

Flexible, source level dynamic linking and re-linking ^{*}

– Work in Progress –

Sophia Drossopoulou, Susan Eisenbach,
Imperial College London

Abstract

We give a formal semantics for dynamic linking and re-linking of code. The semantics is at source language level, and allows linking at a finer grain than current Java or C# implementations: Besides supporting the loading and verification of classes interleaved with program execution, it also allows type-safe removal and replacement of classes, fields and methods. Such extended features support unanticipated software evolution.

1 Introduction

Dynamic linking was introduced into programming languages to support implicit linking of the newest version of libraries. Partially linked code links further code on the fly, as needed, and thus, most recent updates are automatically available to end-users, without the requirement of further compilations or re-linking.

The remit has recently been extended, to support requirements from unanticipated software evolution, whereby code may need to be *updated during* program execution. To address such requirements, new features of the JVM allow the replacement of a class by a class of the same signature, as a “fix-and-continue” feature [5], while in [11] dynamic software updates support type safe dynamic reloading of code whose type may have changed, while the system is running. Also, [2] suggest dynamic linking of modules, with the support of multiple versions of the same module. Finally, [1] suggest a calculus for dynamic linking independently of the programming language and environment.

We present the semantics for a high level language similar to Java or C# with extended dynamic linking. In contrast to earlier formalizations of dynamic linking [12, 8], our model is purely at source language level. Namely, although dynamic linking is usually described at a “low level” [10], in terms of generated *.class files, its effects are visible to the source language programmer, as shown *e.g.*, through a sequence of examples in [6, 9]. This is why, in our view, an explanation at source language level is very important. In order to simplify the exposition, we do not deal with overloading or field shadowing.

In more detail, the model we suggest allows dynamic linking at a very fine grain:

DL1 Method signatures may be loaded before verification of their corresponding bodies,

DL2 Method bodies need be verified only before execution¹,

DL3 Classes may be loaded before their fields/methods,

DL4 Classes, methods and fields may be loaded “piecemeal”,

DL5 Classes, methods, and fields may be removed, if not needed,

DL6 Method bodies may be replaced.

^{*}This work partly supported by DART, EU project IST-2001-33477.

¹This is similar to C# dynamic linking.

$\frac{\chi \vdash \pi' \leq \pi \quad \mathbf{e}, \pi', \sigma, \chi \rightsquigarrow r, \pi'', \chi'}{\mathbf{e}, \pi, \sigma, \chi \rightsquigarrow r, \pi'', \chi'}$	$\frac{}{\mathbf{e}, \pi, \sigma, \chi \rightsquigarrow \text{LnkExc}, \pi, \chi}$
$\frac{\chi \ll \chi'' @ \iota \quad \iota \text{ not reachable from } \sigma \text{ in } \chi \quad \mathbf{e}, \pi, \sigma, \chi'' \rightsquigarrow r, \pi', \chi'}{\mathbf{e}, \pi, \sigma, \chi \rightsquigarrow r, \pi', \chi'}$	$\frac{}{\mathbf{y}, \pi, \sigma, \chi \rightsquigarrow \sigma(\mathbf{y}), \pi, \chi}$ $\frac{}{\mathbf{this}, \pi, \sigma, \chi \rightsquigarrow \sigma(\mathbf{this}), \pi, \chi}$ $\frac{}{r, \pi, \sigma, \chi \rightsquigarrow r, \pi, \chi}$
$\frac{\mathbf{e}, \pi, \sigma, \chi \rightsquigarrow \iota, \pi', \chi' \quad \mathbf{e}', \pi', \sigma, \chi' \rightsquigarrow \kappa, \pi'', \chi'' \quad \chi''(\iota)(\mathbf{f}) = \kappa', \quad \chi''' = \chi''[\iota \mapsto \chi''(\iota)[\mathbf{f} \mapsto \kappa]]}{\mathbf{e}.f := \mathbf{e}', \pi, \sigma, \chi \rightsquigarrow \kappa, \pi'', \chi'''}$	$\frac{}{\mathbf{e}, \pi, \sigma, \chi \rightsquigarrow \iota, \pi', \chi'}$ $\frac{}{\mathbf{e}.f, \pi, \sigma, \chi \rightsquigarrow \chi'(\iota)(\mathbf{f}), \pi', \chi'}$
$\frac{\mathbf{e}, \pi, \sigma, \chi \rightsquigarrow \iota, \pi_1, \chi_1 \quad \mathbf{e}_i, \pi_i, \sigma, \chi_i \rightsquigarrow \kappa_i, \pi_{i+1}, \chi_{i+1} \quad 1 \leq i \leq n \quad \chi_{n+1}(\iota) = (\mathbf{c}, \dots) \quad \mathcal{MB}(\pi_{n+1}, \mathbf{c}, \mathbf{m}) = \mathbf{e}' \quad \phi = \mathbf{this} \mapsto \iota, \mathbf{y}_1 \mapsto \kappa_1, \dots, \mathbf{y}_n \mapsto \kappa_n \quad \mathbf{e}', \pi_{n+1}, \sigma \cdot \phi, \chi_{n+1} \rightsquigarrow \kappa, \pi', \chi'}{\mathbf{e}.m(\mathbf{e}_1, \dots, \mathbf{e}_n), \pi, \sigma, \chi \rightsquigarrow \kappa, \pi', \chi'}$	$\frac{}{\mathcal{F}s(\pi, \mathbf{c}) = \{f_1, \dots, f_r\}}$ $\frac{}{\iota \text{ is fresh in } \chi}$ $\frac{}{\chi' = \chi[\iota \mapsto (\mathbf{c}, f_1 \mapsto \mathbf{0} \dots f_r \mapsto \mathbf{0})]}$ $\frac{}{\mathbf{new} \mathbf{c}, \pi, \sigma, \chi \rightsquigarrow \iota, \pi, \chi'}$

Figure 1: Execution.

