{S}^{\text{st}}, \epsilon\} = \mathcal{B}^{\text{st}}$, $\mathcal{C}\{\mathcal{S}^{\text{lab}}, \mathcal{B}^{\text{st}} + \mathcal{B}^{\text{cs}}\} = \mathcal{B}^{\text{lab}}$, and $\mathcal{C}\{\mathcal{S}^{\text{lab}}, \mathcal{B}^{\text{st}}\} = \epsilon = \mathcal{C}\{\mathcal{S}^{\text{lab}}, \mathcal{B}^{\text{st}} + \mathcal{B}^{\text{cs}'}\}$.

We shall use the assertion $_ \vdash_{\mathcal{B}} _ \diamond$ as a shorthand for $\mathcal{B} \vdash_{\mathcal{B}} _ \diamond$. The first axiom expresses the requirement that empty fragments are compiled into empty fragments, and well-formedness of a non-empty \mathcal{S} fragment in the environment of a \mathcal{B} fragment is equivalent to well-formedness of the corresponding non-empty compilation.

Axiom 1 For any \mathcal{B} :

- $\mathcal{C}\{\epsilon, \mathcal{B}\} = \epsilon$

¹In previous work \mathcal{C} was a partial function; this distinction has important repercussions for the concept of binary compatibility.

1st phase $S^{st} = \text{class Student } \{ \text{int grade; } \}$ $S^{cs} = \text{class CStudent extends Student } \{ \}$ $S^{lab} = \text{class Lab } \{$ CStudent guy; void f() { guy.grade=100; } }
2nd phase $S^{st} = \text{as in 1st phase}$ $S^{cs'} = \text{class CStudent extends Student } \{$ char grade; } $S^{lab} = \text{as in 1st phase}$
3rd phase $S^{st} = \text{as in 1st and 2nd phase}$ $S^{cs'} = \text{as in 2nd phase}$ $S^{lab} = \text{as in 1st and 2nd phase}$ $S^m = \text{class Marker } \{$ CStudent guy; void g() { guy.grade='A'; } }

Figure 2. Students — source fragments

1st phase $B^{st} = \text{class Student } \{ \text{int grade; } \}$ $B^{cs} = \text{class CStudent extends Student } \{ \}$ $B^{lab} = \text{class Lab } \{$ CStudent guy; void f() { guy[Student].grade=100; } }
2nd phase $B^{st} = \text{as in 1st phase}$ $B^{cs'} = \text{class CStudent extends Student } \{$ char grade; } $B^{lab} = \text{as in 1st phase}$
3rd phase $B^{st} = \text{as in 1st and 2nd phase}$ $B^{cs'} = \text{as in 2nd phase}$ $B^{lab} = \text{as in 1st and 2nd phase}$ $B^m = \text{class Marker } \{$ CStudent guy; void g() { guy[CStudent].grade='A'; } }

Figure 3. Students — binary fragments

- $\forall S \neq \epsilon: \mathcal{C}\{\{S, B\}\} \neq \epsilon \iff B \vdash S \diamond \iff B \vdash_B \mathcal{C}\{\{S, B\}\} \diamond$

Notice that the assertion $B_1 \vdash_B B_2 \diamond$ does not imply the existence of an S fragment S such that $B_2 = \mathcal{C}\{\{S, B_1\}\}$. In Java, it is possible for definitions to be well-formed, when considered at binary level, and not to be well-formed, when considered at source level. For example, take B^{AB} :

```

 $B^{AB} = \text{class A } \{ \text{int f() } \{ \text{return 5; } \} \} +$ 
          class B extends A
          { void f() { return; } }.
  
```

A Java linker would link byte-code corresponding to B^{AB} without error. Therefore, $\vdash_B B^{AB} \diamond$ should hold. Nevertheless, the above definitions are *not* well-formed if considered at source level, because the method $f()$ in class B overrides that of A, but has a different result type [12]. Therefore, for

```

 $S^{AB} = \text{class A } \{ \text{int f() } \{ \text{return 5; } \} \} +$ 
          class B extends A
          { void f() { return; } }
  
```

for all Bs: $\mathcal{C}\{\{S^{AB}, B\}\} = \epsilon$.

3.2. Linking

The operator $+$ combines fragments, and is used both at source and at binary level. At either level, we call the $+$ operator *linking*. The source code of the first phase of

the students example consists of $S^{st} + S^{cs} + S^{lab}$; that of the second phase consists of $S^{st} + S^{cs'} + S^{lab}$. The binary code of the first phase consists of $B^{st} + B^{cs} + B^{lab}$; that of the second phase consists of $B^{st} + B^{cs'} + B^{lab}$.

The expression $\mathcal{D}(F)$ should denote the identifiers of all fragments introduced in F . So, $\mathcal{D}(S^{cs} + S^m)$ should be something like $\{ \text{CStudent, Marker} \}$.

