Flexible Models for Dynamic Linking

Sophia Drossopoulou[1], Giovanni Lagorio[2], Susan Eisenbach[1] ⋆

[1] Department of Computing at Imperial College, [2] DISI at the University of Genova

Abstract Dynamic linking supports flexible code deployment: partially linked code links further code on the fly, as needed; and thus, end-users receive updates automatically. On the down side, each program run may link different versions of the same code, possibly causing subtle errors which mystify end-users.
Dynamic linking in Java and C# are similar: The same linking phases are involved, soundness is based on similar ideas, and executions which do not throw linking errors give the same result. They are, however, not identical: the linking phases are combined differently, and take place in different order. Thus, linking errors may be detected at different times in Java and C#.
We develop a non-deterministic model, which includes the behaviour of Java and C#. The non-determinism allows us to describe the design space, to distill the similarities between the two languages, and to use one proof of soundness for both. We also prove that all execution strategies are equivalent in the sense that all terminating executions which do not involve a link error, give the same result.

1 Introduction

Dynamic linking supports flexible code deployment and update: instead of fully linking code before execution, further code is linked on the fly, as needed. Thus, the *newest* version of any imported code is always linked, and the most recent updates are automatically available to users. Dynamic linking was incorporated into operating systems, *e.g.*, Multics[32], Unix, and Windows dynamic link libraries (DLLs), which enable applications to share code, thus saving disk and memory usage. Recently, Java and C# [1] incorporated dynamic linking into the language.

One question connected to dynamic linking is the choice of components to be linked, esp. if there are several with the same name. DLLs and .NET offer sophisticated systems of versioning, side-by-side components, registries *etc*. Difficulties in managing DLLs lead to the term "DLL Hell" [24]. The .NET architecture, with assemblies carrying versioning information claims to have solved this problem [25]. Java, on the other hand, links the first class with given name found in the classpath, and any more sophisticated scheme can be implemented through custom class loaders [22].

Another question connected to dynamic linking is the type safety guarantees given *after* choosing components. Breaking type safety jeopardizes the integrity of memory, and ultimately security [11,23]. DLLs do not attempt to guarantee type safety: type errors may occur and go undetected, or throw exceptions of unrelated nature in unrelated

⋆ work partly supported by DART, European Commission research Directorate IST-01-6-1A

[1] or rather, the common runtime system of .NET which deals with many languages rather than just C#; but we focus on C# in this paper, as it is most easily compared with Java

parts of the code. Conversely, in Java and C# if the components linked turn out to be "incompatible", link related exceptions are thrown, describing the nature of the problem. Thus, although Java and C# do not guarantee choice of compatible components, they guarantee type safety and give error messages that signal the source of the problem.

Our study is concerned with how Java and C# tackle the second question, *i.e.,* how they guarantee type safety. Dynamic linking in Java and C# are similar: The same linking phases are involved, *i.e.*, loading, verification, offset calculation, and layout determination. Soundness is based on similar ideas: *i.e.,* consistency of the layout and virtual tables, verifying intermediate code, and checking before calculating offsets. Executions which do not throw linking errors give the same result.

However, dynamic linking in Java and C# have some differences: The linking phases have different granularity, are combined differently and take place in different order. Linking errors may be detected at different times in Java and C# executions.

We develop a non-deterministic model, which includes the behaviour of Java and C#. We prove preservation properties (which, we consider, were prevalent in the design) prove soundness, and that all executions which do not throw link errors give the same results. Our model is concerned with the interplay of the phases rather than with the particular phases themselves. It is at a higher level than the Java bytecode or the .NET IL. It abstracts from Java multiple loaders and .NET assemblies, and describes the verifier as a type checker, disregarding type inference and data flow analysis issues. It models intermediate code as being interpreted, disregarding the difference between JVM bytecode interpretation, and .NET IL code jit-compilation. It represents dynamic linking *not* necessarily as *is*, but as *perceived* by the source language programmer.

Section 2 introduces Java and C# dynamic linking with an example. Figure 9 lists the main concepts defined. Sections 3 and 4 outline and define the model. Section 5 states properties, soundness, and equivalence. Section 6 concludes. The appendix contains more examples, details for lemmas and proofs.

2 Introduction to the Dynamic Linking Phases

In the presence of dynamic linking, execution can be understood in terms of;

- evaluation, which is not affected by dynamic linking
- loading, which reads classes from the environment
- verification, which checks type-safety of the code
- laying out, which determines object layout and method tables,
- offset calculation, which replaces references to fields and methods in terms of their signature, through the corresponding offsets.

These phases apply to different units of granularity: Loading and laying out apply to classes, whilst verification applies to method bodies, and offset calculation applies to individual member access expressions.

Phases depend on each other: A class can only be laid out after is has been loaded. The offset of a member from a class may only be calculated after that class has been laid out. When verification requires some class to extend a further class it will load the two classes – although [26] suggest a lazier approach of posting constraints instead.

Java	**C#**	**output**
	calc. offset for main	
verify Food ↪ verify main ↪check Meal ≤ Meal ↪check Penne ≤ Penne	jit main ↪ check Meal ≤ Meal ↪load Meal ↪ lay out Meal ↪ check Penne ≤ Penne ↪load Penne; Pasta ↪LoadErr if ¬ Cls ↪ lay out Penne ↪lay out Pasta ↪ calc. offset for eat (Penne)	
calc. offset for main		
execute main	execute main	
		—1—
lay out Meal verify Meal ↪ verify eat (Penne) ↪ check Penne ≤ Pasta ↪load Penne; Pasta ↪LoadErr if ¬ Cls ↪VerifErr, if ¬ Sub ↪ verify chew (Pasta)		
create a new Meal object	create a new Meal object	
		—2—
verify Penne ↪… ↪verify Pasta ↪…		
create a new Penne object	create a new Penne object	
		—3—
calc. offset for eat (Penne)	jit eat (Penne) ↪check Penne ≤ Pasta ↪VerifErr, if ¬ Sub ↪calc. offset for chew (Pasta)	
execute eat (Penne)	execute eat (Penne)	
calc. offset for chew (Pasta)	jit chew (Pasta) ↪calc. offset for **int** cal from Pasta ↪NoFieldErr, if ¬ Fld	
execute chew (Pasta)	execute chew (Pasta)	
		0
		—4—
execute eat (Penne)	execute eat (Penne)	
execute chew (Pasta)	execute chew (Pasta)	
calc. offset for **int** cal from Pasta ↪NoFieldErr, if ¬ Fld		
		100

Table 1. Execution of the program example – with verification

```
class Meal {
    void eat (Penne p){ chew (p); }
    void chew (Pasta p) {
        if (p ==null) print (0);
        else print (p.cal);
    }
}
```

```
class Food {
    public static void main (String[] args) {
        print ("— 1 —"); Meal m = new Meal ();
        print ("— 2 —"); Penne p = new Penne ();
        print ("— 3 —"); m.eat (null);
        print ("— 4 —"); m.eat (p);
    }
}
```

Table 2. Example program

The phases are organized slightly differently in Java than in C#: In Java, offset calculation takes place per instruction, and only before the particular member is accessed, whereas in C# offset calculation takes place per method, and is combined with verification, to give jit-compilation. In Java all methods of a class are verified together, whereas in C# methods are jit-compiled only before being executed. The example from table 2 serves to illustrate these points. It could be expressed in Java and in C#, modulo syntactic modifications[2] and consists of classes Meal and Food, compiled in an environment containing compiled versions of Pasta and Penne:

class Pasta **{ int** cal = 100; **}** **class** Penne **extends** Pasta **{ }**

These classes satisfy the following three requirements:

Cls: Classes Pasta and Penne are present

Sub: Penne is a subclass of Pasta

Fld: Pasta contains a field cal of type **int**

which are crucial for main in Food. E.g., Sub guarantees successful verification of the eat method body, and Fld guarantees successful field access. If Cls, Sub and Fld hold, execution will be successful, and Java and C# will give the same output.

However, the versions of Pasta and Penne available at runtime might differ from those above, and might not satisfy Cls, Sub or Fld: Firstly, Pasta or Penne may not be available, *i.e.*, $\neg$ Cls. Secondly, Penne may not be a subtype of Pasta. *i.e.*, $\neg$ Sub. Thirdly, Pasta may not contain a field **int** cal, *i.e.*, $\neg$ Fld.

These situations will lead to linking errors, detected by the corresponding linking phases. Because these take place at different times in Java and C#, the errors will be reported at different times. This is shown in table 1. The third column contains the output, *e.g.*, — 1 —. The first and second column contain the linking phases as they occur in Java or in C#, with their dependencies indicated through the $\hookrightarrow$ symbol, *e.g.*, in Java, verification of class Meal requires verification of method eat, which in its turn checks that Pasta $\leq$ Pasta, and Penne $\leq$ Penne.

