

This is a preprint version of chapter four of

```
@incollection{Oheimb-Nipkow-Java-LNCS,  
author      = {Oheimb, David von and Nipkow, Tobias},  
title       = {Machine-checking the {J}ava Specification:  
               Proving Type-Safety},  
booktitle   = {Formal Syntax and Semantics of {J}ava},  
editor      = {Jim Alves-Foss},  
url         = {http://isabelle.in.tum.de/Bali/doc/Springer98.html},  
publisher   = {Springer},  
series      = {LNCS},  
volume      = {1523},  
pages       = {119-156}  
year        = {1999}  
}
```

Machine-checking the Java Specification: Proving Type-Safety*

David von Oheimb** and Tobias Nipkow

Fakultät für Informatik, Technische Universität München

<http://www.in.tum.de/~oheimb>

<http://www.in.tum.de/~nipkow>

Abstract. In this article we present BALI, the formalization of a large (hitherto sequential) sublanguage of Java. We give its abstract syntax, type system, well-formedness conditions, and an operational evaluation semantics. Based on these definitions, we can express soundness of the type system, an important design goal claimed to be reached by the designers of Java, and prove that BALI is indeed type-safe.

All definitions and proofs have been done formally in the theorem prover Isabelle/HOL. Thus this article demonstrates that machine-checking the design of non-trivial programming languages has become a reality.

1 Introduction

BALI is a large subset of Java [GJS96]. This article presents its formalization and the proof of a key property, namely the soundness of its type system — specified and verified in the theorem prover Isabelle/HOL [Pau94].

On the face of it, this article is mostly about BALI, its abstract syntax, type system, well-formedness conditions, and operational semantics, formalized as a hierarchy of Isabelle theories, and the structure of the machine-checked proof of type soundness and its implications. Although these technicalities do indeed take up much of the space, there is a meta-theme running through the article, which we consider even more important: the technology for producing machine-checked programming language designs has arrived.

We emphasize that by ‘machine-checked’ we do not just mean that it has passed some type checker, but that some non-trivial properties of the language have been established with the help of a (semi-automatic) theorem prover. The latter process is still not a piece of cake, but it has become more than just feasible. Therefore any programming language intended for serious applications should strive for such a machine-checked design. The benefits are not just greater reliability, but also greater maintainability because the theorem prover keeps track of the impact that changes have on already established properties.

* This is a completely revised and extended version of [NO98].

** Research supported by DFG SPP *Deduktion*.

Note that the type-safety of Java is not sufficient to guarantee secure execution of bytecode programs on the Java Virtual Machine, because the bytecode might be tampered with, produced by a faulty compiler, or not be related to any Java source program at all. This was the main reason for introducing the Bytecode Verifier in the JVM, which checks the integrity, in particular type-correctness, of any bytecode before execution. Of course similar security problems arise for any other high-level languages as well. Nevertheless, the investigation of type-safety at source level is worthwhile: it checks whether the language design is sound, which means that at least all the necessary conditions expressible at that level are fulfilled. In particular static typing loses much of its *raison d'être* if the language is not type-safe.

1.1 Related work

The history of type soundness proofs goes back to the subject reduction theorem for typed λ -calculus but starts in earnest with Milner's slogan 'Well-typed expressions do not go wrong' [Mil78] in the context of ML. Milner uses a denotational semantics, in contrast to most of the later work, including ours. The question of type-safety came to prominence with the discovery of its failure in Eiffel [Coo89]. Ever since, many designers of programming languages (especially OO ones) have been at pains to prove type-safety of their languages (see, for example, the series of papers by Bruce et al. [Bru93,BCM⁺93,BvGS95]).

Directly related to our work is that by Drossopoulou and Eisenbach [DE98] who prove (on paper) type-safety of a subset of Java very similar to BALI. Although we were familiar with an earlier version [DE97] of their work and have certainly profited from it, our work is not a formalization of theirs in Isabelle/HOL but differs in many respects from it, for example in the representation of programs and the use of an evaluation (aka "big-step") semantics instead of a transition (aka "small-step") semantics. Simultaneously with our work, Syme [Sym98] formalized the paper [DE97] as far as possible, uncovering two significant mistakes, both in connection with the use of transition semantics. Syme uses his own theorem prover DECLARE, also based on higher-order logic.

There are two other efforts to formalize aspects of Java in a theorem prover. Dean [Dea97] studies the interaction of static typing with dynamic linking. His simple PVS specification addresses only the linking aspect and requires a formalization of Java (such as our work provides) to turn his lemmas about linking into theorems about the type-safety of dynamically linked programs. Cohen [Coh97] has formalized the semantics of large parts of the Java Virtual Machine, essentially by writing an interpreter in Common Lisp. He used ACL2, the latest version of the Boyer-Moore theorem prover [BM88]. No proofs have been reported yet.

2 Overview

BALI includes the features of Java that we believe to be important for an investigation of the semantics of a practical imperative object-oriented language:

- class and interface declarations with instance fields and methods,
- subinterface, subclass, and implementation relations with inheritance, overriding, and hiding,
- method calls with static overloading and dynamic binding,
- some primitive types, objects (including arrays),
- exception throwing and handling.

This portion of Java is very similar to that covered by [DE98] and [Sym98].

We do not consider Java packages and separate compilation. For the moment, we also leave out several features of Java like class variables and static methods, constructors and finalizers, the visibility of names, and concurrency, but we aim to include at least part of them in later stages of our project. Some constructs are simplified without limiting the expressiveness of the language (see §4.1).

In developing the formalization of BALI and investigating its properties, we aim to meet the following design goals:

- faithfulness to the official language specification,
- succinctness and simplicity,
- maintainability and extendibility,
- adequacy for the theorem prover.

It might be interesting to keep these goals in mind while reading §4 on the formalization of BALI and §5 on our proofs and check how far we have reached them. We comment on our experience in pursuing these goals in §6.

3 The basics of Isabelle/HOL

Before we present the formalization of BALI, we briefly introduce the underlying theorem proving system Isabelle/HOL.

Isabelle/HOL is the instantiation of the generic interactive theorem prover Isabelle [Pau94] with Church’s version of Higher-Order Logic and is very close to Gordon’s HOL system [GM93]. In this article HOL is short for Isabelle/HOL.

The appearance of formulas is standard, e.g. ‘ \longrightarrow ’ is the (right-associative) infix implication symbol. Predicates are functions with Boolean result. Function application is written in curried style. For descriptions we apply Hilbert’s choice operator ε , where $\varepsilon x. P x$ denotes some value x satisfying P , or an arbitrary value if no such x exists.

Logical constants are declared by giving their name and type, separated by ‘ $::$ ’. Primitive recursive function definitions are written as usual. Non-recursive definitions are written with ‘ $\stackrel{\text{def}}{=}$ ’.

Types follow the syntax of ML, except that the function arrow is ‘ \Rightarrow ’. Type abbreviations are introduced simply as equations. A free datatype is defined by listing its constructors together with their argument types, separated by ‘ $|$ ’.

There are the basic types *bool* and *int*, as well as the polymorphic type $(\alpha)set$ of homogeneous sets for any type α . Occasionally we apply the infix ‘image’ operator lifting a function over a set, defined as $f \circ S \stackrel{\text{def}}{=} \{y. \exists x \in S. y = f x\}$.

The product type $\alpha \times \beta$ comes with the projection functions *fst* and *snd*. Tuples are pairs nested to the right, e.g. $(a,b,c) = (a,(b,c))$.

As the list type $(\alpha)list$ is defined via its constructors $[]$ denoting the empty list and the infix ‘cons’ operator ‘#’, it can be introduced by the datatype declaration

$$(\alpha)list = [] \mid \alpha \# (\alpha)list$$

The concatenation operator on lists is written as the infix symbol ‘@’. There is a functional *map* $:: (\alpha \Rightarrow \beta) \Rightarrow (\alpha)list \Rightarrow (\beta)list$ applying a function to all elements of a list, as well as a conversion function *set* $:: (\alpha)list \Rightarrow (\alpha)set$.

We frequently use the datatype

$$(\alpha)option = None \mid Some \alpha$$

It has an unpacking function *the* $:: (\alpha)option \Rightarrow \alpha$ such that *the* $(Some\ x) = x$ and *the* $None = \text{arbitrary}$, where *arbitrary* is an unknown value defined as $\varepsilon x.False$. There is a simple function mapping *o2s* $:: (\alpha)option \Rightarrow (\alpha)set$ converting an optional value to a set, with the characteristic equations $o2s\ (Some\ x) = \{x\}$ and $o2s\ None = \{\}$.

Most of the HOL text shown in this article is directly taken from the input files. However, it has been massaged by hand to hide Isabelle idiosyncrasies, increase readability, and adapt the layout. Minor typos may have been introduced in the process.

We adopt the following typographic conventions: Java keywords like *catch* appear in typewriter font, the names of logical constants like *cfield* appear in sans serif, while type names like *state* and meta-variables like *v* appear in italics.

4 The formalization of Bali

This section presents all aspects of our formalization of BALI¹.

As far as BALI is a subset of Java, it strictly adheres to the Java language specification [GJS96], with several generalizations:

- we allow the result type of a method overriding another method to widen to the result type of the other method instead of requiring it to be identical.
- if a class or an interface inherits more than one method with the same signature, the methods need not have identical return types.
- no check of result types in dynamic method lookup.
- the type of an assignment is determined by the right-hand side, which can be more specific than the left-hand side.

¹ The Isabelle sources are available from the BALI project page <http://www.in.tum.de/~isabelle/bali/>

We found several issues concerning exceptions not specified in [GJS96] and therefore define a reasonable behavior that seems to be consistent with current implementations:

- given a `Null` reference, the `throw` statement raises a `NullPointerException` exception.
- each system exception thrown yields a fresh exception object.
- if there is not enough memory even to allocate an `OutOfMemory` error, program execution simply halts. (Our experiments showed erratic behavior of some implementation in this case, ranging from sudden termination without executing `finally` blocks, over hangup, to repeated invocation of a single exception handler!)

To illustrate our approach, we use the following (artificial) example.

```
class Base {
  boolean vee;
  Base foo(Base x) {
    return x;
  }
}

class Ext extends Base{
  int vee;
  Ext foo(Base x) {
    ((Ext)x).vee=1;
    return null;
  }
}

Base e;
e=new Ext();
try {e.foo(null); }
catch (NullPointerException x) {throw x;}
```

This program fragment consists of two simple but complete class declarations and a block of statements that might occur in any method that has access to these declarations. It contains the following features of BALI:

- class declarations with inheritance, hiding of fields, and overriding of methods (with refined result type),
- return expressions, parameter access,
- sequential composition, expression statements, field assignment, type cast, local accesses, literal values, exception propagation,
- local assignment, instance creation,
- try & catch statement, method call (with dynamic binding), throw statement

4.1 Abstract syntax

First, we describe how we represent the syntax of BALI and which abstractions we have introduced thereby.

Programs A BALI program is a pair of lists of interface and class declarations:

$$prog = (idecl)list \times (cdecl)list$$

Throughout the article, the symbol ‘ T ’ denotes a BALI program, as we use programs as part of the static type context defined in §4.2.

Each declaration is a pair of a name and the defined entity. Some names, like those of predefined classes (including those of system exceptions $xname$), have a predefined meaning and are therefore handled extra. We do not specify the structure of names further, but use the opaque HOL types $tname0$, $mname$, and $ename0$ for user-defined type names, method names, and “expression names” (i.e. field and variable identifiers).

$xname =$	Throwable	
	NullPointer	OutOfMemory ClassCast
	NegArrSize	IndOutBound ArrStore
$tname =$	Object	name of the top of the class hierarchy
	SXcpt $xname$	name of a system exception
	TName $tname0$	other class or interface name
$ename =$	this	special name for this pointer
	EName $ename0$	expression name

An interface ($iface$) contains lists of superinterface names and method declarations. A $class$ specifies the names of an optional superclass and implemented interfaces, and lists of field and method declarations.

$$\begin{aligned}
 iface &= (tname)list \times (sig \times mhead)list \\
 idecl &= tname \times iface \\
 class &= (tname)option \times (tname)list \times (fdecl)list \times (mdecl)list \\
 cdecl &= tname \times class
 \end{aligned}$$

A field declaration ($fdecl$) simply gives the field type (ty , see §4.2). A method declaration ($sig \times mhead$ for interfaces or $mdecl$ for classes) consists of a “signature” [GJS96, 8.4.2] (i.e. the method name and the list of parameter types, excluding the result type) followed by $mhead$, consisting of the list of parameter names and the result type, and (if it appears within a class) the method body ($mbody$). The latter consists of the list of local variables, a statement $stmt$ as body, and a return expression $expr$ (see below). As in [DE98], the separate return expression saves us from dealing with return statements occurring in arbitrary positions within the method body. Such statements may be replaced by

assignments to a suitable result variable followed by a control transfer to the end of the method body, using the result variable as return expression. We provide a dummy result type and value for void methods. For simplicity, up to now each method has exactly one parameter; multiple parameters can be simulated by a single parameter object with multiple fields.

<i>field</i>	= <i>ty</i>	field type
<i>fdecl</i>	= <i>ename</i> × <i>field</i>	
<i>sig</i>	= <i>mname</i> × <i>ty</i>	method name and parameter type
<i>mhead</i>	= <i>ename</i> × <i>ty</i>	parameter name and result type
<i>lvar</i>	= <i>ename</i> × <i>ty</i>	local variable name and type
<i>mbody</i>	= (<i>lvar</i>) <i>list</i> × <i>stmt</i> × <i>expr</i>	local vars, block, and return expression
<i>methd</i>	= <i>mhead</i> × <i>mbody</i>	method (of a class)
<i>mdecl</i>	= <i>sig</i> × <i>methd</i>	

In the abstract syntax given above, the formalization of our example program fragment looks like this:

```

BaseC def (Base, (Some Object,
                [],
                [(vee, PrimT boolean)],
                [(foo,Class Base),(x,Class Base), ([],Skip,x)]))
ExtC def (Ext, (Some Base,
               [],
               [(vee, PrimT int)],
               [(foo,Class Base),(x,Class Ext), ([],
               Expr({ClassT Ext}(Class Ext)x.vee:=Lit (Intg 1)),
               Lit Null)]))
classes def [ObjectC,
            SXcptC Throwable,
            SXcptC NullPointerException, SXcptC OutOfMemory, SXcptC ClassCast,
            SXcptC NegArrSize, SXcptC IndOutBound, SXcptC ArrStore,
            BaseC, ExtC]
tprg def ([],classes)
test def Expr(e:=new Ext);
      try Expr(e.foo({Class Base}Lit Null))
      catch((SXcpt NullPointerException) x) (throw x)

```

where *Base* stands for TName *Base_*, *Ext* for TName *Ext_*, and similarly for *vee*, *x*, and *e*. The constants *Base_*, *Ext_*, etc. are all distinct. The sequence of statements *test* could have been embedded in *tprg*, which we have left out for simplicity.