Linking binary code in actual systems may involve several steps, *e.g.* verification of format, resolution of references, and several checks, often applied in an interleaved manner. We are not interested in these steps themselves and we consider that all checks should take place when testing well-formedness of the fragment resulting from the linking process. Thus, the case where linking fragments B_1 and B_2 should flag an error can be modeled by $\vdash_B B_1 + B_2 \diamond$ *not* holding.

Therefore, linking is based on concatenation – even though it may involve some more actions. This implies the following requirements: Linking introduces the identifiers that are separately introduced by the sub-fragments. A fragment consists of the linking of “simple” fragments each introducing one identifier ($\#(\mathcal{D}(F_i)) = 1$). Compilation introduces the same identifiers as the original.

Axiom 2 For fragments F, F', S, B :

- $\mathcal{D}(F + F') = \mathcal{D}(F) \cup \mathcal{D}(F')$
- $F = F' \implies \text{for } n = \#(\mathcal{D}(F)), \exists F_1, \dots, F_n, F'_1, \dots, F'_n:$

$F = F_1 + \dots + F_n$, $F' = F'_1 + \dots + F'_n$,
and for $i \in \{1 \dots n\}$: $F_i = F'_i$, $\#(\mathcal{D}(F_i)) = 1$

- $B \vdash S \diamond \implies \mathcal{D}(\mathcal{C}\{S, B\}) = \mathcal{D}(S)$

Note that compilation after linking, *i.e.* $\mathcal{C}\{S_1 + S_2, B\}$, need *not* be equivalent to linking after compilation, *i.e.* to $\mathcal{C}\{S_1, B\} + \mathcal{C}\{S_2, B\}$. For example, $\mathcal{C}\{S^{1ab}, B^{st}\} = \epsilon$, whereas $\mathcal{C}\{S^{1ab} + S^{cs}, B^{st}\} = B^{1ab} + B^{cs}$.

Fragments “containing” other fragments are said to subsume them:

Definition 2 F subsumes F_1 if $F = F_1 + F_2$ for some F_2 .

For example, $S^{1ab} + S^{st} + S^{cs'} + S^m$ subsumes $S^{st} + S^{cs'}$.

3.3. Disjoint fragments

In some cases we expect pairs of fragments to be *disjoint*, *i.e.* to introduce different entities:

Definition 3 For fragments F_1, F_2 :

- F_1, F_2 are disjoint iff $\mathcal{D}(F_1) \cap \mathcal{D}(F_2) = \emptyset$.

Thus, $S^{cs} + S^{st}$ and S^m are disjoint, whereas $S^{st} + S^{cs}$ and $S^m + S^{cs'}$ are not. The constituent sub-fragments of a well-formed fragment are disjoint.

Axiom 3 For fragments S_1, S_2, B, B_1, B_2 :

- $B \vdash S_1 + S_2 \diamond \implies S_1$ and S_2 disjoint
- $B \vdash_B B_1 + B_2 \diamond \implies B_1$ and B_2 disjoint

3.4. Locality

In general, one expects properties that can be established in a certain environment, to hold for larger environments as well. For instance, one expects an expression which has a type in a certain environment to have the same type in any larger environment. Such properties were proven, in [6], and also used in [5].

In particular, for fragments we expect the following locality properties: Linking disjoint, well-formed S or B fragments produces well-formed fragments. If B is well-formed in environment B_1 then it remains so in any well-formed larger environment $B_1 + B_2$. Checking a binary in the empty environment is equivalent to checking it with itself as the environment. Finally, if S is well-formed in environment B_1 and in the larger environment $B_1 + B_2$, then compilation in the two environments produces identical results.

Axiom 4 For fragments S, S_1, S_2, B, B_1, B_2 :

- $B \vdash S_1 \diamond, B \vdash S_2 \diamond, S_1, S_2$ disjoint $\implies B \vdash S_1 + S_2 \diamond$

- $B \vdash_B B_1 \diamond, B \vdash_B B_2 \diamond, B_1, B_2$ disjoint $\implies B \vdash_B B_1 + B_2 \diamond$

- $B_1 \vdash_B B \diamond, \vdash_B B_1 + B_2 \diamond \implies B_1 + B_2 \vdash_B B \diamond$

- $\epsilon \vdash_B B \diamond \iff B \vdash_B B \diamond$

- $B_1 \vdash S \diamond$ and $B_1 + B_2 \vdash S \diamond \implies \mathcal{C}\{S, B_1\} = \mathcal{C}\{S, B_1 + B_2\}$

3.5 Strong locality

Strong locality requires the compilation of a well-formed S fragment to be identical to its compilation in a larger well-formed environment, *ie*:

$$B_1 \vdash S \diamond \text{ and } B_1 + B_2 \vdash_B B \diamond \implies \mathcal{C}\{S, B_1\} = \mathcal{C}\{S, B_1 + B_2\}.$$

This property actually corresponds to recasting the third point from axiom 4 for S fragments, and it is weaker than the fifth point from axiom 4.