DL1-DL4 represent a refinement of dynamic linking as currently implemented. **DL5-DL6** are even more “dynamic”, and come at a higher implementation cost, as they require re-checking the entire program. **DL5-DL6** have been suggested in order to support on-the-fly code replacement, and thus, unanticipated software evolution.

Furthermore, the fact that we model dynamic linking at source language level, *imposes* the requirement for **DL1** and **DL2**: Namely, when verifying a method body, \mathbf{e} , we need to find the signatures of any methods called within \mathbf{e} ; for this, we load the method signatures for all called methods. But, we do *not* verify their corresponding method bodies, because if we did, we would obtain total, eager, verification *before* execution.²

Notation All mappings in our model are partial, and ϵ represents undefined. A function f' *extends* another function f at a , formally $f' \ll f @ a$, if f' is defined on one more element than f , *i.e.*, iff $f' \downarrow_{\text{dom}(f)} = f$ and $\text{dom}(f') = \text{dom}(f) \uplus \{a\}$.

2 Syntax, runtime model, and execution

Expressions in our minimal language are method call, field assignment, field access, object creation, the receiver **this**, a parameter identifier \mathbf{y}_i ($i \in \mathbb{N}$) or null indicated by $\mathbf{0}$. Class names are indicated by \mathbf{c} , \mathbf{c}' *etc.*, method names are \mathbf{m} , \mathbf{m}' , *etc.*, and field identifiers are \mathbf{f} , \mathbf{f}' .

$$\mathbf{e} \in \text{Exp} ::= \mathbf{e}.m(\mathbf{e}^*) \mid \mathbf{e}.f := \mathbf{e} \mid \mathbf{e}.f \mid \mathbf{new} \mathbf{c} \mid \mathbf{this} \mid \mathbf{y}_i \mid \mathbf{0}$$

²This problem does not appear in Java and C# dynamic linking, because bytecode is “enriched” with the method signature for called methods; therefore, verification takes place under the assumption that the classes contain the corresponding methods, and if they do not, then an offset calculation error is thrown. In that sense, verification in our model is more eager than in Java and C#.

We represent dynamic linking through "extensions" to the underlying program π ³. Therefore, the operational semantics rewrites tuples of expressions, programs, stacks and heaps into tuples of results, programs and heaps.

$$\begin{array}{ll}
\sim & : \quad \text{Exp} \times \text{Prg} \times \text{Stack} \times \text{Heap} \longrightarrow \mathcal{R} \times \text{Prg} \times \text{Heap} \\
\sigma \in \text{Stack} & = \quad (\text{StackFrm})^* \\
\phi \in \text{StackFrm} & = \quad (\{\mathbf{this}\} \uplus \text{Param}) \longrightarrow \mathbb{N} \\
\chi \in \text{Heap} & = \quad \mathbb{N}^+ \longrightarrow \text{Obj} \\
o \in \text{Obj} & = \quad \text{ClsId} \times \text{FieldMap} \\
fm \in \text{FieldMap} & = \quad \text{FldId} \longrightarrow \mathbb{N} \\
\iota \in \mathbb{N}^+ & \\
\kappa \in \mathbb{N} & \\
r \in \mathcal{R} & = \quad \mathbb{N} \cup \{\text{nullPtrExc}, \text{StuckExc}, \text{LnkExc}\}
\end{array}$$

Stacks, σ , are sequences of stack frames, ϕ , which map **this** to an address, and the parameters y_1, \dots, y_n to addresses or null. We use the shorthand $\sigma(z)$, for $\phi_n(z)$, when ϕ_n is the top of σ (i.e., $\sigma = \phi_1 \dots \phi_n$). Heaps, χ , map addresses to objects. The notation $o = (c, fm)$, stands for an object of class c , with fields described by the mapping fm . For such an object, the field lookup $o(f)$ is a shorthand for $o \downarrow_2(f)$, and field update $o[f \mapsto \kappa]$ is a shorthand for $(o \downarrow_1, o \downarrow_2[f \mapsto \kappa])$. Results, r , are either addresses, or null (**0**), or exceptions. Addresses, ι , are positive numbers.

Figure 1 contains the operational semantics of the language. Rule *lnk* replaces a program by an extending program *at any time* during execution, while *exc* allows link related exceptions to be thrown *at any time* during execution. We thus have a highly non-deterministic model, similarly to [8]. Rule *gbgCllt* allows for garbage collection – the appendix contains the definition for reachable addresses. Garbage collection is important in our study, as classes may be removed when there are no objects of that class. Garbage collection is defined in a style similar to that from [3].