Representation of lookup tables The representation of declarations as lists gives an implicit finiteness constraint, which turns out to be necessary for the well-foundedness of the subclass and subinterface relation. The list representation also enables an explicit check whether the declared entities are named uniquely, implemented with the function `unique` given below. This function does not check for duplicate definitions, which is harmless.

`unique` :: $(\alpha \times \beta)list \Rightarrow bool$
`unique` $t \stackrel{\text{def}}{=} \forall (x_1, y_1) \in \text{set } t. \forall (x_2, y_2) \in \text{set } t. x_1 = x_2 \longrightarrow y_1 = y_2$

For the lookup of declared entities, we transform declaration lists into abstract tables. They are realized in HOL as “partial” functions mapping names to values:

$(\alpha, \beta)table = \alpha \Rightarrow (\beta)option$

The empty table, pointwise update, extension of one table by another, the function converting a declaration list into a table, and an auxiliary predicate relating entries of two tables, are defined easily:

`empty` :: $(\alpha, \beta)table$
`-[_↦-]` :: $(\alpha, \beta)table \Rightarrow \alpha \Rightarrow \beta \Rightarrow (\alpha, \beta)table$
`-⊕-` :: $(\alpha, \beta)table \Rightarrow (\alpha, \beta)table \Rightarrow (\alpha, \beta)table$
`table_of` :: $(\alpha \times \beta)list \Rightarrow (\alpha, \beta)table$
`- hiding -`
`entails -` :: $(\alpha, \beta)table \Rightarrow (\alpha, \gamma)table \Rightarrow (\beta \Rightarrow \gamma \Rightarrow bool) \Rightarrow bool$

`empty` $\stackrel{\text{def}}{=} \lambda k. \text{None}$
`t[x↦y]` $\stackrel{\text{def}}{=} \lambda k. \text{if } k = x \text{ then Some } y \text{ else } t\ k$
`s ⊕ t` $\stackrel{\text{def}}{=} \lambda k. \text{case } t\ k \text{ of None } \Rightarrow s\ k \mid \text{Some } x \Rightarrow \text{Some } x$

`table_of []` = `empty`
`table_of ((k,x)#t)` = `(table_of t)[k↦x]`

$t \text{ hiding } s \text{ entails } R \stackrel{\text{def}}{=} \forall k\ x\ y. t\ k = \text{Some } x \longrightarrow s\ k = \text{Some } y \longrightarrow R\ x\ y$

For the union of tables, we also need the type of non-unique tables,

$(\alpha, \beta)tables = \alpha \Rightarrow (\beta)set$

together with a union operator and straightforward variants of two of the notions defined above:

`-⊕⊕-` :: $(\alpha, \beta)tables \Rightarrow (\alpha, \beta)tables \Rightarrow (\alpha, \beta)tables$
`Un_tables` :: $((\alpha, \beta)tables)set \Rightarrow (\alpha, \beta)tables$
`- hidings -`
`entails -` :: $(\alpha, \beta)tables \Rightarrow (\alpha, \gamma)tables \Rightarrow (\beta \Rightarrow \gamma \Rightarrow bool) \Rightarrow bool$

$$\begin{aligned}
\text{Un_tables } ts & \stackrel{\text{def}}{=} \lambda k. \bigcup_{t \in ts}. t \ k \\
s \oplus \oplus t & \stackrel{\text{def}}{=} \lambda k. \text{ if } t \ k = \{\} \text{ then } s \ k \text{ else } t \ k \\
t \text{ hidings } s \text{ entails } R & \stackrel{\text{def}}{=} \forall k. \forall x \in t \ k. \forall y \in s \ k. R \ x \ y
\end{aligned}$$

A simple application of type *table* is the translation of programs to tables indexed by interface and class names:

$$\begin{aligned}
\text{iface } :: \text{ prog} \Rightarrow (\text{tname}, \text{iface})\text{table} & \stackrel{\text{def}}{=} \text{table_of} \circ \text{fst} \\
\text{class } :: \text{ prog} \Rightarrow (\text{tname}, \text{class})\text{table} & \stackrel{\text{def}}{=} \text{table_of} \circ \text{snd}
\end{aligned}$$

More interesting are the following functions that traverse the type hierarchy of a program, collecting the methods and fields into a table (the types *tname* and *ref_ty* are defined in §4.2):

$$\begin{aligned}
\text{imethds } :: \text{ prog} \Rightarrow \text{tname} \Rightarrow (\text{sig}, \quad \text{ref_ty} \times \text{mhead})\text{tables} \\
\text{cmethd } :: \text{ prog} \Rightarrow \text{tname} \Rightarrow (\text{sig}, \quad \text{ref_ty} \times \text{methd})\text{table} \\
\text{fields } \quad :: \text{ prog} \Rightarrow \text{tname} \Rightarrow ((\text{ename} \times \text{ref_ty}) \times \text{field} \)\text{list}
\end{aligned}$$

Note that *imethds* collects a non-unique table of method declarations allowing for inheritance of more than one method with the same signature.

As Syme [Sym98] points out, a naive recursive definition of these functions is not possible in HOL because the class hierarchy might be cyclic, which is ruled out for well-formed programs (see §4.3) only. This leads to partial functions, which HOL does not support directly. Syme defines these functions as relations instead. In contrast, we have chosen to define them as proper functions, based on Slind’s work on well-founded recursion [Sli96]. We do not give their definitions, but only the recursion equations, which we derive as easy consequences:

$$\begin{aligned}
\text{wf_prog } \Gamma \wedge \text{iface } \Gamma \ I = \text{Some } (is, ms) & \longrightarrow \\
\text{imethds } \Gamma \ I = \text{Un_tables } ((\lambda J. \text{imethds } \Gamma \ J) \text{ "set } is) \oplus \oplus & \\
\quad (\text{o2s} \circ \text{table_of } (\text{map } (\lambda (s, mh). (s, \text{lfaceT } I, mh)) \ ms)) &
\end{aligned}$$

$$\begin{aligned}
\text{wf_prog } \Gamma \wedge \text{class } \Gamma \ C = \text{Some } (sc, si, fs, ms) & \longrightarrow \\
\text{cmethd } \Gamma \ C = (\text{case } sc \text{ of None} \Rightarrow \text{empty} \mid \text{Some } D \Rightarrow \text{cmethd } \Gamma \ D) \oplus & \\
\quad \text{table_of } (\text{map } (\lambda (s, m). (s, (\text{ClassT } C, m))) \ ms) &
\end{aligned}$$

$$\begin{aligned}
\text{wf_prog } \Gamma \wedge \text{class } \Gamma \ C = \text{Some } (sc, si, fs, ms) & \longrightarrow \\
\text{fields } \Gamma \ C = \text{map } (\lambda (fn, ft). ((fn, \text{ClassT } C), ft)) \ fs \ @ & \\
\quad (\text{case } sc \text{ of None} \Rightarrow [] \mid \text{Some } D \Rightarrow \text{fields } \Gamma \ D) &
\end{aligned}$$

The structure of the three equations is the same: the tables are constructed recursively from the corresponding tables of the superinterfaces or the superclass (if any), which models inheritance, augmented — with overriding — by the newly declared items. All declared items receive an extra label, namely their defining interface or class.

In our example, we obtain

```

fields tprg Base = [((vee, ClassT Base), PrimT boolean)]
fields tprg Ext   = [((vee, ClassT Ext), PrimT int),
                  ((vee, ClassT Base), PrimT boolean)]
cmethd tprg Base = empty[(foo, Class Base)→
                        (ClassT Base, (x, Class Base), ([], Skip, x))]
cmethd tprg Ext  = empty[(foo, Class Base)→
                        (ClassT Ext, (x, Class Ext), ([],
                        Expr({ClassT Ext}(Class Ext)x.vee:Lit (Intg 1)),
                        Lit Null))]

```

Terms We define statements (appearing in method bodies), expressions (appearing in statements), and values (appearing in expressions) as recursive data-types.

Statements are reduced to their bare essentials. We do not formalize syntactic variants of conditionals and loops. Neither do we consider jumps like the `break` statement.

For a more modular description, we divide the `try - catch - finally -` statement into a `try - catch -` statement and a `- finally -` statement, which might be used in any context. Our version of the `try - catch -` statement has exactly one `catch` clause. Multiple `catch` clauses can be simulated with cascaded `if - else -` statements applying the `- instanceof -` operator.

```

stmt = Skip
      | Expr expr
      | stmt; stmt
      | if (expr) stmt else stmt
      | while(expr) stmt
      | throw(expr)
      | try stmt catch(tname ename) stmt
      | stmt finally stmt

```

`Skip` denotes the empty statement. The “expression statement” `Expr` is a conversion from expressions to statements causing evaluation for side effects only. Assignments and method calls, which are expressions because they yield a value, can be turned into statements this way. In contrast to Java, for simplicity we allow this conversion to be applied to any kind of expression.

Concerning expressions, our formalization leaves out the standard unary and binary operators as their typing and semantics is straightforward. The `this` expression is modeled as a special non-assignable local variable named `this`. The `super` construct can be simulated with a `this` expression that is cast to the superclass of the current class. Creation of multi-dimensional arrays can be simulated with nested array creation, while access and assignment to multi-dimensional arrays is nested anyway.

It might be reasonable to introduce the general notion of variables (i.e. left-hand sides of assignments) in order to factor out common behavior of local variables, class instance variables, and array components. But we have chosen not to do so because the semantic treatment of these three variants of variables differs considerably. This decision leads to some redundancy between access and assignment, especially in the semantics for arrays.

$expr = \text{new } tname$	class instance creation
$\text{new } ty[expr]$	array creation
$(ty) expr$	type cast
$expr \text{ instanceof } ref_ty$	type comparison operator
$\text{Lit } val$	literal
$ename$	local/parameter access
$ename := expr$	local/parameter assignment
$\{ref_ty\} expr. ename$	field access
$\{ref_ty\} expr. ename := expr$	field assignment
$expr[expr]$	array access
$expr[expr] := expr$	array assignment
$expr. mname(\{ty\} expr)$	method call

The terms in braces $\{\dots\}$ above are called *type annotations*. Strictly speaking, they are not part of the input language but serve as auxiliary information (computed by the type checker) that is crucial for the static binding of fields and the resolution of method overloading. Distinguishing between the actual input language and the augmented language would lead to a considerable amount of redundancy. We avoid this by assuming that the annotations are added beforehand by a kind of preprocessor. The correctness of the annotations is checked by the typing rules (see §4.2).

The definition of values is straightforward. It relies on the standard HOL types of Boolean values (*bool*) and integers (*int*), whereas the type *loc* of locations, i.e. abstract non-null addresses of objects, is not further specified.

$val = \text{Unit}$	dummy result of void methods
$\text{Bool } bool$	
$\text{Intg } int$	
Null	
$\text{Addr } loc$	

The definitions below give some simple destructor functions for *val* with their characteristic properties.

```

the_Bool :: val ⇒ bool
the_Intg  :: val ⇒ int
the_Addr  :: val ⇒ loc

```

```

the_Bool (Bool b) = b
the_Intg (Intg i) = i
the_Addr (Addr a) = a

```

4.2 Type system

This section defines types, various ordering relations between types, and the typing rules for statements and expressions.

Types We formalize BALI types as values of datatype ty , dividing them into primitive and reference types:

```
prim_ty = void
        | boolean
        | int

ref_ty  = NullT
        | lfaceT tname
        | ClassT tname
        | ArrayT ty

ty      = PrimT prim_ty
        | RefT ref_ty
```

`void` is used as a dummy type for methods without result. In the sequel `NT` stands for `RefT NullT`, `lface I` for `RefT(lfaceT I)`, `Class C` for `RefT(ClassT C)`, and `T[]` for `RefT(ArrayT T)`.

An interface or class type is considered as a proper type only if there is a corresponding declaration for its type name in the current program, which is checked by the following predicates:

```
is_iface :: prog ⇒ tname ⇒ bool
is_class :: prog ⇒ tname ⇒ bool
is_type  :: prog ⇒ ty    ⇒ bool

is_iface Γ tn  $\stackrel{\text{def}}{=} \text{iface } \Gamma \text{ tn} \neq \text{None}$ 
is_class Γ tn  $\stackrel{\text{def}}{=} \text{class } \Gamma \text{ tn} \neq \text{None}$ 
is_type  Γ (PrimT _) = True
is_type  Γ NT       = True
is_type  Γ (lface I) = is_iface Γ I
is_type  Γ (Class C) = is_class Γ C
is_type  Γ (T[])     = is_type  Γ T
```

For all types, a default value is defined via

```
default_val :: ty ⇒ val
default_val (PrimT void    ) = Unit
default_val (PrimT boolean) = Bool False
default_val (PrimT int     ) = Intg 0
default_val (RefT r        ) = Null
```

Type relations The relations between types depend on the interface and class hierarchy of a given program Γ , and are therefore expressed with reference to Γ . The direct subinterface ($_ \vdash _ \prec_i^1 _$), subclass ($_ \vdash _ \prec_c^1 _$), and implementation ($_ \vdash _ \rightsquigarrow^1 _$) relations are of type $prog \times tname \times tname \Rightarrow bool$ and are defined as follows:

$$\begin{aligned} \Gamma \vdash I \prec_i^1 J &\stackrel{\text{def}}{=} \text{is_iface } \Gamma \ I \ \wedge \ \text{is_iface } \Gamma \ J \ \wedge \quad J \in \text{set}(\text{fst}(\text{the}(\text{iface } \Gamma \ I))) \\ \Gamma \vdash C \prec_c^1 D &\stackrel{\text{def}}{=} \text{is_class } \Gamma \ C \ \wedge \ \text{is_class } \Gamma \ D \ \wedge \ \text{Some } D = \text{fst}(\text{the}(\text{class } \Gamma \ C)) \\ \Gamma \vdash C \rightsquigarrow^1 I &\stackrel{\text{def}}{=} \text{is_class } \Gamma \ C \ \wedge \ \text{is_iface } \Gamma \ I \ \wedge \ I \in \text{set}(\text{fst}(\text{snd}(\text{the}(\text{class } \Gamma \ C)))) \end{aligned}$$

The transitive (but not reflexive) closures $_ \vdash _ \prec_i _$ and $_ \vdash _ \prec_c _$ can be defined inductively:

$$\begin{array}{c} \frac{\Gamma \vdash I \prec_i^1 K}{\Gamma \vdash I \prec_i K} \quad \frac{\Gamma \vdash I \prec_i J; \Gamma \vdash J \prec_i K}{\Gamma \vdash I \prec_i K} \\ \frac{\Gamma \vdash C \prec_c^1 E}{\Gamma \vdash C \prec_c E} \quad \frac{\Gamma \vdash C \prec_c D; \Gamma \vdash D \prec_c E}{\Gamma \vdash C \prec_c E} \end{array}$$