The table shows execution both when Cls, Sub, and Fld hold, and when they do not. Thus, if Cls, Sub, and Fld hold, the two executions will print the same output. However:

[2] The concrete example written in both C# and Java complete with instructions on how to produce the behaviour described in this paper is available at [27].

Verification is "lazier" in C#: Thus, $\neg$ Sub would cause a linking error after —1— in Java, and after —3— in C#. Java verification checks all methods of that class, whereas C# verifies each method when jit-compiling it before its first call.

Offset calculation is "lazier" in Java: Thus, $\neg$ Fld would cause a linking error after —3— in C#, and after —4— in Java. References to fields (or methods) are resolved in Java only when the field is actually accessed during execution, whereas in C# references are resolved when the method containing the reference is jit-compiled.

Subtypes are "optimistic" in Java. Thus, $\neg$ Cls could cause a linking error before —1— in C#, but only after —1— in Java. Checking that a class is a subclass of itself causes loading of the class in C#, but does not in Java.

3 Outline of the Model

With the concept of programs, P, we describe code in all its forms, *i.e.*, the "raw" classes as loaded, the method bodies before and after verification/jit-compilation, and the class layout. Ps map identifiers to classes, and addresses to method bodies. Classes contain their superclass names, and they are either "raw", and contain the signatures of fields and methods, and the method bodies; or, they are "laid out", and contain layout tables which map field and method signatures to offsets and virtual method tables which map offsets to addresses. Global contexts, W, represent the context from which "raw" classes may be loaded[3].

Heaps, H, map addresses to objects. Expressions, e, allow for method call, field access and assignment. Execution reads classes from a global context W, and modifies heaps, expressions, and programs. Therefore, it has form: $P, H, e \rightsquigarrow_W P', H', e'$.

Loading, verification and laying out of classes can be understood as enriching the information in the program, and is represented through judgement $W \vdash P' \leq P$. Loading is represented through an extension of P according to the contents of W. The layout tables are required to extend those of the superclass. Verification and jit-compilation is represented through modification of the method body so as to indicate that it has been verified, and possible substitutions of symbolic references by offsets.

Offset calculation has the format $e \rightsquigarrow_P e'$, meaning that symbolic references in e are replaced by offsets in e', according to the layout tables in P.

Verification/jit-compilation is represented through: $P, e \rightsquigarrow_{W,E} P', e', t$ which means that e is verified/jit-compiled into expression e' and has type t. The program P may need to be extended to P', using information from W. The typing needs a typing environment E. Verification may need to check subtypes: $P, t', t \rightsquigarrow_W P'$ means that t' was established as a subtype of t, and in the process, P was extended to P'.

The model is highly non-deterministic, supporting the description of both languages:

Verification is "lazier" in C#. The model requires methods to have been verified/jit-compiled before being called (fourth rule in fig. 3), thus allowing the C# lazy approach. However, verification is part of program extension (fifth rule in figure 2), and program extension may take place at any time during execution (first rule fig. 3), thus allowing

[3] *i.e.*, the file system, or the registry etc; therefore, W can be viewed as an abstraction over class loaders, or the versioning system.

Expressions

$$
\begin{array}{rcll}
e, e' \in \mathit{Exp} & ::= & \texttt{new}\ c \mid & \text{instance creation} \\
& & \iota \mid & \text{address} \\
& & \texttt{y} \mid & \text{parameter} \\
& & e\ ma(e') \mid & \text{method invocation} \\
& & e\ fa = e' \mid & \text{field assignment} \\
& & e\ fa \mid & \text{field access} \\
& & \texttt{this} \mid & \text{this reference} \\
& & \texttt{nllPExc} \mid & \text{null-pointer exception} \\
& & \texttt{lnkExc} & \text{linking related exception} \\
t, t' \in \mathit{Typ} & ::= & c & \text{type (class name)} \\
ma \in \mathit{Ann}^M & ::= & .m[c, t, t'] \mid & \text{unresolved method annotation} \\
& & [\kappa] & \text{resolved method annotation} \\
fa \in \mathit{Ann}^F & ::= & .f[c, t] \mid & \text{unresolved field annotation} \\
& & [\kappa] & \text{resolved field annotation} \\
a \in \mathit{Ann} & ::= & fa \mid & \text{field annotation} \\
& & ma & \text{method annotation}
\end{array}
$$

$$
\begin{array}{rcll}
c \in \mathit{ClassId} & = & \mathit{Id} & \text{class identifiers} \\
f \in \mathit{FieldId} & = & \mathit{Id} & \text{field identifiers} \\
m \in \mathit{MethId} & = & \mathit{Id} & \text{method identifiers} \\
\iota & \in & \mathbb{N} & \text{addresses} \\
\kappa & \in & \mathbb{N} & \text{offsets}
\end{array}
$$

Programs

$$
\begin{array}{rcll}
P \in \mathit{Prg} & = & (\mathit{ClassId} \to (\mathit{ClassRaw} \uplus \mathit{ClassLaidOut})) & \\
& & \times (\mathbb{N} \to \mathit{Body}) & \text{programs}
\end{array}
$$

$$
\begin{array}{rcll}
\mathit{ClassRaw} & = & \mathit{ClassId} \times \Delta^F \times \Delta^M & \\
\delta^F \in \Delta^F & = & \mathit{FieldId} \to \mathit{Typ} & \text{field descriptions} \\
\delta^M \in \Delta^M & = & \mathit{MethId} \times \mathit{Typ} \times \mathit{Typ} \to \mathit{Exp} & \text{method descriptions}
\end{array}
$$

$$
\begin{array}{rcll}
\mathit{ClassLaidOut} & = & \mathit{ClassId} \times \mathcal{T}^F \times \mathcal{T}^M \times \mathcal{T}^C & \\
\tau^F \in \mathcal{T}^F & = & \mathit{FieldId} \times \mathit{Typ} \to \mathbb{N}^+ & \text{field layout tables} \\
\tau^M \in \mathcal{T}^M & = & \mathit{MethId} \times \mathit{Typ} \times \mathit{Typ} \to \mathbb{N} & \text{method layout tables} \\
\tau^C \in \mathcal{T}^C & = & \mathbb{N} \to \mathbb{N} & \text{code tables}
\end{array}
$$

$$
\begin{array}{rcll}
\mathit{Body} & = & (\mathit{Typ} \times \mathit{Typ} \times \mathit{Exp}) & \text{meth. body before jit/verif.} \\
& & \uplus\ \mathit{Exp} & \text{meth. body after jit/verif.}
\end{array}
$$

Global contexts

$$
W \in \mathit{ClassId} \to \mathit{ClassRaw}
$$

Subtypes

$$
\frac{P(c_1)\downarrow_1 = c_2}{P \vdash c_1 \leq c_1} \qquad \frac{P \vdash c_1 \leq c_2 \quad P \vdash c_2 \leq c_3}{P \vdash c_1 \leq c_3}
$$

Figure 1. Expressions, programs, subtypes

the Java approach too. Of course, it also allows further behaviour, *e.g.*, where only some methods are verified/jit-compiled, or where classes are verified upon loading.

Offset calculation is "lazier" in Java. The model combines verification and jit-compilation into one judgment, $P, e \rightsquigarrow_{W,E} P', e', t$, which requires offset calculation for its subexpressions (third to sixth rule in fig. 5). This describes C# jit-compilation. However, offset calculation may also leave the expression unmodified (last rule in fig. 4), and that describes Java verification.

On the other hand, offset calculation may take place during execution (last rule in fig. 3), and the operational semantics for member access requires the offset to have been calculated (fourth and fifth rule in fig. 3). This describes Java offset calculation.

The model allows many more executions, *e.g.*, offsets may be calculated even if not required, and verification/jit-compilation may replace only some of the symbolic references by offsets.

Subtypes are "optimistic" in Java. The model considers any class identifier a subtype of itself (last rule in fig. 5); thus reflecting Java. However, programs may be extended during verification (penultimate rule in fig. 5), thus reflecting C#.

Timing of link-related actions The model allows loading, jit-compilation, verification, and offset calculation to take place at any time (first rule in fig. 3), even if not needed. It also allows throwing linking exceptions (not null pointer exceptions) at any time (second rule in fig. 3), without requiring them to be necessary and without distinguishing the reason. This non-determinism does not reflect practical implementations but simplifies the model considerably.

4 The model

All mappings are partial; $dom(f)$, $rng(f)$ denote the domain and range of function f.

Expressions The syntax is given in figure 1. It allows for classes, subclasses, methods and fields, and describes an imperative language.