The remaining rules are standard for the operational semantics of such a small object oriented language, and are similar to previous work [7]: Rule *var* describes the lookup of parameters, the receiver, or a value. Rules *fld* and *fldAss* describe field lookup and field update. Rule *call* describes method call: a new frame ϕ is pushed onto the stack, which is discarded after execution of the method body e' . Rule *new* describes object creation, where all fields of its class are initialized with **0**. In the appendix we give rules for null pointer exceptions and for stuck execution.

The different kinds of exception and the guarantees of Soundness Our calculus allows three kinds of exception: *Null pointer* exceptions, `nullPtrExc`, are thrown when a null-pointer is de-referenced to access a field or a method; they are propagated to the context as well. *Stuck* exceptions, `stuckExc`, are thrown when an object is de-referenced to access a non-existing field, or when a non-existing method is called; they are also propagated to the context. *Link* exceptions, `LnkExc`, may be thrown at any point during execution, and thus they need not be propagated to the context. They represent link-related exceptions, i.e., verification errors, class-not-found errors, class-circularity-errors, load errors *etc.*⁴

Soundness, as stated in section 4, guarantees that execution will not get stuck, i.e., that that execution will not attempt to access a non-existing field or method. The case of an "absent method" is slightly subtle: If the dynamic class of the receiver of a method call of m contains a method signature for that m , but contains no method body, then that represents the case where the method body could not be verified, and therefore it is a link related error. If, on the other hand, the dynamic class of the receiver contains no method signature for m , then, this corresponds to the case where the class has no such method, and therefore it is a `stuckExc`. More details are in the appendix.

³Note, that "program extensions" stands for all the dynamic linking steps from **DL1- DL6**, i.e., even such that remove classes or methods

⁴Of course, we do not model the `FldAbsent` and `MethodAbsent` errors of Java and C#, since these are related to the "enriched" bytecode, and offset calculations.

<p style="text-align: center;"><i>refl</i></p> $\frac{}{\chi \vdash \pi \leq \pi}$	<p style="text-align: center;"><i>trans</i></p> $\frac{\chi \vdash \pi' \leq \pi'' \quad \chi \vdash \pi'' \leq \pi}{\chi \vdash \pi' \leq \pi}$
<p style="text-align: center;"><i>ldClass</i></p> $\frac{\pi'_s \ll \pi_s @ c \quad \pi_s(\pi'_s(c)) \neq \epsilon}{\chi \vdash (\pi'_s, \pi_f, \pi_{ms}, \pi_{mb}) \leq (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}$	<p style="text-align: center;"><i>ldFld</i></p> $\frac{\pi'_f(c) \ll \pi_f(c) @ f \quad \pi_s(c) \neq \epsilon \quad \chi(l) = (c', _) \implies \pi \not\vdash c' \leq c \quad \pi \vdash c \bowtie c' \implies \pi'_f(c)(f) = \epsilon}{\chi \vdash (\pi_s, \pi'_f, \pi_{ms}, \pi_{mb}) \leq (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}$
<p style="text-align: center;"><i>ldMthSig</i></p> $\frac{\pi'_{ms} \ll \pi_{ms} @ (c, m) \quad \pi_s(\pi'_s(c)) \neq \epsilon \quad \pi \vdash c' \leq c \implies \pi \vdash \pi_{ms}(c', m) \leq \pi'_{ms}(c, m) \quad \pi \vdash c \leq c' \implies \pi \vdash \pi'_{ms}(c, m) \leq \pi_{ms}(c', m)}{\chi \vdash (\pi_s, \pi_f, \pi'_{ms}, \pi_{mb}) \leq (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}$	<p style="text-align: center;"><i>verMthBdy</i></p> $\frac{\pi'_{mb} \ll \pi_{mb} @ (c, m) \quad \pi_{ms}(c, m) = c_1, \dots, c_{n+1}, \quad n \geq 0 \quad \pi'_{mb}(c, m) = e \quad \pi, (\mathbf{this} \mapsto c, \mathbf{y}_1 \mapsto c_1, \dots, \mathbf{y}_n \mapsto c_n) \vdash e : c' \quad \pi \vdash c' \leq c_{n+1}}{\chi \vdash (\pi_s, \pi_f, \pi_{ms}, \pi'_{mb}) \leq (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}$
<p style="text-align: center;"><i>rmFld</i></p> $\frac{\pi_f(c) \ll \pi'_f(c) @ f \quad \pi_{ms}(c', m) = c_1, \dots, c_{n+1}, \quad \pi_{mb}(c', m) = e \implies \pi', (\mathbf{this} \mapsto c', \mathbf{y}_1 \mapsto c_1, \dots, \mathbf{y}_n \mapsto c_n) \vdash e : c'' \quad \pi \vdash c'' \leq c_{n+1}}{\chi \vdash (\pi_s, \pi'_f, \pi_{ms}, \pi_{mb}) \leq (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}$	<p style="text-align: center;"><i>rmMthSig</i></p> $\frac{\pi_{ms} \ll \pi'_{ms} @ (c, m) \quad \pi_{mb}(c, m) = e \quad \pi_{ms}(c', m') = c_1, \dots, c_{n+1}, \quad \pi_{mb}(c', m') = e \implies \pi', (\mathbf{this} \mapsto c', \mathbf{y}_1 \mapsto c_1, \dots, \mathbf{y}_n \mapsto c_n) \vdash e : c'' \quad \pi \vdash c'' \leq c_{n+1}}{\chi \vdash (\pi_s, \pi_f, \pi'_{ms}, \pi_{mb}) \leq (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}$
<p style="text-align: center;"><i>rmMthBdy</i></p> $\frac{\pi_{mb} \ll \pi'_{mb} @ (c, m)}{\chi \vdash (\pi_s, \pi_f, \pi_{ms}, \pi_{mb}) \leq (\pi_s, \pi_f, \pi_{ms}, \pi'_{mb})}$	
<p style="text-align: center;"><i>rplMthBdy</i></p> $\frac{\pi'_{mb} = \pi_{mb}[(c, m) \mapsto e] \quad \pi_{ms}(c, m) = c_1, \dots, c_{n+1}, \quad n \geq 0 \quad \pi, (\mathbf{this} \mapsto c, \mathbf{y}_1 \mapsto c_1, \dots, \mathbf{y}_n \mapsto c_n) \vdash e : c' \quad \pi \vdash c' \leq c_{n+1}}{\chi \vdash (\pi_s, \pi_f, \pi_{ms}, \pi'_{mb}) \leq (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}$	<p style="text-align: center;"><i>rmClass</i></p> $\frac{\pi_s \ll \pi'_s @ c \quad \chi(l) = (c', _) \implies \pi \not\vdash c \leq c' \quad \pi_s(c') \neq c \text{ for all } c' \quad \pi_{ms}(c, m) = \pi_f(c, f) = \epsilon \text{ for all } m, f \quad \pi_{ms}(c', m) = c_1, \dots, c_{n+1}, \quad \pi_{mb}(c, m) = e \implies \pi', (\mathbf{this} \mapsto c', \mathbf{y}_1 \mapsto c_1, \dots, \mathbf{y}_n \mapsto c_n) \vdash e : c'' \quad \pi \vdash c'' \leq c_{n+1}}{\chi \vdash (\pi'_s, \pi_f, \pi_{ms}, \pi_{mb}) \leq (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}$