There is also a kind of transitive closure of $_ \vdash _ \rightsquigarrow^1 _$ defined as

$$\frac{\Gamma \vdash C \rightsquigarrow^1 J}{\Gamma \vdash C \rightsquigarrow J} \quad \frac{\Gamma \vdash C \rightsquigarrow^1 I; \Gamma \vdash I \prec_i J}{\Gamma \vdash C \rightsquigarrow J} \quad \frac{\Gamma \vdash C \prec_c^1 D; \Gamma \vdash D \rightsquigarrow J}{\Gamma \vdash C \rightsquigarrow J}$$

The key relation is widening: $\Gamma \vdash S \preceq T$, where S and T are of type ty , means that S is a syntactic subtype of T , i.e. in any expression context (especially assignments and method invocations) expecting a value of type T , a value of type S may occur. Note that this does not necessarily mean that type S behaves like type T , but only that it has a syntactically compatible set of fields and methods. The widening relation is defined inductively as given below. Note that some rules carry the additional premise that `Object` is a proper class, which will be ensured for well-formed programs.

$$\begin{array}{c} \frac{\text{is_type } \Gamma \ T}{\Gamma \vdash T \preceq T} \quad \frac{\text{is_type } \Gamma \ (\text{RefT } R)}{\Gamma \vdash \text{NT} \preceq \text{RefT } R} \\ \frac{\Gamma \vdash I \prec_i J}{\Gamma \vdash \text{lface } I \preceq \text{lface } J} \quad \frac{\text{is_iface } \Gamma \ I; \text{is_class } \Gamma \ \text{Object}}{\Gamma \vdash \text{lface } I \preceq \text{Class } \text{Object}} \\ \frac{\Gamma \vdash C \prec_c D}{\Gamma \vdash \text{Class } C \preceq \text{Class } D} \quad \frac{\Gamma \vdash C \rightsquigarrow I}{\Gamma \vdash \text{Class } C \preceq \text{lface } I} \\ \frac{\Gamma \vdash \text{RefT } S \preceq \text{RefT } T}{\Gamma \vdash (\text{RefT } S)[] \preceq (\text{RefT } T)[]} \quad \frac{\text{is_type } \Gamma \ T; \text{is_class } \Gamma \ \text{Object}}{\Gamma \vdash T[] \preceq \text{Class } \text{Object}} \end{array}$$

To allow for type casting we also have the casting relation, where $\Gamma \vdash S \preceq_? T$ means that a value of type S may be cast to type T :

$$\begin{array}{c} \frac{\Gamma \vdash S \preceq T}{\Gamma \vdash S \preceq_? T} \quad \frac{\Gamma \vdash C \prec_c D}{\Gamma \vdash \text{Class } D \preceq_? \text{Class } C} \quad \frac{\text{is_class } \Gamma \ C; \text{is_iface } \Gamma \ I}{\Gamma \vdash \text{Class } C \preceq_? \text{lface } I} \\ \frac{\Gamma \vdash \text{RefT } S \preceq_? \text{RefT } T}{\Gamma \vdash (\text{RefT } S)[] \preceq_? (\text{RefT } T)[]} \quad \frac{\text{is_class } \Gamma \ \text{Object}; \text{is_type } \Gamma \ T}{\Gamma \vdash \text{Class } \text{Object} \preceq_? T[]} \end{array}$$

$$\frac{\text{is_iface } \Gamma \ J; \neg \Gamma \vdash I \prec_i J; \text{imethds } \Gamma \ I \ \text{hidings imethds } \Gamma \ J \ \text{entails} \ (\lambda(m_1, (pn_1, rT_1)) \ (m_2, (pn_2, rT_2))). \ \Gamma \vdash rT_1 \preceq rT_2}{\Gamma \vdash \text{lface } I \preceq? \text{lface } J} \quad \frac{\text{is_iface } \Gamma \ I; \text{is_class } \Gamma \ C}{\Gamma \vdash \text{lface } I \preceq? \text{Class } C}$$

Typing rules Now we come to type-checking itself, which is expressed as a set of constraints on the types of expressions, relative to a type environment.

An *environment* consists of a global part, namely a program Γ , and a local part (written ‘ Λ ’), namely the types of the local variables including the current class, i.e. the type of **this**:

$$\begin{aligned} \text{lenv} &= (\text{ename}, \text{ty}) \ \text{table} \\ \text{env} &= \text{prog} \times \text{lenv} \end{aligned}$$

$$\begin{aligned} \text{prg} &:: \text{env} \Rightarrow \text{prog} \stackrel{\text{def}}{=} \lambda(\Gamma, \Lambda). \ \Gamma \\ \text{lcl} &:: \text{env} \Rightarrow \text{lenv} \stackrel{\text{def}}{=} \lambda(\Gamma, \Lambda). \ \Lambda \end{aligned}$$

The well-typedness of statements and the typing of expressions are defined inductively relative to an environment. The typing of expressions is unique, as can be shown easily by rule induction.

$$\begin{aligned} _ \vdash _ :: \diamond &:: \text{env} \Rightarrow \text{stmt} \Rightarrow \text{bool} \\ _ \vdash _ :: _ &:: \text{env} \Rightarrow \text{expr} \Rightarrow \text{ty} \Rightarrow \text{bool} \end{aligned}$$

The type-checking rules for most statements are standard:

$$\begin{aligned} &\frac{}{E \vdash \text{Skip} :: \diamond} \quad \frac{E \vdash e :: T}{E \vdash \text{Expr } e :: \diamond} \quad \frac{E \vdash c_1 :: \diamond; \ E \vdash c_2 :: \diamond}{E \vdash c_1; \ c_2 :: \diamond} \\ &\frac{E \vdash e :: \text{PrimT } \text{boolean}; \ E \vdash c_1 :: \diamond; \ E \vdash c_2 :: \diamond}{E \vdash \text{if}(e) \ c_1 \ \text{else} \ c_2 :: \diamond} \\ &\frac{E \vdash e :: \text{PrimT } \text{boolean}; \ E \vdash c :: \diamond}{E \vdash \text{while}(e) \ c :: \diamond} \quad \frac{E \vdash c_1 :: \diamond; \ E \vdash c_2 :: \diamond}{E \vdash c_1 \ \text{finally} \ c_2} \end{aligned}$$

Note the use of the widening relation in the following two rules to ensure that a value thrown or caught as an exception is indeed an exception object.

$$\begin{aligned} &\frac{E \vdash e :: \text{Class } tn; \ \text{prg } E \vdash \text{Class } tn \preceq \text{Class } (\text{SXcpt } \text{Throwable})}{E \vdash \text{throw } e :: \diamond} \\ &\frac{(\Gamma, \Lambda) \vdash c_1 :: \diamond; \ \Gamma \vdash \text{Class } tn \preceq \text{Class } (\text{SXcpt } \text{Throwable}); \ \Lambda \ vn = \text{None}; \ (\Gamma, \Lambda[vn \rightarrow \text{Class } tn]) \vdash c_2 :: \diamond}{(\Gamma, \Lambda) \vdash \text{try } c_1 \ \text{catch}(tn \ vn) \ c_2 :: \diamond} \end{aligned}$$

The `try _ catch _` statement is the only one that involves a change of the type environment, namely to include typing information for the exception parameter. The name of this parameter is required to be new in the local environment.

The typing rules for the first few of the expressions are straightforward, except for the confusing direction of the casting relation in the type cast rule:

$$\frac{\text{is_class (prg } E) C}{E \vdash \text{new } C::\text{Class } C} \quad \frac{\text{is_type (prg } E) T; E \vdash i::\text{PrimT int}}{E \vdash \text{new } T[i]::T}$$

$$\frac{E \vdash e::T; \text{prg } E \vdash T \preceq? T'}{E \vdash (T') e::T'} \quad \frac{\text{typeof } (\lambda a. \text{None}) x = \text{Some } T}{E \vdash \text{Lit } x::T}$$

$$\frac{E \vdash e::\text{RefT } T; \text{prg } E \vdash \text{RefT } T \preceq? \text{RefT } T'}{E \vdash e \text{ instanceof } T'::\text{PrimT boolean}}$$

The rule for `Lit` prohibits addresses as literal values, which is implemented by supplying `λa. None` as the “dynamic type” argument in the call of the function

`typeof :: (loc ⇒ ty option) ⇒ val ⇒ ty option`

`typeof dt Unit = Some (PrimT void)`
`typeof dt (Bool b) = Some (PrimT boolean)`
`typeof dt (Intg i) = Some (PrimT int)`
`typeof dt Null = Some (RefT NullT)`
`typeof dt (Addr a) = dt a`

This function is reused below with a more interesting value for the parameter `dt`, namely a function to compute the dynamic type of a reference.

The typings of all three assignment variants are quite similar, except that for local variables additionally an assignment to `this` is forbidden. In any case, as a generalization to the Java specification, the type of the assignment is determined by the right-hand (as opposed to the left-hand) side.

$$\frac{\text{lcl } E \text{ vn} = \text{Some } T; \text{is_type (prg } E) T}{E \vdash \text{vn}::T}$$

$$\frac{E \vdash \text{vn}::T; E \vdash v::T'; \text{prg } E \vdash T' \preceq T; \text{vn} \neq \text{this}}{E \vdash \text{vn}:=v::T'}$$

$$\frac{E \vdash e::\text{Class } C; \text{cfield (prg } E) C \text{ fn} = \text{Some (fd,fn)}}{E \vdash \{\text{fd}\}e.\text{fn}::\text{fn}}$$

$$\frac{E \vdash \{\text{fd}\}e.\text{fn}::T; E \vdash v::T'; \text{prg } E \vdash T' \preceq T}{E \vdash \{\text{fd}\}e.\text{fn}:=v::T'}$$

$$\frac{E \vdash a::T[]; E \vdash i::\text{PrimT int}}{E \vdash a[i]::T}$$

$$\frac{E \vdash a[i]::T; E \vdash v::T'; \text{prg } E \vdash T' \preceq T}{E \vdash a[i]:=v::T'}$$

$$\frac{E \vdash e::\text{RefT } T; E \vdash p::pT; \text{max_spec (prg } E) T (mn,pT) = \{((md,(pn,rT)),pT')\}}{E \vdash e.mn(\{pT'\}p)::rT}$$

The function $\text{cfield} :: \text{prog} \Rightarrow \text{tname} \Rightarrow (\text{ename}, \text{ref_ty} \times \text{field})\text{table}$, defined as $\text{cfield } \Gamma \ C \stackrel{\text{def}}{=} \text{table_of } ((\text{map } (\lambda((n,d),t). (n,(d,t))) (\text{fields } \Gamma \ C)))$, is a variant of `fields`. It implements a field lookup that is based on the field name alone in contrast to a combination of field name and defining class. Thus in the above typing rule for field access, equal field names hide each other, while at run-time all fields are accessible, using the defining class as an additional search key.

The type annotations $\{\dots\}$ in the above rules for field access and method call are used to implement static binding for fields and to resolve overloaded method names statically. Technically speaking, the typing rules serve as constraints on these annotations during type-checking, but one can also think of the annotations being filled with schematic variables that are instantiated with their correct values in the type-checking process, as is demonstrated in the example overleaf. The value of each annotation is uniquely determined by the value of a function in the premise of the field access and method call rule:

A field access $\{fd\}e.fn$ is annotated with the defining class of the field found when searching the class hierarchy for the name fn (using `cfield`), starting from the static type `Class C` of e . The annotation $\{fd\}$ will be used at run-time to access the field via the pair (fn,fd) .

A method call $e.mn(\{pT'\}p)$ is type-correct only if the function `max_spec` determining the set of “maximally specific” [GJS96, 15.11.2] methods for reference type T (as defined below) yields exactly one method entry. In this case, the method call is annotated by pT' , which is the argument type of the most specific method mn applicable according to the static types T of e and pT of p . Thus any static overloading of the method name mn has been resolved and the dynamic method lookup at run-time will be based on the signature (mn,pT') .

$$\begin{aligned} \text{max_spec} &:: \text{prog} \Rightarrow \text{ref_ty} \Rightarrow \text{sig} && \Rightarrow ((\text{ref_ty} \times \text{mhead}) \times \text{ty}) \ \text{set} \\ \text{appl_methds} &:: \text{prog} \Rightarrow \text{ref_ty} \Rightarrow \text{sig} && \Rightarrow ((\text{ref_ty} \times \text{mhead}) \times \text{ty}) \ \text{set} \\ \text{mheads} &:: \text{prog} \Rightarrow \text{ref_ty} \Rightarrow \text{sig} && \Rightarrow (\text{ref_ty} \times \text{mhead}) \ \text{set} \\ \text{more_spec} &:: \text{prog} \Rightarrow (\text{ref_ty} \times \text{mhead}) \times \text{ty} \Rightarrow (\text{ref_ty} \times \text{mhead}) \times \text{ty} \Rightarrow \text{bool} \end{aligned}$$

$$\begin{aligned} \text{max_spec } \Gamma \ T \ \text{sig} &\stackrel{\text{def}}{=} \{m \mid m \in \text{appl_methds } \Gamma \ T \ \text{sig} \wedge \\ &\quad (\forall m' \in \text{appl_methds } \Gamma \ T \ \text{sig}. \\ &\quad \quad \text{more_spec } \Gamma \ m' \ m \longrightarrow m' = m)\} \\ \text{appl_methds } \Gamma \ T \ (mn, pT) &\stackrel{\text{def}}{=} \{(m, pT') \mid m \in \text{mheads } \Gamma \ T \ (mn, pT') \wedge \\ &\quad \Gamma \vdash pT \preceq pT'\} \end{aligned}$$

$$\begin{aligned} \text{mheads } \Gamma \ \text{NullT} &= \lambda \text{sig}. \{\} \\ \text{mheads } \Gamma \ (\text{lfaceT } I) &= \text{imethds } \Gamma \ I \\ \text{mheads } \Gamma \ (\text{ClassT } C) &= \text{o2s} \circ \text{option_map } (\lambda(d,(h,b)).(d,h)) \circ \text{cmethd } \Gamma \ C \\ \text{mheads } \Gamma \ (\text{ArrayT } T) &= \lambda \text{sig}. \{\} \\ \text{more_spec } \Gamma \ ((md,mh),pT) \ ((m',mh'),pT') &\stackrel{\text{def}}{=} \Gamma \vdash \text{RefT } md \preceq \text{RefT } m' \wedge \\ &\quad \Gamma \vdash pT \preceq pT' \end{aligned}$$

where

$$\begin{aligned} \text{option_map} &:: (\alpha \Rightarrow \beta) \Rightarrow (\alpha \ \text{option} \Rightarrow \beta \ \text{option}) \\ \text{option_map } f &\stackrel{\text{def}}{=} \lambda y. \text{case } y \ \text{of } \text{None} \Rightarrow \text{None} \mid \text{Some } x \Rightarrow \text{Some } (f \ x) \end{aligned}$$