Strong locality is satisfied by the Java subset we have formalized [6]. In the original version of this paper, we required this property as an axiom, and this axiom was central to the argumentation of that version.

However, because of its treatment of packages, strong locality does *not* hold in full Java.² This realization led to the adoption of a slightly different formalization of binary compatibility, *cf.* sections 5 and 5.1.

3.6 Faithfulness of the model

The above concludes the axiomatic description of the basic model for compilation and linking. We believe that it describes concisely most people’s expectations of compilation and linking. The axioms are satisfied by the Java subset we have studied.

However, the question as to the faithfulness of the model to Java is open, as there does not exist a full formal specification of Java. Even if such a formal description existed, it would still be debatable how far this description corresponded to the developers’ intentions, the language definition [12] and the language implementations.

On the other hand, Sun Java developers [2] have studied the previous version of this paper, have given us feedback, and pointed out a discrepancy with Java, which we have taken into account. Thus our confidence that this model is appropriate for Java has grown.

The fact that the Java developers responded to this calculus, and not to full blown formalizations of parts of the language, is, we believe, a strong indication that the fragment calculus represents the appropriate level of abstraction for the description of issues of separate compilation and linking.

²This was reported to us by Gilad Bracha [10].

4. Updating

Java binary compatibility is concerned with the effects of modifications to source and binary code. Therefore we need operators to describe such modifications: $_ \oplus _$ describes the effect of updating some code, whereas $_ \oplus_c _$ describes the effect of updating B code through compilation of S code.

Definition 4 For fragments F_1, F_2 :

- $F_1 \oplus F_2 = F_0 + F_2$,

where F_0 such that $\exists F_3$ with

$$F_1 = F_0 + F_3, \mathcal{D}(F_3) \subseteq \mathcal{D}(F_2), F_0, F_2 \text{ disjoint.}$$

For example $(S^{\text{st}} + S^{\text{cs}} + S^{\text{lab}}) \oplus S^{\text{cs}'} = S^{\text{st}} + S^{\text{cs}'} + S^{\text{lab}}$. Also, $(B^{\text{st}} + B^{\text{lab}}) \oplus B^{\text{m}} = B^{\text{st}} + B^{\text{lab}} + B^{\text{m}}$.

We can show that \oplus is well-defined using the first two points from axiom 2. Updating is associative but not commutative. For disjoint fragments updating is equivalent to linking. Furthermore,

Lemma 1 For fragments F, F_1, F_2, F_3, F_4 :

- F disjoint from $F_2 \implies (F_1 \oplus F) + F_2 = (F_1 + F_2) \oplus F$
- $F_1 = F_3 + F, F_2 = F_4 + F$, and F_3, F_4 disjoint $\implies F_1 \oplus F_2 = F_2 \oplus F_1$
- F_1, F_2 disjoint, and F_1, F_3 disjoint $\implies F_1 + (F_2 \oplus F_3) = (F_1 + F_2) \oplus F_3$
- F_1, F_2 disjoint $\implies F \oplus (F_1 + F_2) = F \oplus F_1 \oplus F_2$.

The expression $B \oplus_c S$ denotes updating B by the compilation of S in environment B .

Definition 5 For fragments B and S , we define:

- $B \oplus_c S = B \oplus \mathcal{C}\{S, B\}$

The second phase of our example updates $B^{\text{st}} + B^{\text{cs}} + B^{\text{lab}}$ through compilation of $S^{\text{cs}'}$, giving:

$$(B^{\text{st}} + B^{\text{cs}} + B^{\text{lab}}) \oplus_c S^{\text{cs}'} = B^{\text{st}} + B^{\text{cs}'} + B^{\text{lab}}$$

The third phase compiles the new fragment S^{m} updating the result of the previous phase, giving:

$$(B^{\text{st}} + B^{\text{cs}'} + B^{\text{lab}}) \oplus_c S^{\text{m}} = B^{\text{st}} + B^{\text{cs}'} + B^{\text{lab}} + B^{\text{m}}$$

Because the arguments of $_ \oplus_c _$ come from different domains, the concepts of commutativity and associativity do not apply. We use $_ \oplus_c _$ in a left-associative manner. In general, $\mathcal{C}\{S, B\} \oplus_c S' \neq \mathcal{C}\{S \oplus S', B\}$. The left hand side represents separate compilation of fragments whereas the right hand side represents compilation of the fragments

together. As we mentioned earlier, in Java these can be different, and it is possible for one compilation to succeed, and the other not to. This happens in the second phase of the students example, where $\mathcal{C}\{S^{\text{lab}}, B^{\text{st}} + B^{\text{cs}}\} \oplus_c S^{\text{cs}'} = B^{\text{lab}} + B^{\text{cs}'}$, whereas $\mathcal{C}\{S^{\text{lab}} \oplus S^{\text{cs}'}, B^{\text{st}} + B^{\text{cs}}\} = \epsilon$.