Figure 2: Program extension

3 Programs and program extension

A program maps a class to its superclass, a class and a field identifier to its type, and class and method identifier to its signature, and method body:

$$\begin{array}{lll}
\pi \in \text{Prg} = & \text{ClsId} \longrightarrow \text{ClsId} & \text{maps class to its superclass} \\
& \times \text{ClsId} \longrightarrow \text{FldId} \longrightarrow \text{ClsId} & \text{maps class and field name to type} \\
& \times \text{ClsId} \times \text{MthId} \longrightarrow (\text{ClsId}^+) & \text{maps class and meth. name to type} \\
& \times \text{ClsId} \times \text{MthId} \longrightarrow \text{Exp} & \text{maps class and meth. name to body}
\end{array}$$

A program is therefore a tuple; we use the notation $(\pi_s, \pi_f, \pi_{ms}, \pi_{mb})$ to indicate the four components, and we use the naming convention that π_s , π_f , π_{ms} , and π_{mb} stand for mappings of the appropriate signatures.

We now define field and method look up functions:

$$\begin{array}{ll}
\mathcal{F} & : \text{Prg} \times \text{ClsId} \times \text{FldId} \longrightarrow \text{ClsId} \\
\mathcal{F}_s & : \text{Prg} \times \text{ClsId} \times \text{FldId} \longrightarrow \mathcal{P}(\text{FldId}) \\
\mathcal{MS} & : \text{Prg} \times \text{ClsId} \times \text{MthId} \longrightarrow \text{ClsId}^* \\
\mathcal{MB} & : \text{Prg} \times \text{ClsId} \times \text{MthId} \longrightarrow \text{Exp}
\end{array}$$

where, for $\pi = (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})$, these functions are defined as follows:

$$\begin{array}{ll}
\mathcal{F}(\pi, c, f) & = \begin{cases} \pi_f(c)(f) & \text{if } \pi_f(c)(f) \neq \epsilon \\ \mathcal{F}(\pi, c', f) & \text{if } \pi_f(c)(f) = \epsilon, \text{ and } c \neq \pi_s(c) = c' \\ \epsilon & \text{otherwise.} \end{cases} \\
\mathcal{F}_s(\pi, c) & = \begin{cases} \{ f \mid \mathcal{F}(\pi, c, f) \neq \epsilon \} & \text{if } \pi_s(c) \neq \epsilon \\ \epsilon & \text{otherwise} \end{cases} \\
\mathcal{MS}(\pi, c, m) & = \begin{cases} \pi_{ms}(c, m) & \text{if } \pi_{ms}(c, m) \neq \epsilon \\ \mathcal{MS}(\pi, c', m) & \text{if } \pi_{ms}(c, m) = \epsilon, \text{ and } c \neq \pi_s(c) = c' \\ \epsilon & \text{otherwise.} \end{cases} \\
\mathcal{MB}(\pi, c, m) & = \begin{cases} \pi_{mb}(c, m) & \text{if } \pi_{mb}(c, m) \neq \epsilon \\ \mathcal{MB}(\pi, c', m) & \text{if } \pi_{mb}(c, m) = \epsilon, \text{ and } c \neq \pi_s(c) = c' \\ \epsilon & \text{otherwise.} \end{cases}
\end{array}$$

The judgments $\pi_s \vdash c' \leq c$ and $\pi \vdash c' \leq c$ mean that c' is a subclass of c , while the judgments $\pi_s \vdash c' \not\leq c$ and $\pi \vdash c' \not\leq c$ mean that the two classes are distinct with one a subclass of the other. The four judgments are defined in the appendix.