The well-typedness of our example code *test* is derived as given below. For formatting reasons, the derivation tree is cut at several positions, whereby the positions are marked with the labels of the cut subtrees. Irrelevant values in formulas are replaced by $_$. We use the following abbreviations:

$\Gamma = \text{tprg}$
 $\Lambda = \text{empty}[e \rightarrow \text{Class Base}]$
 $\text{SNP} = \text{SXcpt NullPointer}$

During the derivation, the schematic variable $?pT'$ is instantiated with *Class Base*, as a result of the function `max_spec`.

$$\begin{array}{c}
 \frac{\Lambda e = \text{Some } (\text{Class Base}) \quad \text{is_type } \Gamma (\text{Class Base})}{(\Gamma, \Lambda) \vdash e :: \text{Class Base}} \quad \frac{\text{is_class } \Gamma \text{ Ext}}{(\Gamma, \Lambda) \vdash \text{new Ext} :: \text{Class Ext}} \quad \Gamma \vdash \text{Class Ext} \preceq \text{Class Base} \\
 \frac{e \neq \text{this} \quad (\Gamma, \Lambda) \vdash \text{new Ext} :: \text{Class Ext} \quad \Gamma \vdash \text{Class Ext} \preceq \text{Class Base}}{(\Gamma, \Lambda) \vdash \text{Expr}(e = \text{new Ext}) :: \diamond} \quad (\text{LAss}) \\
 \Lambda e = \text{Some } (\text{Class Base}) \quad \text{is_type } \Gamma (\text{Class Base}) \quad \text{typeof } (\lambda a. \text{None}) \text{ Null} = \text{Some NT} \quad \text{max_spec } \Gamma (\text{ClassT Base}) (\text{foo}, \text{NT}) = \\
 \frac{(\Gamma, \Lambda) \vdash e :: \text{Class Base} \quad (\Gamma, \Lambda) \vdash \text{Lit Null} :: \text{NT} \quad \{((-, -, \text{Class Base}), ?pT')\}}{(\Gamma, \Lambda) \vdash \text{Expr}(e.\text{foo}\{\text{?pT'}\} \text{Lit Null}) :: \text{Class Base}} \quad (\text{Call}) \\
 \frac{(\Gamma, \Lambda) \vdash \text{Expr}(e.\text{foo}\{\text{?pT'}\} \text{Lit Null}) :: \text{Class Base}}{(\Gamma, \Lambda) \vdash \text{Expr}(e.\text{foo}\{\text{?pT'}\} \text{Lit Null}) :: \diamond} \\
 \frac{\Lambda[x \rightarrow \text{Class SNP}] \quad x = \text{Some } (\text{Class SNP}) \quad \text{is_type } \Gamma (\text{Class SNP})}{(\Gamma, \Lambda[x \rightarrow \text{Class SNP}]) \vdash x :: \text{Class SNP}} \quad \Gamma \vdash \text{Class SNP} \preceq \text{Class } (\text{SXcpt Throwable}) \quad (\text{Throw}) \\
 \frac{(\Gamma, \Lambda[x \rightarrow \text{Class SNP}]) \vdash \text{throw } x :: \diamond}{(\Gamma, \Lambda) \vdash \text{try Expr}(e.\text{foo}\{\text{?pT'}\} \text{Lit Null}) \text{ catch } (\text{SNP } x) (\text{throw } x) :: \diamond} \quad (\text{LAss}) \\
 \frac{(\text{Call}) \quad \Gamma \vdash \text{Class SNP} \preceq \text{Class } (\text{SXcpt Throwable}) \quad \Lambda x = \text{None} \quad (\text{Throw})}{(\Gamma, \Lambda) \vdash \text{try Expr}(e.\text{foo}\{\text{?pT'}\} \text{Lit Null}) \text{ catch } (\text{SNP } x) (\text{throw } x) :: \diamond} \quad (\text{LAss}) \\
 \frac{(\Gamma, \Lambda) \vdash \text{Expr}(e = \text{new Ext}); \text{try Expr}(e.\text{foo}\{\text{?pT'}\} \text{Lit Null}) \text{ catch } (\text{SNP } x) (\text{throw } x) :: \diamond}{(\Gamma, \Lambda) \vdash \text{Expr}(e = \text{new Ext}); \text{try Expr}(e.\text{foo}\{\text{?pT'}\} \text{Lit Null}) \text{ catch } (\text{SNP } x) (\text{throw } x) :: \diamond}
 \end{array}$$

4.3 Well-formedness

A program must satisfy a number of well-formedness conditions concerning global properties of all declarations. The conditions are expressed as predicates on field, method, interface, class, and whole program declarations.

$$\begin{aligned}
\text{wf_fdecl} &:: \text{prog} \Rightarrow && \text{fdecl} \Rightarrow \text{bool} \\
\text{wf_mhead} &:: \text{prog} \Rightarrow \text{sig} \quad \times \text{mhead} \Rightarrow \text{bool} \\
\text{wf_mdecl} &:: \text{prog} \Rightarrow \text{tname} \Rightarrow \text{mdecl} \Rightarrow \text{bool} \\
\text{wf_idecl} &:: \text{prog} \Rightarrow && \text{idecl} \Rightarrow \text{bool} \\
\text{wf_cdecl} &:: \text{prog} \Rightarrow && \text{cdecl} \Rightarrow \text{bool} \\
\text{wf_prog} &:: \text{prog} \Rightarrow && \text{bool}
\end{aligned}$$

A field declaration is well-formed iff its type exists:

$$\text{wf_fdecl } \Gamma (fn, ft) \stackrel{\text{def}}{=} \text{is_type } \Gamma ft$$

A method declaration is well-formed only if its argument and result types are defined and the name of the parameter is not `this`. Additionally, if the declaration appears in a class, the names of the local variables must be unique and may not contain the special name `this` nor hide the parameter, all types of the local variables must exist, the method body has to be well-typed (in the static context of its parameter type and the current class), and its result expression must have a type that widens to the result type:

$$\begin{aligned}
\text{wf_mhead } \Gamma ((mn, pT), (pn, rT)) &\stackrel{\text{def}}{=} \text{is_type } \Gamma pT \wedge \text{is_type } \Gamma rT \wedge pn \neq \text{this} \\
\text{wf_mdecl } \Gamma C ((mn, pT), (pn, rT), lvars, blk, res) &\stackrel{\text{def}}{=} \\
&\text{let } ltab = \text{table_of } lvars; E = (\Gamma, ltab[\text{this} \mapsto \text{Class } C][pn \mapsto pT]) \\
&\text{in wf_mhead } \Gamma ((mn, pT), (pn, rT)) \wedge \\
&\text{unique } lvars \wedge ltab \text{ this} = \text{None} \wedge ltab pn = \text{None} \wedge \\
&(\forall (vn, T) \in \text{set } lvars. \text{is_type } \Gamma T) \wedge \\
&E \vdash blk :: \diamond \wedge \exists T. E \vdash res :: T \wedge \Gamma \vdash T \preceq rT
\end{aligned}$$

Even more complex conditions are required for well-formed interface and class declarations. The name of a well-formed interface declaration is not a class name. All superinterfaces exist and are not subinterfaces at the same time. All methods newly declared in the interface are named uniquely and are well-formed. Furthermore, any method overriding a set of methods defined in some superinterfaces has a result type that widens to all their result types:

$$\begin{aligned}
\text{wf_idecl } \Gamma (I, (is, ms)) &\stackrel{\text{def}}{=} \neg \text{is_class } \Gamma I \wedge \\
&(\forall J \in \text{set } is. \text{is_iface } \Gamma J \wedge \neg \Gamma \vdash J \prec_i I) \wedge \\
&\text{unique } ms \wedge (\forall m \in \text{set } ms. \text{wf_mhead } \Gamma m \wedge \\
&\text{let } mtab = \text{Un_tables } ((\lambda J. \text{imethds } \Gamma J) \text{ ``set } is) \text{ in} \\
&(\text{o2s } \circ \text{table_of } ms) \text{ hidings } mtab \text{ entails} \\
&(\lambda (pn, rT) (m, (pn', rT')). \Gamma \vdash rT \preceq rT')
\end{aligned}$$

Similarly, the name of a well-formed class declaration is not an interface name. All implemented interfaces exist, and for any method of such an interface, the class provides an implementing method with a possibly narrower return type. All fields and methods newly declared in the class are named uniquely and are well-formed. If the class is not `Object`, it refers to an existing superclass, which is not a subclass of the current class. Furthermore, any method overriding a method of the superclass has a compatible result type:

$$\begin{aligned}
\text{wf_cdecl } \Gamma \ (C, (sc, si, fs, ms)) &\stackrel{\text{def}}{=} \neg \text{is_iface } \Gamma \ C \wedge \\
&(\forall I \in \text{set } si. \text{is_iface } \Gamma \ I \wedge \\
&\quad \forall s. \forall (m_1, (pn_1, rT_1)) \in \text{imethds } \Gamma \ I \ s. \\
&\quad \exists m_2 \ pn_2 \ rT_2 \ b. \text{cmethd } \Gamma \ C \ s = \text{Some } (m_2, (pn_2, rT_2), b) \wedge \\
&\quad \Gamma \vdash rT_2 \preceq rT_1) \wedge \\
&\text{unique } fs \wedge (\forall f \in \text{set } fs. \text{wf_fdecl } \Gamma \ f) \wedge \\
&\text{unique } ms \wedge (\forall m \in \text{set } ms. \text{wf_mdecl } \Gamma \ C \ m) \wedge \\
&(\text{case } sc \text{ of None } \Rightarrow C = \text{Object} \\
&\quad | \text{Some } D \Rightarrow \text{is_class } \Gamma \ D \wedge \neg \Gamma \vdash D \prec_c C \wedge \\
&\quad \text{table_of } ms \text{ hiding cmethd } \Gamma \ D \text{ entails} \\
&\quad (\lambda((pn_1, rT_1), b) (m, ((pn_2, rT_2), b')). \Gamma \vdash rT_1 \preceq rT_2)
\end{aligned}$$

Finally, all interfaces and classes declared in a well-formed program are named uniquely and are in turn well-formed. For uniformity, this includes the predefined class declarations of `Object` and the (flat) hierarchy of system exceptions.

$$\begin{aligned}
\text{ObjectC} &\stackrel{\text{def}}{=} (\text{Object} \ , \ (\text{None} \ , \ [], [], [])) \\
\text{SXcptC } xn &\stackrel{\text{def}}{=} \text{let } sc = \text{if } xn = \text{Throwable} \text{ then Object else SXcpt Throwable in} \\
&\quad (\text{SXcpt } xn, (\text{Some } sc, [], [], []))
\end{aligned}$$

$$\begin{aligned}
\text{wf_prog } \Gamma &\stackrel{\text{def}}{=} \text{let } is = \text{set } (\text{fst } \Gamma); \ cs = \text{set } (\text{snd } \Gamma) \\
&\text{in ObjectC } \in \ cs \wedge \forall xn. \text{SXcptC } xn \in \ cs \wedge \\
&\quad \text{unique } (\text{fst } \Gamma) \wedge \forall i \in is. \text{wf_idecl } \Gamma \ i) \wedge \\
&\quad \text{unique } (\text{snd } \Gamma) \wedge \forall c \in cs. \text{wf_cdecl } \Gamma \ c)
\end{aligned}$$

Our example program `tprg` is well-formed. Here is a heavily abstracted derivation tree of our proof of this fact.

$$\begin{array}{c}
\text{wf_mdecl } tprg \ \text{Base} \ ((foo, \text{Class } \text{Base}), \\
(x, \text{Class } \text{Base}), [], \text{Skip}, x) \quad \neg(tprg \vdash \text{Object} \prec_c \text{Base}) \\
\hline
\text{wf_cdecl } tprg \ \text{BaseC} \\
\\
\text{wf_mdecl } tprg \ \text{Ext} \ ((foo, \text{Class } \text{Base}), \\
(x, \text{Class } \text{Ext}), [], \text{Expr } (\{\text{ClassT } \text{Ext}\}(\text{Class } \text{Ext}) \\
x. \text{vee} := \text{Lit } (\text{Intg } 1)), \text{Lit } \text{Null}) \quad \neg(tprg \vdash \text{Base} \prec_c \text{Ext}) \\
\hline
\text{wf_cdecl } tprg \ \text{BaseC} \\
\hline
\text{wf_cdecl } tprg \ \text{BaseC} \ \text{wf_cdecl } tprg \ \text{ExtC} \ \text{Base} \neq \text{Ext} \\
\hline
\text{wf_tprg } tprg
\end{array}$$

4.4 Operational semantics

We formalize the semantics of Java in operational style with evaluation rules. This is the natural choice since the language specification itself is given in an operational evaluation-oriented style, which allows for a direct formalization and its straightforward validation. Furthermore, a denotational semantics would require much more difficult mathematical tools, and an axiomatic semantics would be problematic to validate and to use for reasoning on the language as a whole. We prefer an evaluation semantics to a transition semantics in order to obtain a concise description, because we consider a transition semantics less readable and rather low-level, which in particular holds for a formulation as an Abstract State Machine (ASM) like in [BS98].

In this section, we describe the notions of a state and its components and give the evaluation rules for statements and expressions.

State A *state* consists of an optional exception (of type *xcpt*), a heap, and a current invocation frame, which is the values of the local variables (including method and exception parameters and the `this` pointer):

$$\begin{aligned} \text{state} &= (\text{xcpt})\text{option} \times \text{st} \\ \text{st} &= \text{heap} \times \text{locals} \\ \text{heap} &:: \text{st} \Rightarrow \text{heap} \stackrel{\text{def}}{=} \lambda(h,l). h \\ \text{locals} &:: \text{st} \Rightarrow \text{locals} \stackrel{\text{def}}{=} \lambda(h,l). l \end{aligned}$$

Remember that tuples associative to the right, so if for some state σ we have an equation like $\sigma = (x, \sigma')$, then x is the (optional) exception component alone, while the second projection σ' of the state has (tuple) type *st*, i.e. represents a “small” state excluding the exception entry.