We use an augmented high level language, near to source code. The augmentations are memory offsets, and type annotations, used to disambiguate fields or methods with the same name. For example, the expression p.cal [Pasta,**int**] denotes the field called cal of p, of type **int**, and declared in class Pasta. This symbolic reference will be replaced during offset calculation; *e.g.*, if **int** cal has offset 3 in class Pasta then the expression will be rewritten to p[3].

Values are addresses, which are natural numbers denoted by ι, ι' *etc*; the null pointer is $\mathbf{0}$. `nllPExc` is the exception raised when a field is accessed or a method is called on $\mathbf{0}$. Also, `lnkExc` stands for, and does not distinguish between, any link related exception, ie verification errors, class not found, class circularities, absence of fields and methods.

Programs reflect the internal representation of code. They are described in figure 1. They map identifiers to raw ($ClassRaw$) or to laid out classes ($ClassLaidOut$), and addresses to method bodies.[4] Raw classes correspond to *.class or *.dll files. They consist

[4] Throughout this paper, we extract components out of tuples of mappings implicitly when needed, *e.g.*, $P(c)$ is a shorthand for $P{\downarrow_1}(c)$, and $P(\iota)$ is a shorthand for $P{\downarrow_2}(\iota)$.

$$\frac{}{W \vdash P \leq P} \qquad \frac{W \vdash P' \leq P'' \quad W \vdash P'' \leq P}{W \vdash P' \leq P}$$

$$\frac{\begin{array}{l} P = c = P' \\ P(c) = \epsilon \\ P'(c) = W(c) = \langle c_s, \_, \_ \rangle \\ P(c_s) \neq \epsilon \end{array}}{W \vdash P' \leq P} \qquad \frac{\begin{array}{l} P = c = P' \\ P(c) = \langle c_s, \delta^F, \delta^M \rangle,\ P'(c) = \langle c_s, \tau^F, \tau^M, \tau^C \rangle \\ P(c_s) = \langle \_, \tau^F_s, \tau^M_s, \tau^C_s \rangle \\ \tau^F \text{ injective},\ dom(\tau^F) = \{\langle f, t \rangle \mid \delta^F(f) = t\} \\ rng(\tau^F) \cap FdOffs(P, c_s) = \emptyset \\ \tau^M \leq \tau^M_s \text{ wrt } dom(\delta^M), \quad \tau^C \preceq \tau^C_s \text{ wrt } \tau^M(dom(\delta^M)) \\ \delta^M(m, t, t') = e \implies \exists \iota : \\ \quad \tau^C(\tau^M(m, t, t')) = \iota, P(\iota) = \epsilon,\ P'(\iota) = (t, t', e) \end{array}}{W \vdash P' \leq P}$$

$$\frac{\begin{array}{l} P = \iota = P' \\ P(\iota) = \langle t_r, t_p, e \rangle, \quad P'(\iota) = e' \\ \exists c \in dom(P') : \\ \quad P(c') = \langle \_, \_, \_, \tau^C_1 \rangle,\ \iota \in rng(\tau^C_1) \implies P \vdash c' \leq c \\ \quad P', e \rightsquigarrow_{W, \{\texttt{this} \mapsto c, \texttt{y} \mapsto t_p\}} P', e', t \\ \quad P' \vdash t \leq t_r \end{array}}{W \vdash P' \leq P}$$

Figure 2. Program extension

of the superclass name, the field descriptions ($\delta^F \in \Delta^F$) consisting of field identifiers and types, and method descriptions ($\delta^M \in \Delta^M$) consisting of method identifier, argument type, return type and method body. Laid out classes consist of a field layout table ($\tau^F \in \mathcal{T}^F$), which determines the offset for a field with given identifier and type, the method layout table ($\tau^M \in \mathcal{T}^M$) which maps method signatures to offsets, and the virtual table ($\tau^C \in \mathcal{T}^C$), which maps offsets to addresses of method bodies. [5]

Method bodies which have not been checked consist of a signature and expression, $Typ \times Typ \times Exp$. Bodies which have been checked consist of an expression, Exp.

Execution modifies the current program, expression and heap. It therefore has the form

$$P, H, e \rightsquigarrow_W P', H', e'$$

expressing that the global context may be used for program extension. It is defined through small step semantics in figure 3.

Heaps, H, map addresses to objects, which are memory blocks consisting of class identifier, and values for the fields. Values are object addresses, or **0**. Heaps have form:

$$H : \mathbb{N}^+ \rightarrow \mathbb{N} \uplus ClassId.$$

If $H(\iota) = c \in ClassId$ then ι points to an object of class c. The fields of that object are stored at some offset, κ, from ι. An address ι is fresh in H iff $\forall \kappa : H(\iota + \kappa) = \epsilon$.

[5] Appendix B contains an example clarifying descriptions and layout tables in the presence of method inheritance and field overriding.

$$\frac{W \vdash P' \leq P}{P, H, e \rightsquigarrow_W P', H, e} \qquad \frac{}{P, H, e \rightsquigarrow_W P, H, \texttt{lnkExc}}$$

$$\frac{FdOffs(P, c) = \{\kappa_1, \ldots, \kappa_n\}, \quad \iota \text{ free in } H}{P, H, \texttt{new } c \rightsquigarrow_W P, H[\iota \mapsto c, \iota{+}\kappa_1 \mapsto \mathbf{0}, \ldots \iota{+}\kappa_\mathbf{n} \mapsto \mathbf{0}], \iota}$$

$$\frac{\begin{array}{l} H(\iota) = c \\ P(c) = \langle \_, \_, \_, \tau^C \rangle \\ P(\tau^C(\kappa)) = e \end{array}}{P, H, \iota[\kappa](\iota') \rightsquigarrow_W P, H, e[\iota/\texttt{this}, \iota'/\texttt{y}]} \qquad \begin{array}{l} \iota \neq 0 \\ \hline P, H, \iota[\kappa] \rightsquigarrow_W P, H, H(\iota + \kappa) \\ P, H, \iota[\kappa] = \iota' \rightsquigarrow_W P, H[\iota + \kappa \mapsto \iota'], \iota' \end{array}$$

$$\frac{}{\begin{array}{l} P, H, \mathbf{0}[\kappa] \rightsquigarrow_W P, H, \texttt{nllPExc} \\ P, H, \mathbf{0}[\kappa] = \iota \rightsquigarrow_W P, H, \texttt{nllPExc} \\ P, H, \mathbf{0}[\kappa](\iota) \rightsquigarrow_W P, H, \texttt{nllPExc} \end{array}} \qquad \begin{array}{l} \dfrac{P, H, e \rightsquigarrow_W P', H', e'}{P, H, \sqsubset e \sqsupset^{exe} \rightsquigarrow_W P', H', \sqsubset e' \sqsupset^{exe}} \\ \\ \dfrac{z = \texttt{nllPExc}, \text{ or } z = \texttt{lnkExc}}{P, H, \sqsubset z \sqsupset^{exe} \rightsquigarrow_W P', H', z} \end{array}$$

$$\frac{a \rightsquigarrow_P a'}{P, H, \sqsubset a \sqsupset^{off} \rightsquigarrow_W P, H, \sqsubset a' \sqsupset^{off}}$$

Figure 3. Execution.

The following heap, H_0, contains a Penne object at *2*, and a Food object at *4*:

$H_0(2)$ = Penne start Penne object $\quad H_0(3)$ = *55* field **int** cal from Pasta
$H_0(4)$ = Food start Food object $\quad H_0(\iota)$ = ϵ for all other ι 's

Thus, as in [7], heaps are modelled at a lower level than in verifier studies [30,15,26], where objects are indivisible entities, and where there are no address calculations. Our lower level model can describe the potential damage when executing unverified code.

Program Extension We need auxiliary concepts of mapping extension ($g' \leq g$ wrt A, $g' \preceq g$ wrt A), and program equality up to class or address ($P =c= P'$, $P =\iota= P'$):

Definition 1 *For injective mappings g, g', set A, and for P, P', and ι, and c :*

- $g' \leq g$ *wrt* A, *iff* $dom(g') = dom(g) \cup A$, *and* $\forall y \in dom(g) : \ g'(y) = g(y)$. [6]
- $g' \preceq g$ *wrt* A, *iff* $dom(g') = dom(g) \cup A$, *and* $\forall y \in dom(g) \setminus A : \ g'(y) = g(y)$.
- $P =\iota= P'$ *iff* $\forall c : \ P(c) = P'(c)$, *and* $\forall \iota' \in dom(P) \setminus \{\iota\} : \ P(\iota') = P'(\iota')$.
- $P =c= P'$ *iff* $\forall c' \neq c : \ P(c') = P'(c')$, *and* $\forall \iota \in dom(P) : \ P(\iota) = P'(\iota)$.