The rules in figure 2 define when a program extends another program:

refl and *trans* say that the extension relationship is reflexive and transitive.

ldClass describes loading of a class c , provided that the superclass of the loaded class has a superclass defined in the old program, *i.e.*, for $\pi'_s(c) = c'$, we need $\pi_s(c') \neq \epsilon$. Note, that the class is loaded as an "empty shell" and its fields and methods may be loaded later, "piecemeal", through rules *ldFld* and *ldMthSig*.

ldFld describes loading the type of a field f for a class c . It requires that no objects of that class or its subclasses exist on the heap⁵, and that the field f has not been defined in any of the subclasses or superclasses of c ⁶.

⁵If we dropped this requirement, then at a later stage we could verify a method body that used the new field, although we might have objects in the heap, which did not contain that field

⁶In dynamic linking as in current JVMs and C# implementations, the situation is simpler. The requirement that there are no objects of subclass of c is implicitly guaranteed because classes are fully loaded before the first object is created, and the requirement that there is no other field f in any of the sub- or superclasses of c need not be imposed because field resolution takes the class defining the field into account.

ldMthSig describes loading the signature of a method m in a class c . It requires any signature for m in a subclass of c to be a "subsignature" of the newly loaded signature, and any signature for m in a superclass of c to be a "supersignature" of the newly loaded signature,⁷ The "subsignature judgment $\pi \vdash c'_1 \dots c'_n \leq c_1 \dots c_n$ requires the argument types of the first to be supertypes of the corresponding argument types of the second, and the return type of the first to be a subtype of the return type of the second - or the second signature to be empty; it is defined in the appendix. Notice that it does *not* require the classes appearing in the signatures of methods or the types of fields to belong to classes that have already been loaded.

verMthBdy describes the verification of a method body. It requires the signature of the method to have already been loaded, and the body of the method to verify, *i.e.*, to be type correct, assuming that the parameters y_1, \dots, y_n have types c_1, \dots, c_n , as prescribed by the signature, and return a result of a subtype of c_{n+1} , the return type of the signature. In figure 3 we define verification.

rmFld allows the removal of a field f from a class c , provided that no method body verification depended on the existence of that field, *i.e.*, all method bodies verify in π' , after the removal of the field. Note, that in contrast to rule *ldFld*, we do *not* require that there should be no objects of class c on the heap, since this does not affect type soundness⁸.

rmMthSig allows the removal of method signature for m for class c , provided that no method body verification depended on the existence of that method signature, *i.e.*, all method bodies verify in π' .

rmMthBdy allows the removal of method body for m for class c . Note, that this allows the removal of methods which are currently active on the stack. This is not problematic in our setting with large step operational semantics, which, in some sense, "copies" method bodies upon method call.

rplMthBdy allows replacing the method body for m in class c by expression e , provided that e satisfies the method signature.

rmCls allows the removal of class c , provided that this class is not the superclass of any other class c' , that there exist no objects of that class on the heap, that c has no fields or methods defined, and that c was not needed in the verification of any method.

Rules *ldCls*, *ldFld*, *ldMthSig* and *ldMthBody* correspond to features **DL1-DL4**, and thus, they reflect dynamic linking as in current Java and C# implementations, albeit with finer steps, thus requiring more detailed checks than current implementations. Rule *rmCls* corresponds to current JVM hotswapping.

Rules *rmFld*, *rmMthSig*, *rmMthBody*, and *rmCls* correspond to **DL5-DL6**. Because they allow the removal of entities, they make heavy requirements on checking - in effect they require re-checking the complete current program, and, in some cases, they also require scanning the complete heap. The expense of these checks is justified by the fact that runtime removal and replacement are rarely applied, but when they are applied, they are necessitated by the demand to upgrade safely a, possibly safety-critical, system on the fly.

Rules *rmFld*, *rmMthSig*, *rmMthBody*, and *rmCls* combined with rules *ldCls*, *ldFld*, *ldMthSig* and *ldMthBody* have the overall effect of replacement of entities. However, although any replacement can be represented by a sequence of such primitive removal and loading operations, there exist simple replacements, which require a long sequence of such

⁷This requirement is unnecessary in Java and C# implementations, because of dynamic linking is at byte-code level, in which method calls are annotated with the corresponding signature.

⁸This has been pointed out by the second anonymous referee!