An *exception* is a reference to an instance of some exception class, which is a subclass of `Throwable`. Normally, when an exception is thrown, a fresh exception object is allocated and its location returned to represent the exception. But in the case of system exceptions, we defer their allocation (and just record their names) until an enclosing `catch` block references it. This helps to avoid the subtleties of (conditional) side effects on the heap and out-of-memory conditions. Thus we model exceptions as follows.

$$\begin{aligned} \text{xcpt} &= \text{XcptLoc } \text{loc} \\ &\quad | \text{SysXcpt } \text{xname} \end{aligned}$$

A *heap* maps locations to objects, while local variables map names to values:

$$\begin{aligned} \text{heap} &= (\text{loc} \quad , \text{obj})\text{table} \\ \text{locals} &= (\text{ename}, \text{val})\text{table} \end{aligned}$$

In our model there is no need to explicitly maintain a stack of invocation frames containing local variables and return addresses for method calls. In this way we also abstract over the finiteness of stack space. On the other hand, we explicitly model the possibility of memory allocation on the heap to fail if there is no free location (i.e. some a with $(\text{heap } \sigma) a = \text{None}$) available. Memory allocation is loosely, yet deterministically, defined by the function

$\text{init_vars} \quad :: (\alpha \times \text{ty})\text{list} \Rightarrow (\alpha, \text{val})\text{table}$
 $\text{init_Obj} \quad :: \text{prog} \Rightarrow \text{tname} \Rightarrow \text{obj}$
 $\text{init_Arr} \quad :: \text{ty} \Rightarrow \text{int} \Rightarrow \text{obj}$
 $\text{init_vars} \quad \stackrel{\text{def}}{=} \text{table_of} \circ \text{map} (\lambda(n, T). (n, \text{default_val } T))$
 $\text{init_Obj } \Gamma C \stackrel{\text{def}}{=} \text{Obj } C (\text{init_vars } (\text{fields } \Gamma C))$
 $\text{init_Arr } T i \stackrel{\text{def}}{=} \text{Arr } T (\lambda j. \text{ if } 0 \leq j \wedge j < i \text{ then Some } (\text{default_val } T) \text{ else None})$
 $\text{raise_if} \quad :: \text{bool} \Rightarrow \text{xname} \Rightarrow (\text{xcpt})\text{option} \Rightarrow (\text{xcpt})\text{option}$
 $\text{np} \quad \quad :: \text{val} \Rightarrow (\text{xcpt})\text{option} \Rightarrow (\text{xcpt})\text{option}$
 $\text{raise_if } c \text{ xn } xo \stackrel{\text{def}}{=} \text{ if } c \wedge (xo = \text{None}) \text{ then Some } (\text{SysXcpt } xn) \text{ else } xo$
 $\text{np } v \quad \quad \stackrel{\text{def}}{=} \text{raise_if } (v = \text{Null}) \text{ NullPointer}$

The definition of `raise_if` deserves a comment: `raise_if c xn xo` either propagates an already thrown exception `xo` or raises the system exception `xn` if `c` is true. As an application, `np v` checks for a null pointer access through the value `v` and throws a `NullPointer` exception in this case, but any other exception that has already occurred takes precedence.

Evaluation rule format Internally, the evaluation rules are given as mutually inductive sets of tuples. These sets define relations, which we present as predicates of the following form.

- $\Gamma \vdash \sigma - c \rightarrow \sigma' :: \text{prog} \Rightarrow \text{state} \Rightarrow \text{stmt} \Rightarrow \text{state} \Rightarrow \text{bool}$
means that the execution of statement `c` transforms state `σ` into `σ'`.
- $\Gamma \vdash \sigma - e \triangleright v \rightarrow \sigma' :: \text{prog} \Rightarrow \text{state} \Rightarrow \text{expr} \Rightarrow \text{val} \Rightarrow \text{state} \Rightarrow \text{bool}$
means that expression `e` evaluates to `v`, transforming `σ` into `σ'`.

Although defined as relations (for technical reasons), the semantics given below can be shown to be functional, i.e. deterministic.

Strictly speaking it is not necessary to include an exception in the start state of a computation. Similarly, an expression needs only return either a value or an exception, but not both. However, the symmetry achieved by our slightly redundant model simplifies the rules considerably. In particular, we can avoid case distinctions on whether exceptions occur in intermediate states, which would cause the rules to be split. Suppose for example that $\Gamma \vdash \sigma - c \rightarrow \sigma'$ had the signature $\text{prog} \Rightarrow \text{st} \Rightarrow \text{stmt} \Rightarrow \text{state} \Rightarrow \text{bool}$, i.e. all rules assume that there is no exception in the start state. Then the rule(s) for sequential composition would look like

$$\frac{\Gamma \vdash \sigma_0 - c_1 \rightarrow (\text{None}, \sigma_1); \Gamma \vdash \sigma_1 - c_2 \rightarrow \sigma_2}{\Gamma \vdash \sigma_0 - c_1; c_2 \rightarrow \sigma_2}$$

$$\frac{\Gamma \vdash \sigma_0 - c_1 \rightarrow (\text{Some } xs, \sigma_1)}{\Gamma \vdash \sigma_0 - c_1; c_2 \rightarrow (\text{Some } xs, \sigma_1)}$$

As a consequence of the design decisions just mentioned, there is exactly one rule for each syntactic construct. Additionally there are general rules defining that exceptions simply propagate when a series of statements is executed or a series of expressions is evaluated:

$$\frac{}{\Gamma \vdash (\text{Some } xc, \sigma) - c \rightarrow (\text{Some } xc, \sigma)}$$

$$\frac{}{\Gamma \vdash (\text{Some } xc, \sigma) - e \triangleright \text{arbitrary} \rightarrow (\text{Some } xc, \sigma)}$$

All other rules can assume that in their concerning initial state no exception has been thrown. For such states, we define the abbreviation Norm σ , which stands for (None, σ) .

Execution of statements The rules for the statements not explicitly involving exceptions are obvious:

$$\frac{}{\Gamma \vdash \text{Norm } \sigma - \text{Skip} \rightarrow \text{Norm } \sigma} \quad \frac{\Gamma \vdash \text{Norm } \sigma_0 - c_1 \rightarrow \sigma_1; \Gamma \vdash \sigma_1 - c_2 \rightarrow \sigma_2}{\Gamma \vdash \text{Norm } \sigma_0 - c_1; c_2 \rightarrow \sigma_2}$$

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 - e \triangleright v \rightarrow \sigma_1}{\Gamma \vdash \text{Norm } \sigma_0 - \text{Expr } e \rightarrow \sigma_1} \quad \frac{\Gamma \vdash \text{Norm } \sigma_0 - e \triangleright v \rightarrow \sigma_1; \Gamma \vdash \sigma_1 - \text{if the_Bool } v \text{ then } c_1 \text{ else } c_2 \rightarrow \sigma_2}{\Gamma \vdash \text{Norm } \sigma_0 - \text{if}(e) c_1 \text{ else } c_2 \rightarrow \sigma_2}$$

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 - \text{if}(e) (c; \text{while}(e) c) \text{ else Skip} \rightarrow \sigma_1}{\Gamma \vdash \text{Norm } \sigma_0 - \text{while}(e) c \rightarrow \sigma_1}$$

If no other exceptions have occurred while evaluating its argument and testing for a null reference (using `np`), the `throw` statement copies the evaluated location into the exception component of the state:

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 - e \triangleright a' \rightarrow (x_1, \sigma_1); x_1' = \text{np } a' x_1; x_1'' = (\text{if } x_1' = \text{None} \text{ then } (\text{Some } (\text{XcptLoc } (\text{the_Addr } a')))) \text{ else } x_1')}{\Gamma \vdash \text{Norm } \sigma_0 - \text{throw } e \rightarrow (x_1'', \sigma_1)}$$

For the semantics of the `try - catch -` statement we have to distinguish whether some exception is thrown and then caught by the `catch` clause or not. In the first case, i.e. there is an exception of appropriate dynamic type to be handled, the `catch` clause is executed with its exception parameter set to the caught exception. In the second case the `catch` clause is skipped. Because of technical limitations of the inductive definition package of Isabelle/HOL, even in this case we have to provide an occurrence of the execution relation, which in effect simply sets σ_2 to (x_1', σ_1') .

$$\frac{\begin{array}{l} \Gamma \vdash \text{Norm } \sigma_0 - c_1 \rightarrow \sigma_1; \Gamma \vdash \sigma_1 - \text{salloc} \rightarrow (x_1', \sigma_1'); \\ \text{case } x_1' \text{ of None} \Rightarrow \sigma_1'' = (x_1', \sigma_1') \wedge c_2' = \text{Skip} \\ \quad | \text{Some } xc \Rightarrow \text{let } a = \text{Addr } (\text{the_XcptLoc } xc) \text{ in} \\ \quad \quad \text{if } \Gamma, \sigma_1' \vdash a \text{ fits Class } tn \\ \quad \quad \text{then } \sigma_1'' = \text{Norm } (\text{lupd}[vn \mapsto a] \sigma_1') \wedge c_2' = c_2 \\ \quad \quad \text{else } \sigma_1'' = (x_1', \sigma_1') \wedge c_2' = \text{Skip}; \end{array}}{\Gamma \vdash \sigma_1'' - c_2' \rightarrow \sigma_2} \\ \hline \Gamma \vdash \text{Norm } \sigma_0 - (\text{try } c_1 \text{ catch}(tn \text{ } vn) c_2) \rightarrow \sigma_2$$

On the one hand, the exception parameter of the `catch` clause must represent the exception thrown in the `try` block by a reference to its exception object. As on the other hand we defer the allocation of system exceptions when evaluating expressions, we have to ensure that even for such exceptions a suitable exception object is allocated on the heap of σ_1' , replacing the `SysXcpt` entry by an `XcptLoc` entry in x_1' . This is achieved by the auxiliary relation $\Gamma \vdash \sigma \text{--salloc--} \sigma' :: \text{prog} \Rightarrow \text{state} \Rightarrow \text{state} \Rightarrow \text{bool}$. If no system exception has been thrown, the relation behaves like the identity on the state, and otherwise allocates an exception object and modifies the state accordingly. Note that this allocation step is impossible — and therefore program execution halts — if there is no free address left.

$$\frac{\Gamma \vdash \text{Norm } \sigma \text{--salloc--} \text{Norm } \sigma}{\Gamma \vdash (\text{Some } (\text{XcptLoc } a), \sigma) \text{--salloc--} (\text{Some } (\text{XcptLoc } a), \sigma)}$$

$$\frac{\text{new_Addr } (\text{heap } \sigma) = \text{Some } (a, x); \quad xobj = \text{init_Obj } \Gamma (\text{SXcpt } (\text{if } x = \text{None} \text{ then } xn \text{ else } \text{OutOfMemory}))}{\Gamma \vdash (\text{Some } (\text{SysXcpt } xn), \sigma) \text{--salloc--} (\text{Some } (\text{XcptLoc } a), \text{hupd}[a \mapsto xobj] \sigma)}$$

The `finally` statement is similar to the sequential composition, but executes its second clause regardless whether an exception has been thrown in its first clause or not. If an exception occurs in either clause, it is (re-)raised after the statement, and if both parts throw an exception, the first one takes precedence.

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 \text{--}c_1 \text{--} (x_1, \sigma_1); \quad \Gamma \vdash \text{Norm } \sigma_1 \text{--}c_2 \text{--} (x_2, \sigma_2); \quad x_2' = (\text{if } x_1 \neq \text{None} \wedge x_2 = \text{None} \text{ then } x_1 \text{ else } x_2)}{\Gamma \vdash \text{Norm } \sigma_0 \text{--}(c_1 \text{ finally } c_2)\text{--} (x_2', \sigma_2)}$$

Evaluation of expressions In contrast to the statement rules, almost all evaluation rules for expressions deserve some comments.

Creating a new class instance means picking a free address a and updating the heap at that address with an object, the fields of which are initialized with default values according to their types. Note that the rule is not applicable — and therefore execution halts — if `new_Addr` fails.

$$\frac{\text{new_Addr } (\text{heap } \sigma) = \text{Some } (a, x)}{\Gamma \vdash \text{Norm } \sigma \text{--new } C \triangleright \text{Addr } a \rightarrow \text{x_case } x (\text{hupd}[a \mapsto \text{init_Obj } \Gamma C] \sigma) \sigma}$$

The same applies for the creation of a new array, where additionally an exception is raised if the length of the array is negative:

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 \text{--}e \triangleright i' \text{--} (x_1, \sigma_1); \quad i = \text{the_Intg } i'; \quad \text{new_Addr } (\text{heap } \sigma_1) = \text{Some } (a, x); \quad x_1' = \text{raise_if } (i < 0) \text{ NegArrSize } (\text{if } x_1 = \text{None} \text{ then } x \text{ else } x_1)}{\Gamma \vdash \text{Norm } \sigma_0 \text{--new } T[e] \triangleright \text{Addr } a \rightarrow \text{x_case } x_1' (\text{hupd}[a \mapsto \text{init_Arr } T i] \sigma_1) \sigma_1}$$

A type cast merely returns its argument value, but raises an exception if the dynamic type happens to be unsuitable:

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 - e \triangleright v \rightarrow (x_1, \sigma_1); \quad x_1' = \text{raise_if}(\neg \Gamma, \sigma_1 \vdash v \text{ fits } T) \text{ ClassCast } x_1}{\Gamma \vdash \text{Norm } \sigma_0 - (T) e \triangleright v \rightarrow (x_1', \sigma_1)}$$

The type comparison operator checks if the type of its argument is assignable to the given reference type:

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 - e \triangleright v \rightarrow \sigma_1; \quad b = (v \neq \text{Null} \wedge \Gamma, \text{snd } \sigma_1 \vdash v \text{ fits RefT } T)}{\Gamma \vdash \text{Norm } \sigma_0 - e \text{ instanceof } T \triangleright \text{Bool } b \rightarrow \sigma_1}$$

The result of a literal expression is simply the given value:

$$\overline{\Gamma \vdash \text{Norm } \sigma - \text{Lit } v \triangleright v \rightarrow \text{Norm } \sigma}$$

An access to a local variable (or the `this` pointer) reads from the local state component:

$$\overline{\Gamma \vdash \text{Norm } \sigma - vn \triangleright \text{the } (\text{locals } \sigma \text{ } vn) \rightarrow \text{Norm } \sigma}$$

An assignment to a local variable updates the state, unless the evaluation of the subexpression raises an exception:

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 - e \triangleright v \rightarrow (x, \sigma_1); \quad \sigma_1' = (\text{if } x = \text{None} \text{ then } \text{lupd}[vn \mapsto v] \sigma_1 \text{ else } \sigma_1)}{\Gamma \vdash \text{Norm } \sigma_0 - vn := e \triangleright v \rightarrow (x, \sigma_1')}$$

A field access reads from a field of the given object, taking into account the type annotation which yields the defining class of the field as determined statically. It also checks for null pointer access.