The following lemma, used later to prove lemma 5, describes the result of interleaving compilation and linking. If B_1 and S are disjoint, then B_1 remains unaffected when updating $B_1 + B_2$ through compilation of S , but may be taken into account when compiling S .

Lemma 2 For fragments S, B_1, B_2 , with B_1 disjoint B_2 , and B_1 disjoint S :

- $(B_1 + B_2) \oplus_c S = B_1 + (B_2 \oplus \mathcal{C}\{S, B_1 + B_2\})$

Proof by lemma 1, first and third point of axiom 2. \square

5. Binary compatibility

The Java language specification [12] describes binary compatible changes as follows:

“A change to a type is binary compatible with (equivalently, does not break compatibility with) pre-existing binaries if pre-existing binaries that previously linked without error will continue to link without error.”

Our notion of *binary compatible change* aims to capture the above. It restricts source code modifications in terms of properties of the resulting compilation. Therefore, its formalization will have the general form:

S is a binary compatible change of B iff:

$$\vdash_B \dots B \dots \diamond \implies \vdash_B \dots B \oplus_c S \dots \diamond$$

During the process of formalization we realized that the definition from [12] is *not* unambiguous, because it does not answer the following questions.

Q1 Should source code which cannot be compiled be considered a binary compatible change?

Q2 How many binaries are meant?

Q3 What should be the environment for the compilation of the modified code?

Regarding question Q1, in the previous version of this paper we considered changes which did not compile to be binary *incompatible*. This tied in with the fact that we described compilation through a *partial* function. Consideration of questions Q2 and Q3 led to four alternative interpretations, which are described in section 5.1. We had postulated strong locality, and thus could prove composition properties such as those in section 6.

Without strong locality, the composition properties do not hold for link compatibility as defined in section 5.1. However, such properties are part of the rationale for the concept of binary compatibility [2, 20]. Therefore, when the lack of strong locality was reported to us, we felt that our approach required revision.

Thus, we explore the approach whereby a change which does not compile *is* binary compatible. This allows us to re-establish the composition properties and prove them more simply.

This approach motivates our current description of compilation through a *total* function \mathcal{C} , where $\mathcal{C}\{\mathcal{S}, \mathcal{B}\} = \epsilon$ if $\mathcal{B} \vdash \mathcal{S} \diamond$ does not hold. Moreover, source code which cannot be compiled does not produce binaries, and therefore does not modify the original binaries; so it can be understood as the empty change, and hence binary compatible.

We now consider the remaining questions Q2 and Q3 and discuss the three alternative interpretations I1, I2 and I3.

Definition 6 (I1-I3) For \mathcal{S} fragment \mathcal{S} and \mathcal{B} fragment \mathcal{B} :

I1 \mathcal{S} is a weak binary compatible change of \mathcal{B} if:

$$\vdash_{\mathcal{B}} \mathcal{B} \diamond \implies \vdash_{\mathcal{B}} \mathcal{B} \oplus_c \mathcal{S} \diamond$$

I2 \mathcal{S} is a local binary compatible change of \mathcal{B} , iff for all \mathcal{B}_0 disjoint from \mathcal{S} :

$$\vdash_{\mathcal{B}} \mathcal{B}_0 + \mathcal{B} \diamond \implies \vdash_{\mathcal{B}} \mathcal{B}_0 + (\mathcal{B} \oplus_c \mathcal{S}) \diamond$$

I3 \mathcal{S} is a binary compatible change of \mathcal{B} , iff for all \mathcal{B}_0 disjoint from \mathcal{S} :

$$\vdash_{\mathcal{B}} \mathcal{B}_0 + \mathcal{B} \diamond \implies \vdash_{\mathcal{B}} (\mathcal{B}_0 + \mathcal{B}) \oplus_c \mathcal{S} \diamond$$

Thus, binary compatibility requires compilation to preserve linking capabilities, but only if it was successful. We now discuss the three interpretations:

Interpretation I1 is too weak. For instance, it allows the removal of a method definition $\mathfrak{f}()$ from a class in \mathcal{B} , provided that $\mathfrak{f}()$ were not called inside \mathcal{B} . However, further libraries which linked with \mathcal{B} might rely on $\mathfrak{f}()$.

Interpretation I2 is too localized. Namely, $\mathcal{S}^{cs'}$ is a local binary compatible change of $\mathcal{B}^{st} + \mathcal{B}^{cs} + \mathcal{B}^{lab}$, and \mathcal{S}^{cs} is not a local binary compatible change of $\mathcal{B}^{st} + \mathcal{B}^m$; these facts are as expected. However, \mathcal{S}^{cs} is trivially a local binary compatible change of \mathcal{B}^m (because $\mathcal{C}\{\mathcal{S}^{cs}, \mathcal{B}^m\} = \epsilon$), even though compilation of \mathcal{S}^{cs} destroys the well-formedness of the environment $\mathcal{B}^{st} + \mathcal{B}^{cs'} + \mathcal{B}^m$.