A program P' extends another program P, if P' contains more information (through loading of classes), or more refined information (through verification, jit-compilation or layout calculation) than P. This relationship has the format

$$W \vdash P' \leq P$$

[6] Note, that the sets A and $dom(g)$ need *not* be disjoint.

c.f. figure 2, and is defined in the global context of a W which expresses the environment (possibly a file system) from which classes are loaded. [7]

In more detail, $W \vdash P' \leq P$ if: 1) P' is in the reflexive, transitive closure of the relation. 2) P' and P are identical up to c, a raw class read from W whose superclass (c_s) is already in P. 3) P' and P are identical up to class c, and a) the field layout of c extends that of c_s and fields introduced by c get fresh offsets, b) the method layout of c extends that of c_s, c) all methods in c which override (have the same signature as) methods in c_s are mapped to new addresses. 4) P' and P are identical up to address ι, and P (ι') contains the verified/jit-compiled version of the method at P (ι).

The first rule of figure 3 says that programs may be extended at any time. The second rule allows linking exceptions to be thrown at any time. This is, of course, highly non-deterministic, and allows linking phases or errors even if unnecessary.

Evaluation is the part of execution that is not directly affected by dynamic linking. It is described by the third through eighth rule in figure 3. It requires the following auxiliary function which collects the field offsets from all superclasses:

$$\mathit{FdOffs}(P, c) = \bigcup_{P \vdash c \leq c'} rng(P(c') \downarrow_2)$$

Creation of a new object of class c, `new` c, allocates fresh addresses for the fields of c at the corresponding offsets, initializing them with $\mathbf{0}$.

Method call, $\iota[\kappa](\iota')$, looks up the method body e in the dynamic class of the receiver ι, using the offset κ, and executes that body after replacing `this` by the actual receiver ι, and the parameter `y` by the argument ι'. Therefore, evaluation only applies to expressions which do *not* contain `this`, or `y`. The format of the call $\iota[\kappa](\iota')$ (rather than $\iota.m[c, t_r, t_p](\iota')$) means that the offset has been calculated. The requirement $P(c) = \langle \_, \_, \_, \tau^C \rangle$ (rather than $P(c) = \langle \_, \_, \_ \rangle$) means that the class c has been laid out. The requirement that $P(\tau^C(\kappa)) = e$ (rather than $P(\tau^C(\kappa)) = \langle \_, \_, \_ \rangle$) means that the particular method has been verified/jit-compiled.

Field lookup retrieves the contents of the heap at the given offset, whereas field assignment updates the heap at the given offset, as in the fifth rule. Method call and field access for $\mathbf{0}$ throw a `nllPExc`, as described in the sixth rule of the table.

Execution is propagated to its context, as described in the seventh rule. Both link related, and unrelated exceptions (*i.e., z*) are propagated out of their contexts, as described in the eighth rule. Execution contexts allow a succinct description of propagation:

$$\begin{array}{rcl} \sqsubset\cdot\sqsupset^{exe} ::= \sqsubset\cdot\sqsupset^{exe} ma(e) & | & \iota\, ma(\sqsubset\cdot\sqsupset^{exe}) \;\; | \\ \sqsubset\cdot\sqsupset^{exe} fa = e & | & \iota\, fa = \sqsubset\cdot\sqsupset^{exe} \;\; | \;\; \sqsubset\cdot\sqsupset^{exe} fa \end{array}$$

Offset Calculation replaces a symbolic reference through an offset, and has format

$$a \rightsquigarrow_P a'$$

[7] The particular environment is not needed for the proof of soundness - it was omitted *e.g.*, in the model in [7], but is needed when formulating and proving equivalence of strategies.

$$\frac{P(c) = \langle \_, \tau^F, \_, \_ \rangle \quad \tau^F(f,t) = \kappa}{.f[c,t] \rightsquigarrow_P [\kappa]} \qquad \frac{P(c) = \langle \_, \_, \tau^M, \_ \rangle \quad \tau^M(m,t_r,t_p) = \kappa}{.m[c,t_r,t_p] \rightsquigarrow_P [\kappa]} \qquad \frac{}{a \rightsquigarrow_P a}$$

Figure 4. Offset calculation.

where a represents a field or method annotation. Figure 4 says that for fields, we look up the name of the field and its type in the class, whilst for methods we look up the name, argument type and result type in the class[8]. The last rule allows a to be left unmodified.

The last rule in 3 allows offset calculation to happen during execution, as in Java. For this, we have defined appropriate notion of offset calculation contexts as

$\sqsubset\cdot\sqsupset^{\mathit{off}} ::= e \sqsubset\cdot\sqsupset^{\mathit{off}} \mid e \sqsubset\cdot\sqsupset^{\mathit{off}} = e \mid e \sqsubset\cdot\sqsupset^{\mathit{off}}(e)$

Offset calculation also happens during jit-compilation, (figure 5) thus modelling C#. Combining this with the rule that leaves offsets unmodified we model Java verification which does not calculate the offsets.

Verification and Jit-Compilation We describe the similarities between Java verification and C# jit-compilation through the verification/jit-compilation judgment:

$$P, e \rightsquigarrow_{W,E} P', e', t$$

defined in figure 5, which transforms an expression e to e', type checks e to have type t, and may extend the program P to P'. The process takes place in an environment E which maps `this` and the parameter `y` to types, *i.e.*, $E : \{ \texttt{this}, \texttt{y} \} \rightarrow \mathit{Typ}$, and in the global context W.

The parameter `y` and the receiver `this` have the type given in the environment E. Verification/jit-compilation of an object creation expression requires c to be a class, and gives it type c. The value **0** has any class type c.

Method call requires the receiver and argument to be well-typed, and to be of subtypes of c and t_p, the receiver and argument types stored in the symbolic method annotation $.m[c, t_r, t_p]$. The method call has type t_r, the result type of the annotation. The symbolic annotation may be replaced by an offset, thus modeling C# jit-compilation. Offset calculation also allows for the identity, thus modeling Java verification. Similar explanations apply to the rules which access fields.

Finally, verification may require classes to be loaded, and the offset calculation may require layout information about some classes. This is described through the sixth rule, which allows extension of the program at any time.

Verification/jit-compilation may need to check that a type is a subtype of another type, and while doing so may need to load further classes, as in judgment:

$$P, t_1, t_2 \rightsquigarrow_W P'$$

also given in figure 5. Notice, that this judgment allows any identifier to be a subtype of itself even if not loaded - this follows the "optimistic" Java approach.

[8] Thus, a class may inherit or define several methods with same names, argument type but different result type, and it may inherit fields with same name as its own fields.

$$\frac{}{P, \texttt{this} \rightsquigarrow_{W,E} P, \texttt{this}, E(\texttt{this})} \qquad \frac{P, c, c \rightsquigarrow_W P'}{P, \texttt{new}\ c \rightsquigarrow_{W,E} P', \texttt{new}\ c, c}$$

$$P, \texttt{y} \rightsquigarrow_{W,E} P, \texttt{y}, E(\texttt{y}) \qquad P, 0 \rightsquigarrow_{W,E} P', 0, c$$

$$\frac{\begin{array}{l} P, e_1 \rightsquigarrow_{W,E} P_1, e_1', t_1 \\ P_1, e_2 \rightsquigarrow_{W,E} P_2, e_2', t_2 \\ P_2, t_1, c \rightsquigarrow_W P_3 \\ P_3, t_2, t_f \rightsquigarrow_W P' \\ .f[c, t_f] \rightsquigarrow_{P'} fa \end{array}}{P, e_1.f[c, t_f] = e_2 \rightsquigarrow_{W,E} P', e_1'\, fa = e_2', t_f} \qquad \frac{\begin{array}{l} P, e \rightsquigarrow_{W,E} P_1, e', t_e \\ P_1, t_e, c \rightsquigarrow_W P' \\ .f[c, t_f] \rightsquigarrow_{P'} fa \end{array}}{P, e.f[c, t_f] \rightsquigarrow_{W,E} P', e'\, fa, t_f}$$

$$\frac{\begin{array}{l} P, e_1 \rightsquigarrow_{W,E} P_1, e_1', t_1 \\ P_1, e_2 \rightsquigarrow_{W,E} P_2, e_2', t_2 \\ P_2, t_1, c \rightsquigarrow_W P_3 \\ P_3, t_2, t_p \rightsquigarrow_W P' \\ .m[c, t_r, t_p] \rightsquigarrow_{P'} ma \end{array}}{P, e_1.m[c, t_r, t_p](e_2) \rightsquigarrow_{W,E} P', e_1'\, ma(e_2'), t_r} \qquad \frac{\begin{array}{l} W \vdash P'' \leq P \\ P'', e \rightsquigarrow_{W,E} P', e', t \end{array}}{P, e \rightsquigarrow_{W,E} P', e', t}$$

$$\frac{\begin{array}{l} W \vdash P' \leq P \\ P' \vdash t' \leq t \end{array}}{P, t', t \rightsquigarrow_W P'} \qquad \frac{}{P, t, t \rightsquigarrow_W P}$$

Figure 5. Verification and Jit-compilation.