<p style="text-align: center;"><i>verBasic</i></p> <hr style="width: 100%;"/> $\begin{array}{l} \pi, \gamma \vdash \mathbf{this} : \gamma(\mathbf{this}) \\ \pi, \gamma \vdash \mathbf{y} : \gamma(\mathbf{y}) \\ \pi, \gamma \vdash \mathbf{new} \ c : c \\ \pi, \gamma \vdash \mathbf{0} : c \end{array}$	<p style="text-align: center;"><i>verFldAss</i></p> <hr style="width: 100%;"/> $\begin{array}{l} \pi, \gamma \vdash \mathbf{e} : c \\ \mathcal{F}(\pi, c, \mathbf{f}) = c' \\ \pi, \gamma \vdash \mathbf{e}' : c'' \\ \pi \vdash c'' \leq c' \end{array}$ <hr style="width: 100%;"/> $\pi, \gamma \vdash \mathbf{e.f} := \mathbf{e}' : c''$
<p style="text-align: center;"><i>verFld</i></p> <hr style="width: 100%;"/> $\begin{array}{l} \pi, \gamma \vdash \mathbf{e} : c \\ \mathcal{F}(\pi, c, \mathbf{f}) = c' \\ \pi, \gamma \vdash \mathbf{e.f} : c' \end{array}$	<p style="text-align: center;"><i>verMthCl</i></p> <hr style="width: 100%;"/> $\begin{array}{l} \pi, \gamma \vdash \mathbf{e}_i : c'_i \quad 0 \leq i \leq n \\ \mathcal{MS}(\pi, c_0, \mathbf{m}) = c_1, \dots, c_{n+1} \quad n \geq 0 \\ \pi \vdash c'_i \leq c_i \quad 1 \leq i \leq n \end{array}$ <hr style="width: 100%;"/> $\pi, \gamma \vdash \mathbf{e}_0.\mathbf{m}(\mathbf{e}_1, \dots, \mathbf{e}_n) : c_{n+1}$

Figure 3: Verification.

$$\begin{array}{l} \pi = (\pi_s, \pi_f, \pi_{ms}, \pi_{mb}) \\ \pi_s(c) = c' \neq \epsilon \implies \pi_s(c') \neq \epsilon \\ \pi_s(\mathbf{Object}) = \mathbf{Object} \\ \pi \vdash c \leq c' \text{ and } \pi \vdash c' \leq c \implies c = c' \\ \pi_f(c, \mathbf{f}) \neq \epsilon \implies \pi_s(c) \neq \epsilon \\ \pi_{ms}(c, \mathbf{m}) \neq \epsilon \implies \pi_s(c) \neq \epsilon \\ \pi_{mb}(c, \mathbf{m}) \neq \epsilon \implies \pi_{ms}(c, \mathbf{m}) \neq \epsilon \\ \pi \vdash c' \leq c, \pi_f(c', \mathbf{f}) \neq \epsilon \implies \pi_f(c, \mathbf{f}) = \epsilon \\ \pi \vdash c' \leq c, \pi_{ms}(c', \mathbf{m}) \neq \epsilon \implies \pi \vdash \pi_{ms}(c', \mathbf{m}) \leq \pi_{ms}(c, \mathbf{m}) \\ \pi_{mb}(c, \mathbf{m}) = \epsilon \implies \\ \quad \pi_{ms}(c, \mathbf{m}) = c_1, \dots, c_{n+1}, \quad n \geq 0 \\ \quad \pi, (\mathbf{this} \mapsto c, \mathbf{y}_1 \mapsto c_1, \dots, \mathbf{y}_n \mapsto c_n) \vdash \mathbf{e} : c' \\ \quad \pi \vdash c' \leq c_{n+1} \end{array}$$

$$\vdash \pi$$

Figure 4: Well-formed programs

primitive operations *e.g.*, , replacing the signature of a method which is called by other methods⁹. Therefore, although the choice of the primitive operations is attractive at a theoretical level, it is not necessarily the most natural.

4 Soundness

A program is well formed, judgment $\vdash \pi$ in figure 4, if the class hierarchy forms a tree, no field identifier is re-defined in a subclass and superclass, any method overriding a method of a superclass has signature which is a subsignature of the overridden method, and all methods bodies have a signature, and they verify to that signature. We expect to be able to prove that extension, and hence also evaluation, preserves well-formedness of the program.

A heap and stack conform to a program and environment, judgment $\pi, \gamma \vdash \chi, \sigma$ in figure 5, if the receiver and arguments point to objects which conform to their type declared in the environment. An object conforms to a type c if it is of class c' and c' is a subclass of c , and it contains appropriate values for all fields of the class c .

Soundness of the system guarantees that execution of a well typed expression in an appropriate stack and heap and in context of a well formed program will not get stuck, *i.e.*, will not

⁹This has been pointed out by the second anonymous referee!

$$\begin{array}{c}
\frac{\pi \vdash c' \leq c \quad \chi(\iota) = (c', _)}{\pi, \chi \vdash \iota \triangleleft c} \\
\hline
\pi, \chi \vdash \mathbf{0} \triangleleft c
\end{array}
\qquad
\begin{array}{c}
\frac{\chi(\iota) = (c, fm) \quad \mathcal{F}(\pi, c, f) = c' \implies \pi, \chi \vdash fm(f) \triangleleft c'}{\pi, \chi \vdash \iota} \\
\chi(\iota) \neq \epsilon \implies \pi, \chi \vdash \iota \\
\pi, \chi \vdash \sigma(\mathbf{this}) \triangleleft \gamma(\mathbf{this}) \\
\pi, \chi \vdash \sigma(y) \triangleleft \gamma(y) \\
\hline
\pi, \gamma \vdash \chi, \sigma
\end{array}$$