Therefore, we adopt interpretation I3. With this, $\mathcal{S}^{cs'}$ is a binary compatible change of $\mathcal{B}^{st} + \mathcal{B}^{cs} + \mathcal{B}^{lab}$ and of $\mathcal{B}^{cs} + \mathcal{B}^{lab}$; \mathcal{S}^{cs} is not a local binary compatible change of $\mathcal{B}^{st} + \mathcal{B}^m$, and not a binary compatible change of \mathcal{B}^m .

5.1 Further interpretations

The interpretations I1, I2, and I3 were motivated by the absence of the strong locality property[10]. However, the

absence of the strong locality property is due to an idiosyncratic treatment of packages, and may not be a desirable feature for Java anyway. It may not even be a feature of future “web-centered” languages. Furthermore, a treatment which considers an \mathcal{S} fragment which does not compile, to be binary compatible, may be counter-intuitive.

Thus, it is worthwhile exploring a different approach to question Q1, whereby an \mathcal{S} fragment which does not compile is considered binary incompatible. This approach can be based on a treatment of compilation as a *partial* function \mathcal{C}_p , where $\mathcal{C}_p\{\mathcal{S}, \mathcal{B}\} = \mathcal{C}\{\mathcal{S}, \mathcal{B}\}$ iff $\mathcal{B} \vdash \mathcal{S} \diamond$, and undefined otherwise. We also define \oplus_{c_p} as $\mathcal{B} \oplus_{c_p} \mathcal{S} = \mathcal{B} \oplus \mathcal{C}_p\{\mathcal{S}, \mathcal{B}\}$.

If we consider the remaining questions Q2 and Q3 we obtain four further interpretations:

Definition 7 (I4-I7) For \mathcal{S} fragment \mathcal{S} and \mathcal{B} fragment \mathcal{B}

I4 \mathcal{S} is a weak link compatible change of \mathcal{B} if:

$$\vdash_{\mathcal{B}} \mathcal{B} \diamond \implies \vdash_{\mathcal{B}} \mathcal{B} \oplus_{c_p} \mathcal{S} \diamond$$

I5 \mathcal{S} is a strong link compatible change of \mathcal{B} , iff for all \mathcal{B}_0 disjoint from \mathcal{S} :

$$\vdash_{\mathcal{B}} \mathcal{B}_0 + \mathcal{B} \diamond \implies \vdash_{\mathcal{B}} \mathcal{B}_0 + (\mathcal{B} \oplus_{c_p} \mathcal{S}) \diamond$$

I6 \mathcal{S} is a link compatible change of \mathcal{B} , iff for all \mathcal{B}_0 disjoint from \mathcal{S} :

$$\vdash_{\mathcal{B}} \mathcal{B}_0 + \mathcal{B} \diamond \implies \vdash_{\mathcal{B}} (\mathcal{B}_0 + \mathcal{B}) \oplus_{c_p} \mathcal{S} \diamond$$

I7 \mathcal{S} is a link compatible change of \mathcal{B}_1 in the context of \mathcal{B}_2 iff \mathcal{B}_2 is disjoint from \mathcal{S} and \mathcal{B}_1 , and for all \mathcal{B}_0 disjoint from \mathcal{S} :

$$\vdash_{\mathcal{B}} \mathcal{B}_0 + \mathcal{B}_1 + \mathcal{B}_2 \diamond \implies \vdash_{\mathcal{B}} (\mathcal{B}_0 + \mathcal{B}_1 + \mathcal{B}_2) \oplus_{c_p} \mathcal{S} \diamond.$$

The difference between binary compatibility (I1-I3) and link compatibility(I4-I7) is the following: Link compatibility requires preservation of linking capabilities *and* successful compilation, whereas binary compatibility requires preservation of linking capabilities *only if* compilation is successful.

Interpretation I4 is too weak, for the same reasons which make interpretation I1 too weak.

Interpretation I5 is too strong, because it expects \mathcal{B} to contain *all* information necessary for the compilation of \mathcal{S} . Thus, if \mathcal{S} contained code using properties of the predefined class `String`, it would be a strongly link compatible change of `String` – even if \mathcal{S} did not modify `String`, but only used it.

According to interpretation I6, \mathcal{S}^{lab} is a link compatible change of $\mathcal{B}^{st} + \mathcal{B}^{cs}$, and $\mathcal{S}^{cs'}$ is a link compatible change of $\mathcal{B}^{st} + \mathcal{B}^{cs} + \mathcal{B}^{lab}$. Interpretation I6 is weaker than interpretation I5, because it is possible for $(\mathcal{B}_0 + \mathcal{B}) \oplus_{c_p} \mathcal{S}$ to be defined and for $\mathcal{B} \oplus_{c_p} \mathcal{S}$ not to be. This subtlety allows \mathcal{S} to be a link compatible change of a library \mathcal{B} , which imports other libraries, and which cannot be compiled in isolation,

i.e. $\vdash_B B \diamond$ does not hold. Such a library can only be compiled in the presence of further libraries, represented by the fragment B_0 , with which $\vdash_B B_0 + B \diamond$.