5 Soundness and Equivalence of Strategies

The judgment $\vdash P$ defined in fig. 8 guarantees that program P is well formed, *i.e.*, that 1) the class `Object`is defined and has itself as a superclass, 2) all superclasses are present, and the subclass relationship is acyclic except for `Object`, 3) for any laid out class c with superclass c_s the fields and methods have distinct offsets, the methods defined in c_s have the same offsets in c, and 3) all method bodies which are considered as already verified/jit-compiled, *i.e.*, for which $P\ (\iota)$=e, can be verified/jit compiled, albeit without program extension, and therefore in the empty global context, $\emptyset$.

Figure 7 defines conformance. The judgment $P, H \vdash \iota$ expresses that the object stored at ι conforms to its class, c, as stored in H (ι). For all fields of c, the object must contain appropriate values at the corresponding offsets, and no other object may be stored between its fields. The judgment $P \vdash H$ requires all objects to conform to their class, and (implicitly) that the class of any objects stored in H is defined in P. Notice, that **0** conforms to any class, allowing objects with a field initialized to **0**, to belong even to a class that has not been loaded yet.

Types for runtime expressions are given by judgment $P, H \vdash e : t$, from fig. 6, with rules similar to those for verification/jit-compilation, with the difference that heaps *are* taken into account (to give types to addresses), environments are *not* taken into account (runtime expressions do not contain `this`, or `y`), and the program is *not* extended.

$$\frac{}{P, H \vdash \mathbf{0} : c \quad P, H \vdash \texttt{new}\ c : c} \qquad \frac{P, H \vdash \iota \quad P \vdash c \leq c' \quad H(\iota) = c}{P, H \vdash \iota : c'} \qquad \frac{P, H \vdash e : c' \quad P \vdash c' \leq c}{P, H \vdash e.f[c, t] : t} \qquad \frac{P, H \vdash e : c \quad \mathit{TypeOfFd}(P, c, \kappa) = t}{P, H \vdash e[\kappa] : t}$$

$$\frac{P, H \vdash e\,fa : t \quad P, H \vdash e' : t' \quad P \vdash t' \leq t}{P, H \vdash e\,fa = e' : t} \qquad \frac{P, H \vdash e_1 : c_1 \quad P, H \vdash e_2 : t_2 \quad P \vdash c_1 \leq c \quad P \vdash t_2 \leq t_p}{P, H \vdash e_1.m[c, t_r, t_p](e_2) : t_r} \qquad \frac{P, H \vdash e_1 : c_1 \quad P, H \vdash e_2 : t_2 \quad P \vdash t_2 \leq t_p \quad P(c_1) = \langle \_, \_, \tau^M, \_ \rangle \quad \tau^M(\langle ..., t_r, t_p \rangle) = \kappa}{P, H \vdash e_1[\kappa](e_2) : t_r}$$

Figure 6. Types of runtime expressions.

Runtime expressions containing field access offsets are typed using:

$$\mathit{TypeOfFd}(P, c, \kappa) \;=\; t \text{ if } P(c')\downarrow_2 (\_, t) = \kappa \text{ for } P \vdash c \leq c', \quad \epsilon \text{ otherwise}$$

Runtime expressions containing offsets for method call are typed by application of the inverse layout function (in well-formed programs these are injective, hence their inverses are defined).

It is easy to prove that that verification/jit-compilation and execution extend programs, *c.f.* lemma 1 in appendix D. Properties such as subtyping, conformance of the heap, runtime type of an expression, verification of an expression, or well-formedness of a program, established in a program P are preserved in an extending program P' [9]. This is shown in lemma 2 in appendix D.

Thus, we obtain that execution of *any* expression preserves well-formedness of programs. Finally, a verified expression preserves its runtime type, when the receiver and argument have been replaced by appropriate addresses, *c.f.* lemma 3 in appendix D.

Execution of a well-typed expression e does not overwrite objects, creates new objects in the free space, and does *not* affect the type of any expression e'' – even if e'' were a subexpression of e, *c.f.* lemma 4 in appendix D. [10]

Subject reduction guarantees that the heap H' preserves conformance, uninitialized parts of the store are never dereferenced, and the expression preserves its type.

Theorem 1 *If* $P \vdash H$, *and* $\vdash P$, *and*
$P, H \vdash e : t$, *and*
$P, H, e \rightsquigarrow_W P', H', e'$
then
$P' \vdash H'$, *and*
if e' *does not contain an exception, then* $\exists\, t' : \quad P', H' \vdash e' : t', \quad P' \vdash t' \leq t.$

[9] Similar properties were proven in [8], used in [5,7], and explored in our model of binary compatibility [9]. Notice, that for source code, such properties do not always hold [3].

[10] This is required for type soundness in imperative object oriented languages, and was proven, *e.g.*, in [8,31,7]. In the current work it holds only for *well-typed* expressions e.

$$\frac{P \vdash c' \leq c \quad H(\iota) = c'}{P, H \vdash \iota \triangleleft c} \qquad \frac{}{P, H \vdash \mathbf{0} \triangleleft c} \qquad \frac{\begin{array}{l} H(\iota) = c \\ P(c) = \langle \_, \tau^F, \_, \_ \rangle \\ \forall \kappa{:}\, \mathit{TypeOfFd}(P, c, \kappa) = t \implies P, H \vdash \iota + \kappa \triangleleft t \\ 1 \leq \kappa \leq max(\mathit{FdOffs}(P, c)) \implies H(\iota + \kappa) \notin \mathit{ClassId} \end{array}}{P, H \vdash \iota}$$

$$\frac{H(\iota) \in \mathit{ClassId} \implies P, H \vdash \iota}{P \vdash H}$$

Figure 7. Conformance

In theorem 2 we prove that nondeterminism does *not* affect the result of evaluations which do not throw link related exceptions. The global context W needs to be explicitly stated here. The theorem does *not* apply for intermediate results, nor if ν were a link related exception – several counterexamples were shown in section 2.

Theorem 2 *For* e, P, P', P'', H, H', H'', ι , *and* $\nu, \nu' \in \mathbb{N} \cup \{\texttt{nllPExc}\}$, *if*:

$$P, H, e \rightsquigarrow^*_W P', H', \nu,$$
$$P, H, e \rightsquigarrow^*_W P'', H'', \nu',$$

then:

$$\nu = \nu', H' = H'' \quad \textit{up to renaming of addresses.}$$

Finally, in theorem 3 in appendix D we prove that environments which are identical in the parts required for execution, can lead to identical results.

6 Conclusions, related work, and further work

Dynamic linking is a relatively new, very powerful language feature with complex semantics, which needs to be well understood. Our model is simple, especially considering the complexity of the feature, and compared to an earlier model for Java [7].

We have achieved simplicity through many iterations, and through the choice of appropriate abstractions: 1) we do not distinguish the causes of link related exceptions, 2) we allow link-related exceptions to be thrown at *any* time of execution, even when there exist other, legal evaluations, 3) we do not prescribe at which point of execution the program will be extended, and so allow "unnecessary" loading, verification or jit-compilations, 4) we combine in the concept of "program" loaded, verified, and laid out code, 5) we represent programs through mapping rather than texts or data structures. Most of these abstractions were introduced primarily to allow the model to serve for both Java and for C#, and had the agreeable effect of significant simplification.

Non-determinism seems to have been in the the Java designers' minds: the specification [22], sect. 12.1.1 requires resolution errors to be thrown only when linking actions related to the error are required. Through non-determinism we distilled the main ingredients of dynamic linking in both languages, and their dependencies. We prove type soundness, thus obtaining type soundness both for the Java and the C# strategies, and showed that different strategies within the model do not differ widely.