Figure 5: Conformance.

attempt to access non-existing fields or methods, and if it produces a value, this value is guaranteed to conform with the type of the original expression. Soundness makes no guarantees about link and pointer exceptions. We expect to be able to prove the following Theorem:

Theorem 1 (Soundness) *If*

- $\pi, \gamma \vdash \chi, \sigma$, *and* $\vdash \pi$, *and*
- $\pi, \gamma \vdash e : c$ *and* $e, \pi, \sigma, \chi \rightsquigarrow r, \pi', \chi'$

then

- $\vdash \pi'$, *and* $\pi', \gamma \vdash \chi', \sigma$,
- $r \neq \text{StuckExc}$, *and* *if* $r = \kappa$, *then* $\pi', \chi' \vdash \kappa \triangleleft c$

5 Conclusions and further work

We have developed a formal system which models dynamic linking and re-linking at the language level. at a very fine level. The suggested features support unanticipated software evolution.

Further work includes

- completion of the technical details and proofs,
- an investigation in how far the conditions for program extension in this paper are necessary for preservation of program well-formedness,
- a refinement of the model into a series of models reflecting more and more relevant implementation details (*e.g.*, class layout, object layout, offset calculation), and the restrictions they bring into the concept of program extension,
- an investigation in how far the restriction of Java dynamic linking (eg object creation *after loading, no method removal* together with our current model would lead us to the dynamic linking model as in current Java implementations,
- an investigation into ways of expressing dynamic linking operations at a coarser, or more variable granularity,
- extensions to incorporate ideas for unanticipated object reclassification, [4].

Acknowledgements We are very grateful to Giovanni Lagorio, Mariangiola Dezani and Paola Giannini, and the two anonymous FTJP referees for their feedback, suggestions, and careful checking and repairs to the formal model. Furthermore, the second anonymous TJJP referee gave insightful suggestions, and ideas on bridging the gap between theory and practical relevance.

References

- [1] Davide Ancona, Sonia Fagorzi, and Elena Zucca. A Calculus for Dynamic Linking. In *Italian Conference of Theoretical Computer Science*, 2003.
- [2] Gavin Biermann, Michael Hicks, Peter Sewell, and Gareth Stoye. Formalizing Dynamic Software Updating. In *USE Workshop ETAPS*, April 2003.
- [3] Cristiano Calcagno, Peter O’Hearn, and Richard Bornat. Program Logic and Equivalence in the Presence of Gargabe Collection. *Theoretical Computer Science*, 2002.
- [4] Ferruccio Damiani and Paola Giannini. Refined Effects for unanticipated object reclassification, Fickle₃. In *Italian Conference of Theoretical Computer Science*, 2003.
- [5] Mikhail Dimitriev. Hotspot Technology Application for Advanced Profiling. In *ECOOP USE Workshop*, June 2002.
- [6] Sophia Drossopoulou and Susan Eisenbach. Manifestations of Java Dynamic Linking. In *ECOOP Workshop on Unanticipated Software Evolution*, 2002.
- [7] Sophia Drossopoulou, Susan Eisenbach, and Sarfraz Khurshid. Is Java Sound? *Theory and Practice of Object Systems*, 5(1), January 1999.
- [8] Sophia Drossopoulou, Giovanni Lagorio, and Susan Eisenbach. Flexible Models for Dynamic Linking. In *12th European Symposium on Programming*, 2003.
- [9] Susan Eisenbach. Dynamic Linking phases in Java and C#, 2002. <http://www.doc.ic.ac.uk/~sue/foodexample.html>.
- [10] James Goslong, Bill Joy, Guy Steele, and Gilad Bracha. *The JavaTM Language Specification - Second Edition*. Addison-Wesley, 2000.
- [11] Michael Hicks, Jonathan T. Moore, and Scott Nettles. Dynamic Software Updating. In *Programming Language Design and Implementation*. ACM, 2001.
- [12] Zhenyu Qian, Allen Goldberg, and Alessandro Coglio. A Formal Specification of JavaTM Class Loading. In *OOPSLA’2000*, November 2000.

Appendix

Reachable addresses An address ι reaches another address ι' in the context of heap χ , iff for some field f , $\chi(\iota)(f) = \iota'$, or $\chi(\iota)(f)$ reaches ι' . An address ι is reachable from a stack σ in χ , iff there exists a k , so that $1 \leq k \leq n$ and $\sigma = \phi_1 \dots \phi_k \dots \phi_n$, and ϕ_k (**this**) reaches ι , or ϕ_k (y) reaches ι for some y .

Stuck execution and null-pointer exceptions Rules *fldNull*, *fldAssNll*, and *callNull* describe null pointer exceptions. Rules *fldStck*, *fldAssStck*, and *callStck* describe stuck execution, due to a missing field or method body. Rules *propNull*, and *propStck* describe propagation of exceptions.

In particular, as we discussed in section 2, an undefined signature for the particular class, i.e., $\mathcal{MS}(\dots, \dots) = \epsilon$ indicates that the class has no such method, and then execution is stuck, while a defined signature and undefined method body, i.e., $\mathcal{MS}(\dots, \dots) \neg \epsilon = \mathcal{MB}(\dots, \dots)$ indicates a verification error for the method body, and then execution throws a link related exception.