Also, B does not need to contain *all* the type information necessary to type check S ; it only needs to contain *enough* information to ensure type correct compilation of S in the environment of all appropriate fragments B_0 .

Interpretation I6 is the one we had adopted in [7]. However, the reference binaries are still too extensive.³ If for instance, S used features of the predefined class `String`, then in order for S to be a link compatible change of B , B would have *either* to use the same features of `String`, *or* to contain the class declaration of `String`. Since however, S only *uses* the class `String` and does not *modify* it, the distinction of the role of `String` should be reflected in the definition.

Thus, in interpretation I7 we distinguish B_2 , the *context* which may *not* be modified by the compilation of S , from B_1 , which *may*. Therefore, S^m is a link compatible change of ϵ in the context of B^{cs} . Also, an S which uses `String` may be a link compatible change of B , in the context of class `String`.

In fragment systems which satisfy the strong locality property, link compatible changes, as in interpretations I5-I7, enjoy composition properties corresponding to these from the next section.

Link compatibility implies binary compatibility in the sense that any link compatible change of a B is also a binary compatible change of B . Also, strong link compatibility implies link compatibility. We expect that further entailment relationships between the interpretations hold; their proof may require a refinement of the fragment system definition.

6. Composition Properties

We now demonstrate the following five properties:

- **Preservation over larger fragments:** establishes binary compatibility for all fragments containing a fragment for which this property has already been established.
- **Preservation over sequences:** guarantees that combined binary compatible steps preserve their linking capabilities – provided that each step is a binary compatible change of the result of all previous modifications – cf. figure 4.
- **Preservation over libraries:** application of binary compatible changes to different fragments preserves well-formedness – cf. figure 5.

³Having small reference binaries is important, because this allows modifications to be applied to more fragments comprising a large program, as in lemma 5.

- **Lack of folding property:** in general, two binary compatible changes cannot be folded into one.
- **Lack of diamond property:** two binary compatible changes applied to the same fragment, cannot always be reconciled.

These properties are, we believe, crucial in delineating the exact nature of binary compatibility, and are central issues in the design of that feature [20].

Also, these properties affect the way library designers can evolve their libraries: The lack of a folding or a diamond property restricts the ways in which binary compatible changes may be combined. The lack of diamond property means that programmers may not apply *independent* binary compatible changes to the *same* fragment and expect the linking capabilities to be preserved. However, the preservation over libraries allows programmers to apply *independent* binary compatible changes and expect the linking capabilities to be preserved, as long as they were working on *different* fragments. In particular, it means that various libraries may be modified separately, each in binary compatible ways, and still preserve their linking capabilities. This holds, even if these libraries should import each other.

Next we formulate and prove these properties.

Preservation over larger fragments Any binary compatible change is also a binary compatible change of a larger fragment:

Lemma 3 For S , B_1 , B_2 , where S and B_2 are disjoint:

$$\begin{aligned} S \text{ a binary compatible change of } B_1 &\implies \\ S \text{ a binary compatible change of } B_1 + B_2 & \end{aligned}$$

Proof by commutativity and associativity of $+$. □

Preservation over sequences As outlined in figure 4, a sequence of binary compatible steps, S_1, \dots, S_n , applied to fragment B preserves its linking capabilities. In order to establish that a step is binary compatible, we need to know the effect of all prior steps, thus we require that S_{i+1} is binary compatible for $(B_0 + B) \oplus_c S_1 \dots \oplus_c S_i$.

Lemma 4 For B fragments B , B_0 , a sequence of S fragments S_1, \dots, S_n , B_0 disjoint S_i , if

- for $1 \leq i \leq n$: S_{i+1} binary compatible change of B^i where $B^i = (B_0 + B) \oplus_c S_1 \dots \oplus_c S_i$
- $\vdash_B B_0 + B \diamond$

then

- $\vdash_B (B_0 + B) \oplus_c S_1 \dots \oplus_c S_n \diamond$

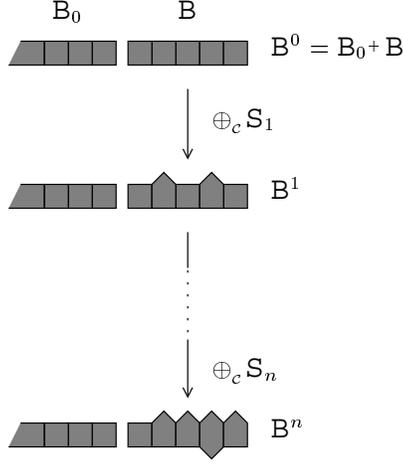


Figure 4. Preservation over sequences

Proof by induction on k ; using that $B^0 = B_0 + B$ and $B^{k+1} = B^k \oplus_c S_{k+1}$, prove that $\vdash_B B^k \diamond$ for all k . Also, $B^n = (B_0 + B) \oplus_c S_1 \dots \oplus_c S_n$. \square

Preservation over libraries Binary compatible modifications S_i applied to fragments B_i which are parts of a program $B_0 + B_1 + \dots + B_n$, preserve the linking capabilities of that program, provided that the modifications are binary compatible for the particular fragments only – *i.e.* require S_i to be a binary compatible change of B_i , which is stronger than requiring S_i to be a binary compatible change of $B_0 + B_1 + \dots + B_n$.