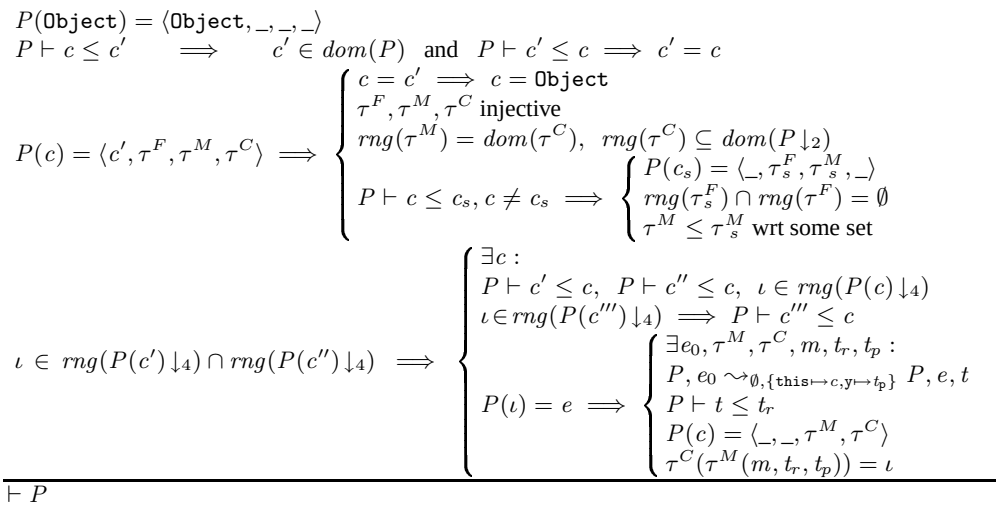

Figure 8. Well-formed programs

Extensive literature is devoted to the Java verifier [30,16]. Dynamic loading in Java is formalized in [19], while problems with security in the presence of multiple loaders are reported in [29], a solution presented in [21], which is found flawed and improved upon in [26]. Type safety for a substantial subset of the .NET IL is proven in [17].

Interest in linking as part of the program lifecycle was kindled through [4]. Separate compilation for Java is discussed in [1]. Module interconnection languages, and mixins [33,2,12,10,13] give explicit control of program composition at source code level.

Dynamic linking gave rise to the concept of binary compatible changes, [14], and [22], sect. 13, *i.e.,* changes which do not introduce more linking errors than the code being replaced; the concept is explored in [9]. Tools that load most recent binary compatible versions of code were developed for Java [28] and C# [20]. Current JVMs go even further, and support replacing a class by a class of the same signature, as a "fix-and-continue" feature [6]. Dynamic software updates [18] support type safe dynamic reloading of code whose type may have changed, while the system is running.

Further work includes a better understanding of binary compatible library developments, extension of the model to allow verification also posting constraints which have to be satisfied upon class loading, as suggested in [26], or to allow field lookup to examine the tables of superclasses as in some of the JVMs , the incorporation of C# assemblies and modules, extensions of the model so as to avoid unnecessary linking steps, and "concretization" of the model so as to obtain Java or C# behaviour.

Acknowledgements

We are indebted to Vladimir Jurisic, Davide Ancona and Elena Zucca for crucial information and suggestions, and to Christopher Anderson and Mark Skipper for feedback.

References

1. Davide Ancona, Giovanni Lagorio, and Elena Zucca. A Fromal Framework for Java Separate Compilation. In *ECOOP 2002*, June 2002.
2. Davide Ancona and Elena Zucca. A calculus of module systems. *Journal Functional Programming*, 12(3):91–132, 2002.
3. Gilad Bracha. Packages break strong locality, February 1999. Private Communication, more at http://www-dse.doc.ic.ac.uk/sue/packages.html.
4. Luca Cardelli. Program Fragments, Linking, and Modularization. In *POPL'97 Proceedings*, January 1997.
5. Drew Dean. The Security of Static Typing with Dynamic Linking. In *Fourth ACM Conference on Computer and Communication Security*, 1997.
6. Mikhail Dimitriev. Hotspot Technology Application for Advanced Profiling. In *ECOOP USE Workhop*, June 2002.
7. Sophia Drossopoulou. An Abstract model of Java dynamic Linking and Loading. In Robert Harper, editor, *Types in Compilation, Third International Workshop, Revised Selected Papers*. Springer, 2001.
8. Sophia Drossopoulou, Susan Eisenbach, and Sarfraz Khurshid. Is Java Sound? *Theory and Practice of Object Systems*, 5(1), January 1999.
9. Sophia Drossopoulou, Susan Eisenbach, and David Wragg. A Fragment Calculus - towards a model of Separate Compilation, Linking and Binary Compatibility. In *LICS Proceedings*, 1999.
10. Dominic Duggan. Sharing in Typed Module Assembly Language. In *Preliminary Proceedings of the Third Workshop on Types in Compilation (TIC 2000)*. Carnegie Mellon, CMU-CS-00-161, 2000.
11. G. Fenton and E. Felton. *Securing Java Getting Down to Business with Mobile Code*. John Wiley and Sons, 1999.
12. Kathleen Fisher, John Reppy, and Jon Riecke. A Calculus for Compiling and Linking Classes. In *ESOP Proceedings*, March 2000.
13. Matthew Flatt, Shiram Khrishnamurthi, and Matthias Felleisen. Classes and Mixins. In *POPL Proceedings*, January 1998.
14. Ira Forman, Michael Conner, Scott Danforth, and Larry Raper. Release-to-Release Binary Compatibility in SOM. In *OOPSLA Proceedings*, October 1995.
15. Stephen N. Freund and J. C. Mitchell. A Formal Framework for the Java Bytecode Language and Verifier. In *OOPSLA Proceeedings*, November 1999.
16. Stephen N. Freund and J. C. Mitchell. A Type System for Object Initialization in the Java Bytecode Language. In *OOPSLA Proceeedings*, October 1998.
17. Andrew Gordon and Don Syme. Typing a multi-language intermediate code. In *Principles of programming Languages 2001*, pages 248–260. ACM Press, 2001.
18. Michael Hicks, Jonathan T. Moore, and Scott Nettles. Dynamic Software Updating. In *Programming Language Design and Implementation*. ACM, 2001.
19. Thomas Jensen, Daniel Le Metyayer, and Tommy Thorn. A Formalization of Visibility and Dynamic Loading in Java. In *IEEE ICCL*, 1998.
20. V. Jurisic. Deja-vu.NET: A Framework for Evolution of Component Based Systems. http://www.doc.ic.ac.uk/˜ajf/Teaching/Projects/DistProjects.html, June 2002.
21. Sheng Liang and Gilad Bracha. Dynamic Class Loading in the Java^{TM} Virtual Machine. In *OOPSLA Proceedings*, October 1998.
22. Tim Lindholm and Frank Yellin. *The Java Virtual Machine*. Addison-Wesley, 1999.
23. Type Safety and Security. In *MSDN Magazine*, 2001.

class A **{**
  A f1;
  A f2;
  C f3;
  A m1(B y) **{** e_1 **}**
  B m1(A y) **{** e_2 **}**
}
	class B **extends** A **{**
	  A f1;
	  B f2;
	  B f4;
	  A m1(B y) **{** e_3 **}**
	  B m2(B y) **{** e_4 **}**
	}

Table 3. Example program demonstrating layout

24. M. Pietrek. Avoiding DLL Hell: Introducing Application Metadata in the Microsoft .NET Framework. In *MSDN Magazine*, msdn.microsoft.com/, October 2000.
25. S. Pratschner. Simplifying Deployment and Solving DLL Hell with the .NET Framework. msdn.microsoft.com/, November 2001.
26. Zhenyu Qian, Allen Goldberg, and Alessandro Coglio. A Formal Specification of Java^{TM} Class Loading. In *OOPSLA'2000*, November 2000.
27. S. Eisenbach. Example Programs in C# and Java that show phases in dynamic linking. http://www-dse.doc.ic.ac.uk/~sue/foodexample.html, September 2002.
28. C. Sadler S. Eisenbach and S. Shaikh. Evolution of Distributed Java Programs. In *IEEE Working Conf on Component Deployment*, June 2002.
29. Vijay Saraswat. Java is not type-safe. Technical report, AT&T Rresearch, 1997. http://www.research. att.comp/~vj/bug.html.
30. Raymie Stata and Martin Abadi. A Type System For Java Bytecode Subroutines. In *POPL'98 Proceedings*, January 1998.
31. Donald Syme. Proving Java Type Sound. In Jim Alves-Foss, editor, *Formal Syntax and Semantics of Java*, volume 1523 of *LNCS*. Springer, 1999.
32. T. Van Vleck(*ed.*). Multics Home. www.multicians.org/, August 2002.
33. Joe Wells and Rene Vestergaard. Confluent Equational Reasoning for Linking with First-Class Primitive Modules. In *ESOP Proceedings*, March 2000.

A Overview of terms and judgments of this paper

In figure 9 we give an overview of terms and judgments defined in that paper.