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 - e \triangleright a' \rightarrow (x_1, \sigma_1); \quad v = \text{the } (\text{snd } (\text{the_Obj } (\text{heap } \sigma_1 (\text{the_Addr } a')))) (fn, T)}{\Gamma \vdash \text{Norm } \sigma_0 - \{T\} e . fn \triangleright v \rightarrow (\text{np } a' \ x_1, \sigma_1)}$$

A field assignment acts accordingly:

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 - e_1 \triangleright a' \rightarrow (x_1, \sigma_1); \quad a = \text{the_Addr } a'; \quad \Gamma \vdash (\text{np } a' \ x_1, \sigma_1) - e_2 \triangleright v \rightarrow (x_2, \sigma_2); \quad (c, fs) = \text{the_Obj } (\text{heap } \sigma_2 \ a); \quad obj = \text{Obj } c \ (fs[(fn, T) := v])}{\Gamma \vdash \text{Norm } \sigma_0 - (\{T\} e_1 . fn := e_2) \triangleright v \rightarrow \text{x_case } x_2 \ (\text{hupd}[a \mapsto obj] \sigma_2) \ \sigma_2}$$

An array access reads a component from the given array, but raises an exception if the index is invalid:

$$\frac{\Gamma \vdash \text{Norm } \sigma_0 - e_1 \triangleright a' \rightarrow \sigma_1; \quad \Gamma \vdash \sigma_1 - e_2 \triangleright i' \rightarrow (x_2, \sigma_2); \quad vo = \text{snd } (\text{the_Arr } (\text{heap } \sigma_2 (\text{the_Addr } a')) (\text{the_Intg } i')); \quad x_2' = \text{raise_if } (vo = \text{None}) \ \text{IndOutOfBounds } (\text{np } a' \ x_2)}{\Gamma \vdash \text{Norm } \sigma_0 - e_1[e_2] \triangleright \text{the } vo \rightarrow (x_2', \sigma_2)}$$

Similarly, an array assignment updates the appropriate component, but first has to check the type of the value to be assigned. Note one subtle difference to field assignment: null pointer access is checked after evaluating the right-hand side, whereas in field assignment the check occurs immediately after calculating the reference.

$$\begin{array}{c}
\Gamma \vdash \text{Norm } \sigma_0 - e_1 \triangleright a' \rightarrow \sigma_1; \quad a = \text{the_Addr } a'; \\
\Gamma \vdash \sigma_1 - e_2 \triangleright i' \rightarrow \sigma_2; \quad i = \text{the_Intg } i'; \\
\Gamma \vdash \sigma_2 - e_3 \triangleright v \rightarrow (x_3, \sigma_3); \\
(T, cs) = \text{the_Arr } (\text{heap } \sigma_3 \ a); \quad \text{obj} = \text{Arr } T \ (cs[i \mapsto v]); \\
x_3' = \text{raise_if } (\neg \Gamma, \sigma_3 \vdash v \text{ fits } T) \ \text{ArrStore } (\\
\text{raise_if } (cs \ i = \text{None}) \ \text{IndOutOfBounds } (\text{np } a' \ x_3)) \\
\hline
\Gamma \vdash \text{Norm } \sigma_0 - (e_1[e_2] := e_3) \triangleright v \rightarrow \text{x_case } x_3' \ (\text{hupd}[a \mapsto \text{obj}]\sigma_3) \ \sigma_3
\end{array}$$

The most complex rule is the one for method invocation: after evaluating e to the target location a' and p to the parameter value pv , the block blk and the result expression res of method mn with argument type T are extracted from the program Γ (using the dynamic type $\text{dyn}T$ of the object stored at a'). For simplicity, we require local variables to be initialized with default values, as the expensive rules for “definite assignment” [GJS96, Ch. 16] merely enable the run-time optimization that variables need not be initialized before being explicitly assigned to. After executing blk and res in the new invocation frame built from the local variables, the parameter pv and a' as the value of `this`, the old invocation frame is restored and the result value v returned:

$$\begin{array}{c}
\Gamma \vdash \text{Norm } \sigma_0 - e \triangleright a' \rightarrow \sigma_1; \\
\Gamma \vdash \sigma_1 - p \triangleright pv \rightarrow (x_2, \sigma_2); \\
\text{dyn}T = \text{fst } (\text{the_Obj } (\text{heap } \sigma_2 \ (\text{the_Addr } a'))); \\
(md, (pn, rT), lvars, blk, res) = \text{the } (\text{cmethd } \Gamma \ \text{dyn}T \ (mn, pT)); \\
\Gamma \vdash (\text{np } a' \ x_2, (\text{heap } \sigma_2, \text{init_vars } lvars[\text{this} \mapsto a'] [pn \mapsto pv])) - blk \rightarrow \sigma_3; \\
\Gamma \vdash \sigma_3 - res \triangleright v \rightarrow (x_4, \sigma_4) \\
\hline
\Gamma \vdash \text{Norm } \sigma_0 - (e.mn(\{pT\}p)) \triangleright v \rightarrow (x_4, (\text{heap } \sigma_4, \text{locals } \sigma_2))
\end{array}$$

Note that all rules are defined carefully in order to be applicable even in not type-correct situations. For example, in any context where a value v is expected to be an address, we do not use a premise like $v = \text{Addr } a$ as this will disable the rule if v happens to be, for example, a null pointer or a Boolean value. Instead, we use an expression like $a = \text{the_Addr } v$, which will yield an arbitrary value if v is not an address, yet will leave the rule applicable. In such cases we could not prove anything useful about a , but during the type soundness proof itself it emerges that for well-formed programs (and statically well-typed statements and expressions) such situations cannot occur. A “defensive” evaluation throwing some artificial exception in case of type mismatches, which would require additional overhead, is therefore not necessary.

Below we give a derivation for the execution of our example code *test*, under the assumptions $\text{new_Addr empty} = \text{Some } (a, \text{None})$ and $\forall \text{obj. new_Addr } (\text{empty}[a \mapsto \text{obj}]) = \text{Some } (b, \text{None})$, which guarantee that there are at least two free locations on the heap.

Given only the start state σ_0 , all other states are computed during the derivation, which results in the sequence

We use the following abbreviations:

$\Gamma = \text{tprg}$

$\text{NP} = \text{NullPointer}$

$\text{blk} = \text{Expr}(\{\text{ClassT } \text{Ext}\}(\text{Class } \text{Ext})x.\text{vee}:=\text{Lit } (\text{Intg } 1))$

$\text{obj1} = \text{Obj } \text{Ext } (\text{empty}[(\text{vee}, \text{ClassT } \text{Base}) \mapsto \text{Bool } \text{False}]$
 $\quad [(\text{vee}, \text{ClassT } \text{Ext}) \mapsto \text{Intg } 0])$

$\text{obj2} = \text{Obj } (\text{SXcpt } \text{NP}) \text{ empty}$

$h = \text{empty}[a \mapsto \text{obj1}]$

$l = \text{empty}[e \mapsto \text{Addr } a]$

$\sigma_0 = \text{Norm } (\text{empty}, \text{empty})$

$\sigma_1 = \text{Norm } (h, l)$

$\sigma_2 = \text{Norm } (h, \text{empty}[\text{this} \mapsto \text{Addr } a][x \mapsto \text{Null}])$

$\sigma_3 = (\text{Some } (\text{SysXcpt } \text{NP}), h, \text{empty}[\text{this} \mapsto \text{Addr } a][x \mapsto \text{Null}])$

$\sigma_4 = (\text{Some } (\text{SysXcpt } \text{NP}), h, l)$

$\sigma_5 = (\text{Some } (\text{XcptLoc } b), h[b \mapsto \text{obj2}], l)$

$\sigma_6 = \text{Norm } (h[b \mapsto \text{obj2}], [x \mapsto \text{Addr } b])$

$\sigma_7 = (\text{Some } (\text{XcptLoc } b), h[b \mapsto \text{obj2}], [x \mapsto \text{Addr } b])$

$\frac{\Gamma \vdash \sigma_2 -x \triangleright \text{Null} \rightarrow \sigma_2 \quad \Gamma, \text{snd } \sigma_2 \vdash \text{Null fits Class } \text{Ext}}{\Gamma \vdash \sigma_2 -(\text{Class } \text{Ext})x \triangleright \text{Null} \rightarrow \sigma_2}$

$\frac{\Gamma \vdash \sigma_2 -(\text{Class } \text{Ext})x \triangleright \text{Null} \rightarrow \sigma_2 \quad \sigma_3 = (\text{Some } (\text{SysXcpt } \text{NP}), \text{snd } \sigma_2) \quad \Gamma \vdash \sigma_3 -\text{Lit } (\text{Intg } 1) \triangleright \rightarrow \sigma_3}{\Gamma \vdash \sigma_2 -(\{\text{ClassT } \text{Ext}\}(\text{Class } \text{Ext})x.\text{vee}:=\text{Lit } (\text{Intg } 1)) \triangleright \rightarrow \sigma_3} (\text{Blk})$

$\frac{\Gamma \vdash \sigma_2 -(\{\text{ClassT } \text{Ext}\}(\text{Class } \text{Ext})x.\text{vee}:=\text{Lit } (\text{Intg } 1)) \triangleright \rightarrow \sigma_3}{\Gamma \vdash \sigma_2 -\text{Expr}(\{\text{ClassT } \text{Ext}\}(\text{Class } \text{Ext})x.\text{vee}:=\text{Lit } (\text{Intg } 1)) \rightarrow \sigma_3} (\text{Blk})$

$\frac{\Gamma \vdash \sigma_1 -e \triangleright \text{Addr } a \rightarrow \sigma_1 \quad \text{cmethd } \Gamma \text{ Ext } (\text{foo}, \text{Class } \text{Base}) = \text{Some } (-, (x, -), [], \text{blk}, \text{Lit } \text{Null})}{\Gamma \vdash \sigma_1 -\text{Lit } \text{Null} \triangleright \text{Null} \rightarrow \sigma_1 \quad \sigma_2 = \text{Norm } (h, \text{empty}[\text{this} \mapsto \text{Addr } a][x \mapsto \text{Null}]) \quad (\text{Blk}) \quad \Gamma \vdash \sigma_3 -\text{Lit } \text{Null} \triangleright \rightarrow \sigma_3}$

$\frac{\Gamma \vdash \sigma_1 -\text{Lit } \text{Null} \triangleright \text{Null} \rightarrow \sigma_1 \quad \sigma_2 = \text{Norm } (h, \text{empty}[\text{this} \mapsto \text{Addr } a][x \mapsto \text{Null}]) \quad (\text{Blk}) \quad \Gamma \vdash \sigma_3 -\text{Lit } \text{Null} \triangleright \rightarrow \sigma_3}{\Gamma \vdash \sigma_1 -(\text{e.foo}(\{\text{Class } \text{Base}\} \text{Lit } \text{Null})) \triangleright \rightarrow \sigma_4} (\text{Call}')$

$\frac{\Gamma \vdash \sigma_1 -(\text{e.foo}(\{\text{Class } \text{Base}\} \text{Lit } \text{Null})) \triangleright \rightarrow \sigma_4}{\Gamma \vdash \sigma_1 -\text{Expr}(\text{e.foo}(\{\text{Class } \text{Base}\} \text{Lit } \text{Null})) \rightarrow \sigma_4} (\text{Call}')$

$\frac{\Gamma \vdash \sigma_6 -x \triangleright \text{Addr } b \rightarrow \sigma_6 \quad \sigma_7 = (\text{Some } (\text{XcptLoc } b), \text{snd } \sigma_6)}{\Gamma \vdash \sigma_6 -\text{throw } x \rightarrow \sigma_7} (\text{Throw})$

$\frac{\text{new_Addr empty} = \text{Some } (a, \text{None})}{\Gamma \vdash \sigma_0 -\text{new } \text{Ext} \triangleright a \rightarrow \text{Norm } (h, \text{empty})}$

$\frac{\Gamma \vdash \sigma_0 -\text{new } \text{Ext} \triangleright a \rightarrow \text{Norm } (h, \text{empty})}{\Gamma \vdash \sigma_0 -(\text{e}:=\text{new } \text{Ext}) \triangleright \rightarrow \sigma_1} (\text{Call}')$

$\frac{\Gamma \vdash \sigma_0 -(\text{e}:=\text{new } \text{Ext}) \triangleright \rightarrow \sigma_1}{\Gamma \vdash \sigma_0 -\text{Expr}(\text{e}:=\text{new } \text{Ext}) \rightarrow \sigma_1} (\text{Call}')$

$\frac{\text{new_Addr } h = \text{Some } (b, \text{None})}{\Gamma \vdash \sigma_4 -\text{saloc} \rightarrow \sigma_5} (\text{Throw})$

$\frac{\Gamma \vdash \sigma_4 -\text{saloc} \rightarrow \sigma_5 \quad \Gamma, \text{snd } \sigma_5 \vdash \text{Addr } b \text{ fits Class } (\text{SXcpt } \text{NP}) \quad (\text{Throw})}{\Gamma \vdash \sigma_1 -(\text{try Expr}(\text{e.foo}(\{\text{Class } \text{Base}\} \text{Lit } \text{Null})) \text{ catch } ((\text{SXcpt } \text{NP}) x) (\text{throw } x)) \rightarrow \sigma_7}$

$\frac{\Gamma \vdash \sigma_0 -\text{Expr}(\text{e}:=\text{new } \text{Ext}) \rightarrow \sigma_1}{\Gamma \vdash \sigma_0 -(\text{Expr}(\text{e}:=\text{new } \text{Ext}); \text{try Expr}(\text{e.foo}(\{\text{Class } \text{Base}\} \text{Lit } \text{Null})) \text{ catch } ((\text{SXcpt } \text{NP}) x) (\text{throw } x)) \rightarrow \sigma_7}$

$\frac{\Gamma \vdash \sigma_0 -\text{Expr}(\text{e}:=\text{new } \text{Ext}) \rightarrow \sigma_1 \quad \Gamma \vdash \sigma_1 -(\text{try Expr}(\text{e.foo}(\{\text{Class } \text{Base}\} \text{Lit } \text{Null})) \text{ catch } ((\text{SXcpt } \text{NP}) x) (\text{throw } x)) \rightarrow \sigma_7}{\Gamma \vdash \sigma_0 -(\text{Expr}(\text{e}:=\text{new } \text{Ext}); \text{try Expr}(\text{e.foo}(\{\text{Class } \text{Base}\} \text{Lit } \text{Null})) \text{ catch } ((\text{SXcpt } \text{NP}) x) (\text{throw } x)) \rightarrow \sigma_7}$

5 The proof of type soundness

In this section we discuss our type soundness theorem together with its crucial lemmas. As we spent almost half of the proof effort deriving properties of the type relations and the structure of well-formed programs, we dedicate to them subsections of their own before introducing helpful notions concerning type soundness, the main theorem itself, and interesting corollaries.