In contrast to preservation over sequences, we do not need to know the effect of another modification in order to establish that S_i is a binary compatible change of B_i , but we take into account the effect of previous modifications. Thus, B_k is transformed to B'_k , where $B'_k = B_k \oplus \mathcal{C}\{S_k, B_0 + B'_1 + \dots + B'_{k-1} + B_{k+1} + \dots + B_n\}$; as in figure 5.

This models the situation where programmers make changes to the particular fragments that belong to them, but *are aware* of each other's actions. When all modified fragments are put together, the resulting program $B_0 + B'_1 + \dots + B'_n$ preserves the linking capabilities of the original program. The order of the fragments is immaterial.

Lemma 5 For $B_0, B_1, \dots, B_n, S_1, \dots, S_n$, where S_i disjoint from B_k , from S_k and from B_0 for all $i \neq k, i, k \in \{1 \dots n\}$, if

- S_i is a binary compatible change of B_i for $1 \leq i \leq n$
- $\vdash_B B_0 + B_1 + \dots + B_n \diamond$

then

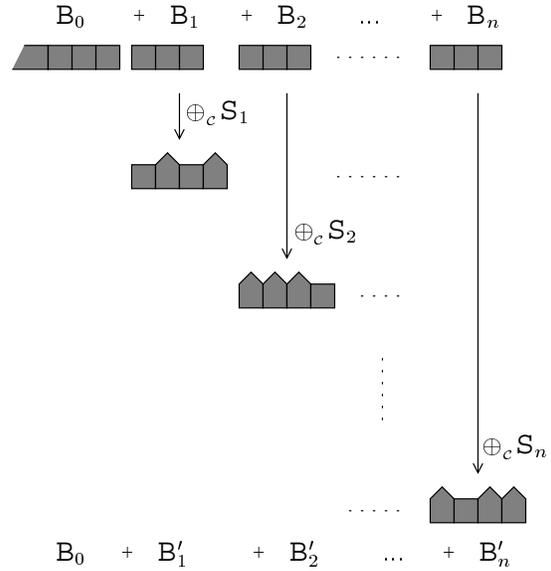


Figure 5. Preservation over libraries

- $\vdash_B B_0 + B'_1 + \dots + B'_n \diamond$
where $B'_k = B_k \oplus \mathcal{C}\{S_k, B_0 + B'_1 + \dots + B'_{k-1} + B_{k+1} + \dots + B_n\}$

Proof Define $B^k = B_0 + B'_1 + \dots + B'_k + B_{k+1} + \dots + B_n$. By induction on k , and using lemma 2, we show that B'_k and B^k are defined, and that $\vdash_B B^k \diamond$ for all k . \square

Interestingly, the other case of the lemma, whereby programmers apply binary compatible changes *unaware* of each other's actions, *cannot* be proven. That is, for $B'_k = B_k \oplus \mathcal{C}\{S_k, B_0 + B_1 + \dots + B_n\}$ we cannot prove that $\vdash_B B_0 + B'_1 + \dots + B'_n \diamond$. This is so, because binary compatibility guarantees preservation of well-formedness if the compilation took place in the *same* environment. In other words, from $\vdash_B B_0 + \dots + B'_k + B_{k+1} + \dots + B_n \diamond$ it is impossible to infer that $\vdash_B B_0 + \dots + B'_k + (B_{k+1} \oplus \mathcal{C}\{S_{k+1}, B_0 + B_1 + \dots + B_n\}) + \dots + B_n \diamond$.

The concepts of transitivity and reflexivity are not applicable to the binary compatibility relationship, because its domain and range do not match. Neither is there a folding or a diamond property:

Lack of folding property For S_1 a binary compatible change of B , and S_2 a binary compatible change of $(B_0 + B) \oplus_c S_1$, it does not necessarily hold that $(B_0 + B) \oplus_c S_1 \oplus_c S_2 = (B_0 + B) \oplus_c (S_1 \oplus S_2)$. As a counter-example, consider $B = B^{st} + B^{cs}$, $S_1 = S^{lab}$ and $S_2 = S^{cs'}$.

Lack of diamond property For S_1 and S_2 binary compatible changes of B , there do not always exist fragments S_3 and S_4 , such that S_3, S_4 disjoint with S_1, S_2 , and S_3 is

a binary compatible change of $B_0 \oplus_c S_1$, and S_4 is a binary compatible change of $B \oplus_c S_2$, and $B \oplus_c S_1 \oplus_c S_3 = B \oplus_c S_2 \oplus_c S_4$. For example, S_1 might be introducing a method f with signature $\text{int} \rightarrow \text{int}$ into a class C , and S_2 introducing another method f with signature $\text{int} \rightarrow \text{char}$ into the same class C .