B An example demonstrating program extension, description and layout tables

The following example aims to demonstrate some fine points about method and field tables. The source code is given in table 3.

We have two classes, A and B, where A has three fields, f1, f2 and f3. Class B shadows f1 with a field of the same type and f2 with a field of a different type – as we shall the types of the shadowed fields do not affect their treatment; also, B introduces a further field f4. Class A introduces the overloaded method m1: there are two versions, one with argument type B, and one with argument type A. The method m1 with argument type A is overridden in B; B also introduces a further method m2.

e	expressions	fig. 1
t	types	fig. 1
ι	addresses	fig. 1
κ	offsets	fig. 1
`nllPExc`	the null-pointer exception	fig. 1
`lnkExc`	link-related exception, *e.g.*, verification, load err. *etc*	fig. 1
fa, ma, a	field, method, or any annotation	fig. 1
δ^F	field descriptions	fig. 1
δ^M	method descriptions	fig. 1
τ^F	field layout tables	fig. 1
τ^M	method layout tables	fig. 1
τ^C	code tables	fig. 1
H	heaps	sec. 4
E	environment giving types to receiver and argument	sec. 4
$\sqsubset\cdot\sqsupset^{exe}$	execution context	sect. 4
$\sqsubset\cdot\sqsupset^{off}$	offset calculation context	sect. 4
$P, H, e \rightsquigarrow_W P', H', e'$	execution in global context W	fig 3
$a \rightsquigarrow_P a'$	offset calculation	fig. 4
$P, e \rightsquigarrow_{W,E} P', e', t$	verification or jit-compilation	fig. 5
$P, t', t \rightsquigarrow_W P'$	t' is a subtype of t, while extending program P to P'	fig. 5
$W \vdash P' \leq P$	program P' extends program P in global context W	fig. 2
$P \vdash t' \leq t$	in program P the type t' is a subtype of t	fig. 1
$\vdash P$	well formed program	fig. 8
$P \vdash H$	well formed heap H for the program P	fig. 7
$P, H \vdash e : t$	runtime expression e has type t in the context of P and H	fig. 6
$P, H \vdash \iota \triangleleft c$	ι conforms class c, or subclass	fig. 7
$g' \leq g$ wrt A	mapping g' injectively extends g into set A, preserving $dom(g)$	def. 1
$g' \preceq g$ wrt A	mapping g' injectively extends g into set A, preserving $dom(g) \setminus A$	def. 1
$P =\iota= P'$	P and P' agree up to address ι	def. 1
$P =c= P'$	P and P' agree up to class c	def. 1
$FdOffs(P, c)$	the set of all offsets allocated for the fields of c in P	page 10
$TypeOfFd(P, c, \kappa)$	the type of the field contained at the offset κ of c in P	page 13

Figure 9. Overview of terms and judgments used in this paper

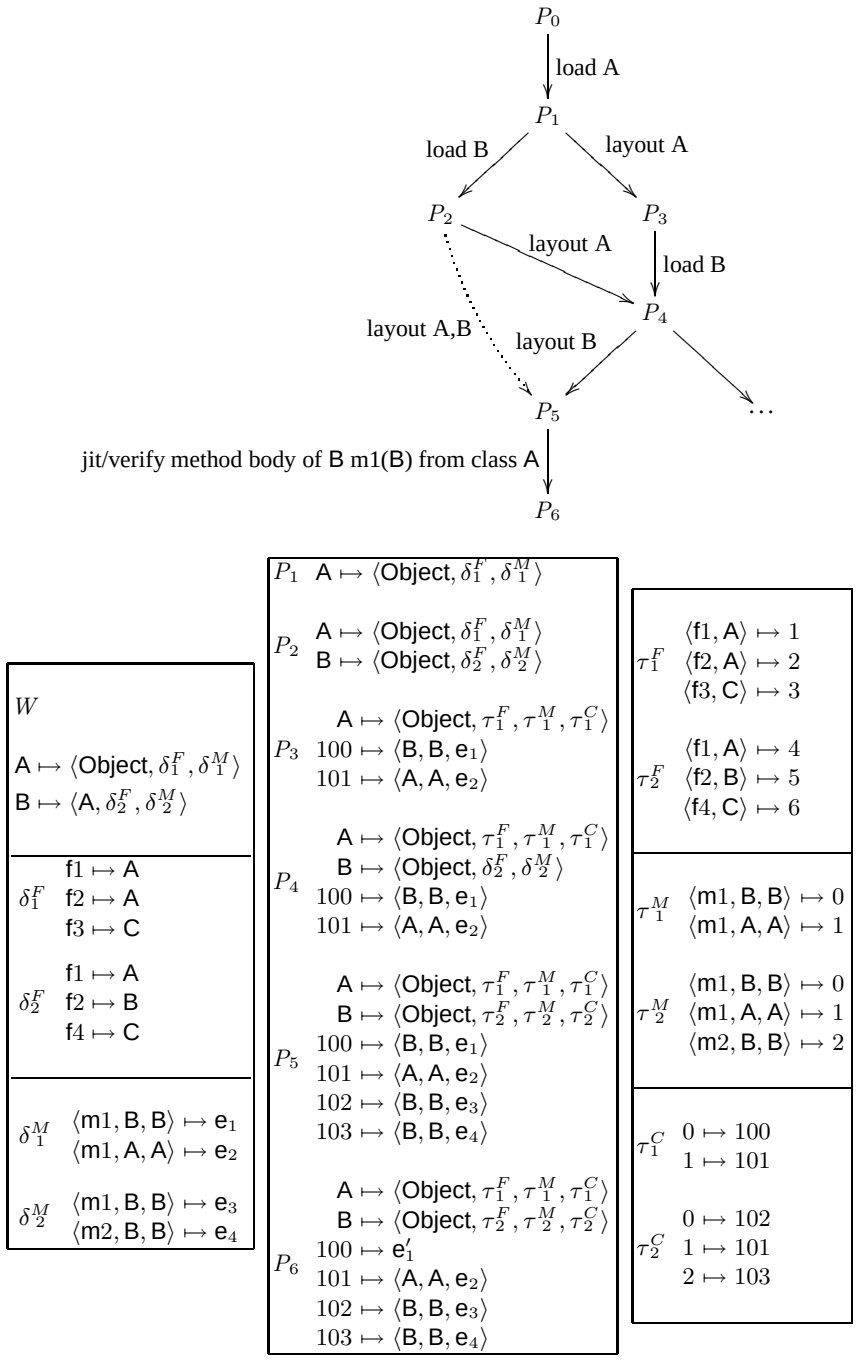

Figure 10. Example demonstrating table layout

Figure 10 shows a global context W which describes these classes. Also, it shows a possible sequence of programs involved in execution, and the contents of these programs.

We start with a program P_0, where A and B have not yet been read in – obviously, P_0 contains Object, but we do not show this for the sake of brevity.

Then, we load A, and obtain P_1, for which $W \vdash P_1 \leq P_0$ holds.

From P_1, by loading B, we obtain $P2$, whereas, if we lay out A, we obtain P_3. Therefore, we have $W \vdash P_2 \leq P_1$ and $W \vdash P_3 \leq P_1$ but $W \nvdash P_2 \leq P_3$ and $W \nvdash P_3 \leq P_2$.

We then have $W \vdash P_4 \leq P_3$ through loading of B, and $W \vdash P_5 \leq P_4$ through laying out of class B.

Finally, from P_5 we obtain P_6 jit/verifying the method body m1 of class A located at address 100. Thus, we have that $P_5, \mathsf{e} \leadsto_{W,\texttt{this}\mapsto\mathsf{A},\mathsf{y}\mapsto\mathsf{B}} P'', \mathsf{e}', t$ and $P'', t, \mathsf{A} \leadsto_W P_6$. So, we also have that $W \vdash P_6 \leq P_5$.

C The example revisited

In this section we revisit the example from section 2, and we show how the non-deterministic semantics of our paper would give the Java and the C# behaviour.

Let us first assume layout tables τ^F_1, τ^M_1, and τ^C_1, with $\tau^F_1 = \epsilon$, and with τ^M_1 (eat,void,Penne)=κ_1, and with τ^M_1 (chew,void,Pasta)=κ_2, and with $\tau^C_1(\kappa_1)=\iota_1$, an with $\tau^C_1(\kappa_2)=\iota_2$.