It is not surprising that many of them are similar to those given by Drossopoulou and Eisenbach [DE98] since the necessity of certain lemmas emerges quite naturally. On the other hand, the proof principles we use are sometimes rather different from those outlined in their earlier paper [DE97], some of which were inadequate.

5.1 Lemmas on the type relations

There are two non-trivial lemmas concerning the type relations of BALI, namely the well-foundedness `wf` of the converse subinterface and subclass relations

$$\text{wf_prog } \Gamma \longrightarrow \text{wf } (\lambda(J, I). \Gamma \vdash I \prec_i J) \\ \wedge \text{wf } (\lambda(D, C). \Gamma \vdash C \prec_c D)$$

and the frequently used transitivity of the widening relation:

$$\text{wf_prog } \Gamma \wedge \Gamma \vdash S \preceq U \wedge \Gamma \vdash U \preceq T \longrightarrow \Gamma \vdash S \preceq T$$

The two relations are well-founded because they are finite and acyclic, where the former is a consequence of representing class and interface declarations as lists, and the latter follows from the irreflexivity of the relations, which in turn follows from the well-formedness of the classes and interfaces implied by the well-formedness of the whole program.

The well-foundedness facts are necessary for deriving the recursion equations for the functions that traverse the type hierarchy of a program (see §4.1) and also give rise to induction principles for the (direct) subinterface and subclass relations, e.g. the rule

$$\frac{\text{wf_prog } \Gamma; P \text{ Object}; \\ \forall C D. C \neq \text{Object} \wedge \Gamma \vdash C \prec_c^1 D \wedge \dots \wedge P D \longrightarrow P C}{\forall E. \text{is_class } \Gamma E \longrightarrow P E}$$

means that for a well-formed program, if some property hold for class `Object` and is preserved by the direct subclass relation, it holds for all classes.

Most lemmas, as well as auxiliary properties for deriving them, typically rely on several well-formedness conditions and are usually proved by rule induction on the type relation involved, or by applying the induction principles just mentioned. For example, the transitivity of $_ \vdash _ \preceq _$ is proved by rule induction on the widening relation. It requires a well-formed program because it uses the properties that every class widens to `Object` and that `Object` has neither a superclass nor a superinterface.

5.2 Lemmas on fields and methods

For the type-safety of field accesses and method calls, characteristic lemmas concerning the field lookup and method lookup are required. They are used to relate the (static) types of fields and methods, as determined at compile-time, to the actual (dynamic) types that occur at run-time.

For example, fields correctly referred to at compile-time must be found at run-time. More formally, if a field access $\{T\}e.fn$, where e is of type $\text{Class } C$, statically refers to a field of type fT defined in the reference type T , then within an instance of some class C' , which may be a subclass of C , the field can be (dynamically) referred to using the same name and its defining class. In particular, there is no dynamic binding for fields. This fact requires the following lemma:

$$\text{wf_prog } \Gamma \wedge \text{cfield } \Gamma \ C \ fn = \text{Some } (T, fT) \wedge \Gamma \vdash \text{Class } C' \preceq \text{Class } C \longrightarrow \\ \text{table_of } (\text{fields } \Gamma \ C') \ (fn, T) = \text{Some } fT$$

Concerning method calls, a similar requirement preventing ‘method not understood’ errors can be formalized: if a method call of the form $e.mn(\{pT\}p)$ with $E \vdash e :: \text{RefT } T$ refers to a method that is statically available for the reference e , the dynamic lookup of the object pointed at by e should yield a method with a compatible result type. The lemma that helps to establish this behavior reads as follows: for a well-formed program, a reference type T , and any class type T_1 that widens to T , if T (statically) supports a method with a given signature, then the (dynamic) type T_1 supports a method with the same signature and whose result type widens to the result type of the first method:

$$\text{wf_prog } \Gamma \wedge (m_1, (pn_1, rT_1)) \in \text{mheads } \Gamma \ T \ sig \wedge \Gamma \vdash \text{Class } T_1 \preceq \text{RefT } T \longrightarrow \\ \exists m_2 \ pn_2 \ rT_2 \ b. \text{cmethd } \Gamma \ T_1 \ sig = \text{Some } (m_2, (pn_2, rT_2), b) \wedge \Gamma \vdash rT_2 \preceq rT_1$$

The proofs of these lemmas are lengthy and require many auxiliary theorems that are proved by induction on the direct subclass relation, by case splitting on the right-hand argument of the widening relation and by rule induction on the subinterface, subclass, and implementation relation.

5.3 Type soundness

Finally, we state and prove the type soundness theorem. We motivate how we express type soundness, comment on the proof of the main theorem, and discuss its consequences.

Goal Type soundness is a relation between the type system and the semantics of a language meaning that all values produced during any program execution respect their static types. This can be formulated as a preservation property: For all state transformations caused by executing a statement or evaluating an expression, if in the original state the contents of all variables “conform” to their respective types, this holds also for any final state. Additionally, if an expression yields some result, this value “conforms” to the type of the expression. Of course, we can only expect all this to hold if we assume a well-formed program and well-typed statements and expressions.

It remains to specify what we mean exactly by ‘conforms’, which is inspired by [DE98]. Relative to a given program Γ and a state σ , a value v conforms to a type T , written $\Gamma, \sigma \vdash v :: T$, iff the dynamic type of v widens to T . Via two auxiliary conformance concepts, this can be lifted to the notion of a whole state σ conforming to an environment E . The proposition $\sigma :: \preceq E$ means that the value of any accessible variable within the state is compatible with its static type. Formally, these four concepts

- $_, - \vdash - :: \preceq - :: \text{prog} \Rightarrow \text{st} \Rightarrow \text{val} \Rightarrow \text{ty} \Rightarrow \text{bool}$
of a value conforming to a type,
- $_, - \vdash - [:: \preceq] - :: \text{prog} \Rightarrow \text{st} \Rightarrow (\alpha, \text{val}) \text{table} \Rightarrow (\alpha, \text{ty}) \text{table} \Rightarrow \text{bool}$
of all values in a table conforming to their respective types,
- $_, - \vdash - :: \preceq \diamond :: \text{prog} \Rightarrow \text{st} \Rightarrow \text{obj} \Rightarrow \text{bool}$
of all components of an object conforming to their respective types, and
- $- :: \preceq - :: \text{state} \Rightarrow \text{env} \Rightarrow \text{bool}$
of a state conforming to an environment

are defined as follows:

$$\begin{aligned}
\Gamma, \sigma \vdash v :: \preceq T &\stackrel{\text{def}}{=} \text{let } \text{dyn_ty} = \text{option_map } \text{obj_ty} \circ \text{heap } \sigma \\
&\quad \text{in } \exists T'. \text{typeof } \text{dyn_ty } v = \text{Some } T' \wedge \Gamma \vdash T' \preceq T \\
\Gamma, \sigma \vdash \text{vs} [:: \preceq] T_s &\stackrel{\text{def}}{=} \forall n T. T_s \ n = \text{Some } T \longrightarrow \\
&\quad (\exists v. \text{vs } n = \text{Some } v \wedge \Gamma, \sigma \vdash v :: \preceq T) \\
\Gamma, \sigma \vdash \text{Obj } C \text{ fs} :: \preceq \diamond &= \Gamma, \sigma \vdash \text{fs} [:: \preceq] \text{table_of } (\text{fields } \Gamma \ C) \\
\Gamma, \sigma \vdash \text{Arr } T \text{ cs} :: \preceq \diamond &= \Gamma, \sigma \vdash \text{cs} [:: \preceq] \text{option_map } (\lambda i. T) \circ \text{cs} \\
(x, \sigma) :: \preceq (\Gamma, \Delta) &\stackrel{\text{def}}{=} \Gamma, \sigma \vdash \text{locals } \sigma [:: \preceq] \Delta \wedge \\
&\quad (\forall a \text{ obj. } \text{heap } \sigma \ a = \text{Some } \text{obj} \longrightarrow \Gamma, \sigma \vdash \text{obj} \quad :: \preceq \diamond) \wedge \\
&\quad (\forall a. \quad x = \text{Some}(\text{XcptLoc } a) \longrightarrow \Gamma, \sigma \vdash \text{Addr } a :: \preceq \text{Class}(\text{SXcpt Throwable}))
\end{aligned}$$

The expression $(\text{option_map } \text{obj_ty} \circ \text{heap } \sigma) \ a$ calculates the dynamic type of the object (if any) at address a on the heap. Note that the conformance relation is defined such that it does not take into account inaccessible variables, i.e. values that occur in the state but not in the corresponding component of the static environment. Among others, this frees us from explicitly deallocating exception parameters after a catch clause.

With the help of the notions just introduced, we can express the propositions we aim to prove as follows. In the context of a well-formed program, the execution of a well-typed statement transforms a state conforming to the environment into another state that again conforms to the environment:

$$E = (\Gamma, \Delta) \wedge \text{wf_prog } \Gamma \wedge E \vdash s :: \diamond \wedge \sigma :: \preceq E \wedge \Gamma \vdash \sigma \text{ --s--} \sigma' \longrightarrow \sigma' :: \preceq E$$

Analogously, the evaluation of a well-typed expression preserves the conformance of the state to the environment where, unless an exception has occurred, the value of the expression conforms to its static type:

$$\begin{aligned}
E = (\Gamma, \Delta) \wedge \text{wf_prog } \Gamma \wedge E \vdash e :: T \wedge \sigma :: \preceq E \wedge \Gamma \vdash \sigma \text{ --e>v--} (x', \sigma') &\longrightarrow \\
(x', \sigma') :: \preceq E \wedge (x' = \text{None} \longrightarrow \Gamma, \sigma' \vdash v :: \preceq T) &
\end{aligned}$$

The validity of these two formulas will result as trivial corollaries from the main theorem, given next.

Main theorem and proof To prove the intended type soundness theorems given above, we utilize rule induction on the derivation on the execution of statements and the evaluation of expressions. As these depend on each other, we must deal with statements and expressions simultaneously. In addition, in order to obtain a suitable induction hypothesis, we have to strengthen the propositions by adding the auxiliary “heap extension” predicate $_ \trianglelefteq _$ (defined below) and introducing universal quantifications explicitly at several positions. As a result, the main theorem looks quite formidable, yet we attempt to cast it into words:

$$\begin{array}{l}
\text{wf_prog } \Gamma \longrightarrow \\
(\Gamma \vdash (x, \sigma) - c \rightarrow (x', \sigma') \longrightarrow \\
\quad \forall A. (x, \sigma) :: \preceq (\Gamma, A) \longrightarrow \\
\quad (\Gamma, A) \vdash c :: \diamond \longrightarrow \\
\quad (x', \sigma') :: \preceq (\Gamma, A) \wedge \sigma \trianglelefteq \sigma') \\
\wedge \\
(\Gamma \vdash (x, \sigma) - e \triangleright v \rightarrow (x', \sigma') \longrightarrow \\
\quad \forall A. (x, \sigma) :: \preceq (\Gamma, A) \longrightarrow \\
\quad \forall T. (\Gamma, A) \vdash e :: T \longrightarrow \\
\quad (x', \sigma') :: \preceq (\Gamma, A) \wedge \sigma \trianglelefteq \sigma' \wedge (x' = \text{None} \longrightarrow \Gamma, \sigma' \vdash v :: \preceq T))
\end{array}$$

For a well-formed program Γ , if the execution of a statement transforms one state into another then for all local environments A , if the the statement is well-typed according to the environment (Γ, A) and the first state conforms to it, so does the second state, and the new heap is an extension of the old one. The same holds for expressions, but additionally the value of the expression conforms to its type, in case there is no exception.

The “heap extension” is a pre-order on states of type $st \Rightarrow st \Rightarrow bool$, where $\sigma \trianglelefteq \sigma'$ means that any object existing on the heap of σ also exists on σ' and has the same type there. (If we considered garbage collection, we would have to restrict this proposition to accessible objects.) The heap extension property holds for any transition of the operational semantics, which turns out to be necessary in our inductive proof.

$$\begin{aligned}
\sigma \trianglelefteq \sigma' &\stackrel{\text{def}}{=} \forall a \text{ obj. } \text{heap } \sigma \ a = \text{Some } \text{obj} \longrightarrow \\
&\quad \exists \text{obj}' . \text{heap } \sigma' \ a = \text{Some } \text{obj}' \wedge \text{obj_ty } \text{obj}' = \text{obj_ty } \text{obj}
\end{aligned}$$

The proof of the main type soundness theorem is by far the heaviest. At the top level, it consists of currently 21 cases, one for each evaluation rule, where

- 8 cases can be solved rather directly (e.g. from the induction hypothesis),
- 7 cases require just simple lemmas on the structure of the state, and
- the remaining 6 cases require extensive reasoning on the characteristic properties of the constructs concerned.

Most of this reasoning is independent of the operational semantics itself and can be tackled separately, which keeps the main proof manageable.

Consequences A corollary of type soundness is that method calls always execute a suitable method, i.e. a ‘method not understood’ run-time error is impossible. This property can be stated more formally: for a well-formed program and a state that conforms to the environment, if an expression of reference type (which plays the role of the target expression for the method call considered) evaluates without an exception to a non-null reference, and if for that (static) type and a given signature a method is available, the dynamic method lookup for the same signature according to the class instance pointed at by the reference value yields a proper method body:

$$E = (\Gamma, A) \wedge \text{wf_prog } \Gamma \wedge E \vdash e :: \text{RefT } T \wedge \sigma :: \preceq E \wedge \Gamma \vdash \sigma -e \triangleright a' \rightarrow \text{Norm } \sigma' \wedge a' \neq \text{Val Null} \wedge \text{dynT} = \text{fst } (\text{the_Obj } (\text{heap } \sigma' (\text{the_Addr } a'))) \wedge \text{mheads } \Gamma \text{ } T \text{ sig} \neq \{\} \longrightarrow \exists m. \text{cmethd } \Gamma \text{ dynT sig} = \text{Some } m$$

This implies that in a well-formed context, in every instance of the evaluation rule for method calls, the function `cmethd` returns a proper method body.

As it stands, the type soundness theorem does not directly say anything about non-terminating computations, which might lead to the conclusion that it is useless for the type-safety of reactive systems and looping programs. Fortunately, the theorem guarantees type-safety even in such cases if one accepts the following meta-level reasoning. An infinite computation can be interrupted after any finite number of computation steps, for example by introducing a counter of steps and raising an exception when a given value has been reached. The theorem implies that the state resulting from interrupting the computation after any finite number of statements executed conforms to the environment. Together with the fact that there is no single non-terminating statement, the whole (infinite) computation can be concluded to be type-safe.

In addition to the evaluation semantics, we plan to define a transition semantics and prove both styles equivalent (for finite computations). The transition semantics will be less concise and abstract, but allows type soundness to be formulated as a subject reduction property, which is more natural for infinite computations. More importantly, it seems to be unavoidable to describe concurrency (and I/O).