7. Conclusions and further work

We gave an axiomatic definition of compilation and linking, and extended it to model binary compatibility. In our view, the contributions of this paper are:

- identification of the appropriate level of abstraction with respect to the description of separate compilation and linking,
- a distillation of the essential features with respect to separate compilation and linking,
- a clarification of the design space for binary compatibility,
- a formal framework for the description of binary compatibility and proof of its properties,
- a strengthening of the properties of binary compatibility proven earlier [7],
- demonstration that the properties of binary compatibility in Java stem from the few features described in the fragment calculus rather than from the rather large set of features of Java [16, 21] and its byte-code [18, 19, 9].

We believe that such a fragment calculus can serve as a basis for the description of the approaches to separate compilation and linking taken by other languages, and as a starting point for an abstract description of dynamic linking and loading in Java [15, 14, 11].

Acknowledgements

We acknowledge the financial support from the EPSRC Grant Ref:GR/L 76709.

We thank David Clarke for feedback, Phil Wadler for heated discussions and many useful suggestions and the anonymous LICS referees for insightful suggestions.

We are very grateful to Gilad Bracha, for detailed comments on the previous version of the paper, and for helping us keep on the right track.

References

[1] E. Boerger and W. Schulte. A Programmer Friendly Modular Definition of the Semantics of Java. In Jim Alves-Foss,

editor, *Formal Syntax and Semantics of Java*, volume 1523. Springer Verlag LNCS, 1999.

[2] G. Bracha. Private Communications, February 1999.

[3] L. Cardelli. Program Fragments, Linking, and Modularization. In *POPL'97 Proceedings*, January 1997.

[4] M. Dausmann, S. Drossopoulou, G. Persch, and G. Winterstein. A Separate Compilation System for Ada. In *Proc. GI Tagung: Werkzeuge der Programmieretechnik*. Springer Verlag Lecture Notes in Computer Science, 1981.

[5] D. Dean. The Security of Static Typing with Dynamic Linking. In *Fourth ACM Conference on Computer and Communication Security*, 1997.

[6] S. Drossopoulou, S. Eisenbach, and S. Khurshid. Is Java Sound? *Theory and Practice of Object Systems*, 5(1), Jan. 1999. available at <http://www-dse.doc.ic.ac.uk/projects/slurp/>.

[7] S. Drossopoulou, D. Wragg, and S. Eisenbach. What is Java Binary Compatibility? In *OOPSLA Proceedings*, October 1998.

[8] I. Forman, M. Conner, S. Danforth, and L. Raper. Release-to-Release Binary Compatibility in SOM. In *OOPSLA'95 Proceedings*, 1995.

[9] S. N. Freund and J. C. Mitchell. A Type System for Object Initialization in the Java Bytecode Language. In *OOPSLA Proceedings*, October 1998.

[10] Gilad Bracha. Packages break Strong Locality. In *more at http://www-dse.doc.ic.ac.uk/sue/packages.html*, February 1999.

[11] A. Goldberg. A Specification of Java Loading and Bytecode Verification. In *5th ACM Conference on Computer and Communications Security*, 1998.

[12] J. Gosling, B. Joy, and G. Steele. *The Java Language Specification*. Addison-Wesley, August 1996.

[13] J. Gosling and H. McGilton. The Java white paper, <http://java.sun.com/java.sun.com/whitePaper/java-whitepaper1.html> 1995.

[14] T. Jensen, D. L. Metyayer, and T. Thorn. A Formalization of Visibility and Dynamic Loading in Java. In *IEEE ICCL*, 1998.

[15] S. Liang and G. Bracha. Dynamic Class Loading in the JavaTM Virtual Machine. In *OOPSLA Proceedings*, October 1998.

[16] T. Nipkow and D. von Oheimb. Java_{right} is type-safe — definitely. In *25th POPL Proceedings*, 1998.

[17] Z. Qian. A Formal Specification of the Java Virtual Machine Instructions. <http://www.informatik.uni-bremen.de/qian/abs-fsjvm.html>, 1997.

[18] Z. Qian. Least Types for Memory Locations in Java Bytecode. In *FOOL 6*. <http://www.cs.williams.edu/kim/FOOL/sched6.html>, 1999.

[19] R. Stata and M. Abadi. A Type System For Java Bytecode Subroutines. In *POPL'98 Proceedings*, January 1998.

[20] G. Steele. Private Communication, January 1998.

[21] D. Syme. Proving Java Type Sound. In Jim Alves-Foss, editor, *Formal Syntax and Semantics of Java*, volume 1523 of LNCS. Springer, 1999.

[22] US Department of Defense. *Reference Manual for the Ada Programming Language*, 1983. ANSI/MIL-STD-1815 A.

[23] N. Wirth. *Programming in Modula-2*. Springer-Verlag, 1982.