Verification is "lazier" in C# We consider the situation where the class Meal has been loaded, but not laid out yet, nor any of its methods have been jit-compiled, or verified. That is, we have a program P_0, where P_0 (Meal)=$\langle \texttt{Object}, \delta^F_S, \delta^M_S \rangle$. We have the configuration $P_0, H, \texttt{new}$ Meal. Here are some possible rewrites:

0	$\leadsto P_0, H, \texttt{lnkExc}$	a link related exception thrown	
1	$\leadsto P_1, H, \texttt{new}$ Meal	Meal has been laid out in P_1: $i.e.,\ P_1(\text{Meal}) = \langle \texttt{Object}, \tau^F_1, \tau^M_1, \tau^C_1 \rangle$ $P_0 = \text{Meal} = P_1$	
2	$\leadsto P_2, H, \texttt{new}$ Meal	Meal has been laid out in P_2: $i.e.,\ P_2(\text{Meal}) = P_1(\text{Meal})$ eat has been jit-compiled/verified, $i.e.,\ P_2(\iota_2) = \texttt{this}....,$ $P_2 = \iota_1 = P_1$	if Sub
3	$\leadsto P_3, H, \texttt{new}$ Meal	Meal has been laid out in P_3, as in P_2... eat has been jit-compiled/verified, as in P_2... chew has been jit-compiled/verified $i.e.,\ P_3(\iota_2) = \textbf{if}...\textbf{else}...$ $P_3 = \iota_2 = P_2$	if Sub

Java execution is described by the 0th and the 3rd execution. Namely, the 0th execution would stand for the case where $\neg$ Sub, and thus the class Meal could not be verified, whereas the 3rd execution stands for the case where the Sub holds, and class Meal was verified. On the other hand, C# only requires for the class to be laid out (no method has yet been called) and this corresponds to the 1st execution.

Of course, the model allows for many more possible rewrites.

So far, we have not seen the situation where C# would require verification, and this is shown in the next example:

Consider execution of the method call m.eat (y). This could rewrite to something like $\iota_{10}[\kappa_1](\iota_{20})$, (remember that eat is at the offset κ_1) assuming that m was at address ι_{10}, and the argument at address ι_{20} in heap H_1. Also, assume that chew has not been jit-compiled yet, *i.e.,* we have P_1, with
$P_1(\iota_2) = \langle$void, Penne, this.chew [Meal,void,Pasta](p)$\rangle$.

Then, we have several possibilities for the configuration $P_1, H_1, \iota_{10}[\kappa_1](\iota_{20})$:

4	$\rightsquigarrow P_0, H,$ `lnkExc`	a link related exception thrown
5	$\rightsquigarrow P_4, H_1, \iota_{10}[\kappa_1](\iota_{20})$	eat has been jit-compiled , in P_4 if Sub $i.e.,\ P_4(\iota_2) =$ this[κ_2](p)

That is, in P_4 the method body for eat has been verified (using Sub), and offset calculation has replaced the the symbolic reference $.chew$[Meal, void, Pasta] with the offset κ_2.

Offset calculation is "lazier" in Java Let us now consider the situation where this.chew (p) is called, which would be something like $\iota_{10}[\kappa_2](\iota_{20})$. We may distinguish two possibilities for the program: In P_3 the method chew has been verified, but symbolic offsets were not replaced, *i.e.,* P_3 (ι_2)= **if**...**else** ... p.cal[Pasta,int].. – this corresponds to Java. In a program P_4, method chew is still "raw", *i.e.,* we have
P_4 (ι_2)= $\langle$void, Pasta, **if**...**else** ... p.cal [int]..$\rangle$.

So, we have the following possibilities:

$P_3, H_1, \iota_{10}[\kappa_2](\iota_{20}) \rightsquigarrow^* P_3, H_1, \iota_{10}$.cal [Pasta,int]

6		$\rightsquigarrow P_3, H_1,$ `lnkExc`	
7		$\rightsquigarrow P_3, H_1, \iota_{10}[\kappa_3]$	if Fld

$P_4, H_1, \iota_{10}[\kappa_2](\iota_{20})$

8	$\rightsquigarrow$	$P_4, H_1,$ `lnkExc`	
9	$\rightsquigarrow$	$P_5, H_1, \iota_{10}[\kappa_2](\iota_{20}) \rightsquigarrow^* P_5, H_1, \iota_{10}[\kappa_3]$	if Fld

Java executions are described by the 6th and 7th possibility: If $\neg$ Fld, then a field not found exception will be thrown; otherwise, the symbolic reference will be replaced by the offset, say κ_3.

C# executions are described by the 8th and 9th possibility: If $\neg$ Fld, then jit-compilation of the method chew will throw a field not found exception; otherwise, it will be successful, and will replace the "raw" method by its jit-compiled version in P_5, so that $P_4 = \iota_2 = P_5$, and $\ldots \vdash P_5 \leq P_4$, and P_5 (ι_2)== **if**...**else** ... p.[κ_3].., [11]

D The lemmas in further detail

In this appendix we give some further details on the lemmas mentioned in section 5.

Lemma 1

- $P, e \rightsquigarrow_{W,E} P', e', t \implies W \vdash P' \leq P$.
- $P, H, e \rightsquigarrow_W P', H', e' \implies W \vdash P' \leq P$.

Proof follows from the definition of the rewrite relationships ... $\rightsquigarrow_{W,E}$ ... and of ... $\rightsquigarrow_W$ ....

Lemma 2 *If $W \vdash P' \leq P$, then*

- $P \vdash t_1 \leq t_2 \implies P' \vdash t_1 \leq t_2$.
- $P \vdash H \implies P' \vdash H$.
- $P, H \vdash e : t \implies P', H \vdash e : t$.
- $P, e \rightsquigarrow_{W,E} P, e', t \implies P', e \rightsquigarrow_{W,E} P', e', t$ [12].
- $\vdash P \implies \vdash P'$.

[11] where κ_3 would be the offset of field int cal in class Pasta

[12] Notice, that the premise $P, e \rightsquigarrow_{W,E} P, e', t$ does not allow extension of the program P. Although the lemma could be generalized to allow for extensions, the current restricted form suffices for the proof of soundness.

Proof by structural induction over the judgments of each of the assertions, *i.e.,* over $P \vdash t_1 \leq t_2$, and $P, H \vdash e : t$ *etc*. The last assertion requires structural induction over the derivation of $W \vdash P' \leq P$, and then a proof that all four requirements of $\vdash P'$ are satisfied.

Lemma 3 *If* $P \vdash H$, *and* $P, e \leadsto_{\emptyset, \{\texttt{this} \mapsto c, \texttt{y} \mapsto t_p\}} P, e', t$, *and* $P, H \vdash \iota \triangleleft c$ *and* $P, H \vdash \iota' \triangleleft t_p$, *then* $P, H \vdash e'[\iota/\texttt{this}, \iota'/\texttt{y}] : t$

Lemma 4 *If* $P \vdash H$, *and* $\vdash P$, *and* $P, H \vdash e : t$, *and* $P, H, e \leadsto_W P', H', e'$, *then*

- $H(\iota) = c \implies H'(\iota) = c$.
- $H'(\iota) = c \implies H(\iota) = c$ *or* ι *free in* H.
- $P \vdash H \implies P' \vdash H'$.
- $P, H \vdash e'' : t'' \implies P', H' \vdash e'' : t''$.

Proof by case analysis over the expression e; the last statement is proven by structural induction over the typing of e''. The requirements $\vdash P$ and $P, H \vdash e : t$ are needed in order to guarantee that that memory is in "appropriate" ways only.

The proof of theorem 1 goes by structural induction over the typing of e.

The proof of theorem 2 requires most of the work: Namely, we first prove that for any execution,

$$P, H, e \leadsto^*_W e', H', \nu,$$

there exists an "equivalent" execution, consisting of three parts, where the first extends the program and leaves expression unaffected so that all necessary information for evaluation is performed, the second evaluates the expression and leaves the program unaffected, and the last extends the program with any bits that were not really necessary for the execution, *i.e.,*

$$P, H, e \leadsto^*_W P'', H, e \leadsto^*_W P'', H', e' \leadsto^*_W P', H', e',$$

where the middle part only involves evaluation or offset calculation steps. We can also show that the contents of the program P'' is a function of the expression e to be executed, and the global context W . We can then show that executions, which only involve evaluation or offset calculation steps, are deterministic up to renaming of addresses.

Theorem 3 (Monotonicity of Execution with respect to environments) *For any* e, P, P', H, H', *and* $\nu \in \mathbb{N} \cup \{\texttt{nllPExc}\}$, *if* :

$$P, H, e \leadsto^*_W P', H', \nu, \quad \textit{and}$$
$$W \mid_{def(P')} = W' \mid_{def(P')},$$

then

$$P, H, e \leadsto^*_{W'} P', H', \nu.$$

We could probably weaken the requirement $W \mid_{def(P')} = W' \mid_{def(P')}$, to say that only the parts required by the expression need to be identical.