6 Experience and statistics

Recalling our design goals stated at the end of §2, we comment how far we have reached them and share some of the lessons learned during the project.

Faithfulness to the official language specification HOL’s expressiveness enables us to formalize the Java specification quite naturally and directly, without facing any severe obstacles. There is almost a one-to-one correspondence between the concepts given in the specification and those defined in BALI. As far as we could tell, all the messy well-formedness conditions inherited from the language specification are actually needed somewhere in the proofs. This inspires confidence in the adequacy of both the specification and our formalization.

We do not yet have tools for automatically generating executable code from our theories, which would be an additional help in validating our formalization. The importance of such a mechanism became very obvious when we uncovered a mistake in our formalization (which was not present in [NO98] but was introduced by modifications) when symbolically executing the example in this article in Isabelle: the list returned by function `fields` was in reverse order. Although the type soundness proof itself was an excellent debugging mechanism which caught many minor and some major mistakes, it failed to detect the wrong order because type soundness is independent of the order in which fields are inherited. In the original language specification we did not find any significant errors, but some omissions and unneeded restrictions, which we lifted.

Succinctness and simplicity Our policy to restrict the number of features considered and to make straightforward simplifications that do not diminish the expressiveness of the language has led to a clear and straightforward formalization. Mixfix syntax and mathematical fonts as offered by Isabelle also contribute greatly to moderately readable definitions and theorems.

The facility to conduct concise proofs strongly depends on the formalization. In our case, the use of the (also more elegant) evaluation semantics saved us from a lot of trouble, while the intricacies of a transition semantics faced by Drossopoulou and Eisenbach [DE97] led to several mistakes that were finally corrected during Syme’s machine-checked proof [Sym97b], but at the expense of additional concepts.

Maintainability and extendibility Unless the language changes drastically, modifications tend to be of a local nature, but only if both the formalization and the proofs are reasonably structured. As always, modularity is the key issue. But when the formalization is extended, even well-structured proofs need to be modified, which remains a tedious job. Higher-level proof scripts and more automation are some of the answers. A dedicated mechanism for change management exploring and fixing the impact of modifications would also help.

We are reasonably happy with the modularity of our work. For instance, Martin Büchi [BW98] has adopted the formalization (including the proofs), extended it to handle compound types, and proved the type-safety of the augmented language, all of which worked very smoothly.

Adequacy for the theorem prover Theorem provers are notoriously sensitive to the precise formulation of definitions and theorems. Thus the two goals of maximal automation of proofs and maximal abstractness of definitions are sometimes in conflict. In a number of cases this meant that although we could start with an abstract definition, we had to derive consequences which were better suited for the available proof procedures. Although we are far from satisfied with the current status of Isabelle’s proof procedures (for example, the handling of assumptions during simplification, or the necessity to expand tuples and similar datatypes by hand), they are basically adequate for the task at hand. Nevertheless, more automation is necessary and feasible by extending the capabilities of Isabelle itself.

Statistics We spent two months (estimated net time) developing and maintaining our formalization, and the Isabelle theory files produced add up to about 1200 lines of well-documented definitions. To conduct and maintain the type soundness proof with all necessary lemmas, it took us roughly three months of work and about 2400 lines of proof scripts.

7 Conclusion

The reader has been exposed to large chunks of a formal language specification and a proof of type soundness and may need to be reminded of the benefits. Even including the slight generalizations mentioned at the beginning of §4, we did not discover a loop-hole in the type system. But we had not seriously expected this either. So what have we gained over and above a level of certainty far beyond any paper-and-pencil proof?

We view our work primarily as an investment for the future. For a start, it can serve as the basis for many other mechanized proofs about Java, e.g. as a foundation for the work by Dean [Dea97] or for compiler correctness. More importantly, we see machine-checked proofs as an invaluable aid in maintaining large language designs (or formal documents of any kind). It is all very well to perform a detailed proof on paper once, but in the face of changes and extensions, the reliability of such proofs begins to crumble. In contrast, we developed the design incrementally, and Isabelle reminded us where proofs needed to be modified. This has shown to be important, for example when we extended BALI with full exception handling. It will continue to help us further: apart from adding the last important Java features missing from BALI, e.g. threads, we also plan to use BALI as a vehicle for experimental extensions of Java such as parameterized types [MBL97,OW97,AFM97].

Despite our general enthusiasm for machine-checked language designs, a few words of warning are in order:

- BALI is still a half-way house: not a toy language any more, but missing many details and some important features of Java.
- The Java type system is, despite subclassing, simpler than that of your average functional language: whereas the type checking rules of Java are almost directly executable, the verification of ML’s type inference algorithm against the type system requires a significant effort [NN98]. The key complication there is the presence of free and bound type variables, which requires complex reasoning about substitutions. VanInwegen [Van97] reports similar difficulties in her formalization of the type system and the semantics of ML.
- Theorem provers, and Isabelle is no exception, require a certain learning effort due to the machine-oriented proof style. Recent moves towards a more human-oriented proof style like Syme’s DECLARE system [Sym97a] promise to lower this hurdle. However, as Harrison [Har97] points out, both proof styles have their merits, and we are currently investigating a combination.

In a nutshell: although machine-checked language designs for the masses are still some way off, this article demonstrates that they have definitely become a viable option for the expert.

Acknowledgments We thank Sophia Drossopoulou, Donald Syme and Egon Börger for the very helpful discussions about their related work. We also thank Wolfgang Naraschewski, Markus Wenzel, Andrew Gordon and several anonymous referees for their comments on earlier reports on our project.

References

- [AFM97] Ole Agesen, Stephen N. Freund, and John C. Mitchell. Adding type parameterization to the Java language. In *ACM Symp. Object-Oriented Programming: Systems, Languages and Applications*, 1997.
- [BCM⁺93] Kim B. Bruce, Jon Crabtree, Thomas P. Murtagh, Robert van Gent, Allyn Dimock, and Robert Muller. Safe and decidable type checking in an object-oriented language. In *Proc. OOPSLA'93*, volume 18 of *ACM SIGPLAN Notices*, pages 29–46, October 1993.
- [BM88] Robert S. Boyer and J Strother Moore. *A Computational Logic Handbook*. Academic Press, 1988.
- [Bru93] Kim B. Bruce. Safe type checking in a statically-typed object-oriented programming language. In *Proc. 20th ACM Symp. Principles of Programming Languages*, pages 285–298. ACM Press, 1993.
- [BS98] Egon Börger and Wolfram Schulte. A programmer friendly modular definition of the dynamic semantics of Java. In Jim Alves-Foss, editor, *Formal Syntax and Semantics of Java*, Lect. Notes in Comp. Sci. Springer-Verlag, 1998. Chapter 11 of this volume.
- [BvGS95] Kim B. Bruce, Robert van Gent, and Angela Schuett. PolyTOIL: A type-safe polymorphic object-oriented language. In W. Olthoff, editor, *ECOOP '95*, volume 952 of *Lect. Notes in Comp. Sci.*, pages 27–51. Springer-Verlag, 1995.
- [BW98] Martin Büchi and Wolfgang Weck. Java needs compound types. Technical Report 182, Turku Center for Computer Science, May 1998. <http://www.abo.fi/~mbuechi/publications/CompoundTypes.html>.
- [Coh97] Richard M. Cohen. The defensive Java Virtual Machine specification. Technical report, Computational Logic Inc., 1997. Draft version.
- [Coo89] William Cook. A proposal for making Eiffel type-safe. In *Proc. ECOOP'89*, pages 57–70. Cambridge University Press, 1989.
- [DE97] Sophia Drossopoulou and Susan Eisenbach. Is the Java type system sound? In *Proc. 4th Int. Workshop Foundations of Object-Oriented Languages*, January 1997.
- [DE98] Sophia Drossopoulou and Susan Eisenbach. Java is type safe — probably. In Jim Alves-Foss, editor, *Formal Syntax and Semantics of Java*, Lect. Notes in Comp. Sci. Springer-Verlag, 1998. Chapter 3 of this volume.
- [Dea97] Drew Dean. The security of static typing with dynamic linking. In *Proc. 4th ACM Conf. Computer and Communications Security*. ACM Press, 1997.
- [GJS96] James Gosling, Bill Joy, and Guy Steele. *The Java Language Specification*. Addison-Wesley, 1996.

- [GM93] M.J.C. Gordon and T.F. Melham, editors. *Introduction to HOL: a theorem-proving environment for higher order logic*. Cambridge University Press, 1993.
- [Har97] John Harrison. Proof style. Technical Report 410, University of Cambridge Computer Laboratory, 1997.
- [MBL97] Andrew C. Myers, Joseph A. Bank, and Barbara Liskov. Parameterized types for Java. In *Proc. 24th ACM Symp. Principles of Programming Languages*, pages 132–145, 1997.
- [Mil78] Robin Milner. A theory of type polymorphism in programming. *J. Comp. Sys. Sci.*, 17:348–375, 1978.
- [NN98] Wolfgang Naraschewski and Tobias Nipkow. Type inference verified: Algorithm W in Isabelle/HOL. In E. Giménez and C. Paulin-Mohring, editors, *Types for Proofs and Programs: Intl. Workshop TYPES '96*, volume 1512 of *Lect. Notes in Comp. Sci.*, pages 317–332. Springer-Verlag, 1998.
- [NO98] Tobias Nipkow and David von Oheimb. *Java^{light}* is type-safe — definitely. In *Proc. 25th ACM Symp. Principles of Programming Languages*, pages 161–170. ACM Press, 1998.
- [OW97] Martin Odersky and Philip Wadler. Pizza into Java: Translating theory into practice. In *Proc. 24th ACM Symp. Principles of Programming Languages*, pages 146–159, 1997.
- [Pau94] Lawrence C. Paulson. *Isabelle: A Generic Theorem Prover*, volume 828 of *Lect. Notes in Comp. Sci.* Springer-Verlag, 1994.
- [Sli96] Konrad Slind. Function definition in higher order logic. In J. von Wright, J. Grundy, and J. Harrison, editors, *Theorem Proving in Higher Order Logics*, volume 1125 of *Lect. Notes in Comp. Sci.*, pages 381–397. Springer-Verlag, 1996.
- [Sym97a] Donald Syme. DECLARE: A prototype declarative proof system for higher order logic. Technical Report 416, University of Cambridge Computer Laboratory, 1997.
- [Sym97b] Donald Syme. Proving Java type soundness. Technical Report 427, University of Cambridge Computer Laboratory, 1997.
- [Sym98] Donald Syme. Proving Java type soundness. In Jim Alves-Foss, editor, *Formal Syntax and Semantics of Java*, Lect. Notes in Comp. Sci. Springer-Verlag, 1998. Chapter 4 of this volume.
- [Van97] Myra VanInwegen. Towards type preservation for core SML. University of Cambridge Computer Laboratory, 1997.

Index

$::$, 3
 $\stackrel{\text{def}}{=}$, 3
 \Rightarrow , 3
 \times , 4
 $|$, 3
 ε , 3
“”, 4
 $[]$, 4
#, 4
 Γ , 6
 \oplus , 8
 $\oplus\oplus$, 8
 $-[-\mapsto-]$, 8
;, 10
 $(-)_-$, 11
 $-:=_-$, 11
 $\{-\}_-$, 11
 $\{-\}_- :=_-$, 11
 $-[-]$, 11
 $-[-] :=_-$, 11
 $-.-(\{-\}_-)$, 11
 $-[]$, 12
 $- \vdash _ \lambda_i^1 _$, 13
 $- \vdash _ \lambda_c^1 _$, 13
 $- \vdash _ \sim^1 _$, 13
 $- \vdash _ \lambda _$, 13
 $- \vdash _ \lambda? _$, 13
 $- \vdash _ :: \diamond$, 14
 $- \vdash _ :: _$, 14
 $- \vdash _ _ \rightarrow _$, 22
 $- \vdash _ _ \text{salloc} \rightarrow _$, 24
 $- \vdash _ _ \triangleright \rightarrow _$, 22
 $_$, $- \vdash _ :: \lambda _$, 30
 $_$, $- \vdash _ [-::\lambda] _$, 30
 $_$, $- \vdash _ :: \lambda \diamond$, 30
 $-::\lambda _$, 30
 $- \triangleleft _$, 31

Addr, 11
appl_methds, 16
ArrayT, 12
ArrStore, 6

Bool, 11
bool, 4
boolean, 12

cdecl, 6

cfield, 16
Class, 12
class, 6
class, 9
ClassCast, 6
ClassT, 12
cmethd, 9
components, 21

default_val, 12

empty, 8
ENAME, 6
ename, 6
ename0, 6
env, 14
Expr, 10

fdecl, 7
field, 7
fields, 21
fields, 9
finally, 10
fits, 21
fst, 4

heap, 20
heap, 20
hiding entails, 8
hidings entails, 8
hupd, 21

idecl, 6
if_else, 10
iface, 12
iface, 6
iface, 9
ifaceT, 12
imethds, 9
IndOutBound, 6
init_Arr, 22
init_Obj, 22
init_vars, 22
instanceof, 11
Int, 11
int, 4
int, 12
is_class, 12

is_iface, 12
 is_type, 12

 lcl, 14
lenv, 14
list, 4
 Lit, 11
locals, 20
 locals, 20
 lupd, 21
lvar, 7

 map, 4
 max_spec, 16
mbody, 7
mdecl, 7
methd, 7
mhead, 7
 mheads, 16
mname, 6
 more_spec, 16

 NegArrSize, 6
 new, 11
 new_Addr, 21
 None, 4
 np, 22
 NT, 12
 Null, 11
 NullPointer, 6
 NullT, 12

 o2s, 4
obj, 21
 obj_ty, 21
 Object, 6
 ObjectC, 19
option, 4
 option_map, 16
 OutOfMemory, 6

 prg, 14
prim_ty, 12
 PrimT, 12
prog, 6

 raise_if, 22
ref_ty, 12
 RefT, 12

set, 4

 set, 4
sig, 7
 Skip, 10
 snd, 4
 Some, 4
st, 20
state, 20
stmt, 10
 SXcpt, 6
 SXcptC, 19
 SysXcpt, 20

table, 8
 table_of, 8
tables, 8
 the, 4
 the_Addr, 11
 the_Arr, 21
 the_Bool, 11
 the_Int, 11
 the_Obj, 21
 this, 6
 throw, 10
 Throwable, 6
 TName, 6
tname, 6
tname0, 6
 try_catch, 10
ty, 12
 typeof, 15

 Un_tables, 8
 unique, 8
 Unit, 11

val, 11
 void, 12

 wf_cdecl, 18
 wf_fdecl, 18
 wf_idecl, 18
 wf_mdecl, 18
 wf_mhead, 18
 wf_prog, 18
 while, 10

 x_case, 21
xcpt, 20
 XcptLoc, 20
xname, 6