$\frac{\text{fldNull}}{\frac{e, \pi, \sigma, \chi \rightsquigarrow \mathbf{0}, \pi, \chi'}{e.f, \pi, \sigma, \chi \rightsquigarrow \text{nllPtrExc}, \pi, \chi'}}$	$\frac{\text{fldAssNull}}{\frac{e, \pi, \sigma, \chi \rightsquigarrow \mathbf{0}, \pi', \chi'}{e', \pi', \sigma, \chi' \rightsquigarrow \kappa, \pi'', \chi''} \quad \frac{e.f := e', \pi, \sigma, \chi \rightsquigarrow \text{nllPtrExc}, \pi'', \chi''}}$
$\frac{\text{callNull}}{\frac{e, \pi, \sigma, \chi \rightsquigarrow \mathbf{0}, \pi_1, \chi_1 \quad e_i, \pi_i, \sigma, \chi_i \rightsquigarrow \kappa_i, \pi_{i+1}, \chi_{i+1} \quad 1 \leq i \leq n}{e.m(e_1, \dots, e_n), \pi, \sigma, \chi \rightsquigarrow \text{nllPtrExc}, \pi', \chi'}}$	$\frac{\text{propNll}}{\frac{e', \pi, \sigma, \chi \rightsquigarrow \text{nllPtrExc}, \pi, \chi \quad e \text{ has subexpression } e'}{e, \pi, \sigma, \chi \rightsquigarrow \text{nllPtrExc}, \pi, \chi}}$
$\frac{\text{fldStck}}{\frac{e, \pi, \sigma, \chi \rightsquigarrow \iota, \pi, \chi' \quad \chi'(\iota)(f) = \epsilon}{e.f, \pi, \sigma, \chi \rightsquigarrow \text{stuckExc}, \pi, \chi'}}$	$\frac{\text{fldAssStck}}{\frac{e, \pi, \sigma, \chi \rightsquigarrow \iota, \pi', \chi' \quad e', \pi', \sigma, \chi' \rightsquigarrow \kappa, \pi'', \chi'' \quad \chi''(\iota)(f) = \epsilon}{e.f := e', \pi, \sigma, \chi \rightsquigarrow \text{stuckExc}, \pi'', \chi''}}$
$\frac{\text{callStck}}{\frac{e, \pi, \sigma, \chi \rightsquigarrow \mathbf{0}, \pi_1, \chi_1 \quad e_i, \pi_i, \sigma, \chi_i \rightsquigarrow \kappa_i, \pi_{i+1}, \chi_{i+1} \quad 1 \leq i \leq n \quad \chi(\iota) = (c, \dots) \quad \mathcal{MS}(\pi, c, m) = \epsilon}{e.m(e_1, \dots, e_n), \pi, \sigma, \chi \rightsquigarrow \text{stuckExc}, \pi', \chi'}}$	$\frac{\text{propStck}}{\frac{e', \pi, \sigma, \chi \rightsquigarrow \text{stuckExc}, \pi, \chi \quad e \text{ has subexpression } e'}{e, \pi, \sigma, \chi \rightsquigarrow \text{stuckExc}, \pi, \chi}}$

Subtypes The judgment $\pi_s \vdash c' \leq c$ means that c' is a (possibly indirect) subclass of c according to π_s , while $\pi_s \vdash c' \not\leq c$ means that $c' \neq c$, and either $\pi_s \vdash c' \leq c$ or $\pi_s \vdash c \leq c'$. Both judgments can be extended to full programs, giving judgments $\pi \vdash c' \leq c$ and $\pi \vdash c' \not\leq c$:

$\frac{\pi_s(c') = c}{\pi_s \vdash c' \leq c}$	$\frac{\pi_s \vdash c' \leq c' \quad \pi_s \vdash c'' \leq c}{\pi_s \vdash c' \leq c}$	$\frac{\pi_s \vdash c' \leq c \quad \pi = (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}{\pi \vdash c' \leq c}$
$\frac{\pi_s \vdash c' \leq c \quad c \neq c'}{\pi_s \vdash c' \not\leq c}$	$\frac{\pi_s \vdash c \leq c' \quad c \neq c'}{\pi_s \vdash c' \not\leq c}$	$\frac{\pi_s \vdash c' \not\leq c \quad \pi = (\pi_s, \pi_f, \pi_{ms}, \pi_{mb})}{\pi \vdash c' \not\leq c}$

Subsignatures The judgment $\pi \vdash c'_1 \dots c'_n \leq c_1 \dots c_n$ says that the signature $c'_1 \dots c'_n$ may safely override signature $c_1 \dots c_n$ in a subclass. The first rule says that any signature may override a non existing definition, while the second rule says that subsignatures are contravariant in argument types and covariant in result types.

$$\frac{}{\pi \vdash c'_1 \dots c'_n \leq c} \quad \frac{\pi \vdash c_i \leq c'_i \quad i \in 1..n-1 \quad \pi \vdash c'_n \leq c_n}{\pi \vdash c'_1 \dots c'_n \leq c_1 \dots c_n}$$