Static single information form for abstract compilation

Davide Ancona and Giovanni Lagorio

DISI, University of Genova, Italy
{davide,lagorio}@disi.unige.it

Abstract. In previous work we have shown that more precise type analysis can be achieved by exploiting union types and static single assignment (SSA) intermediate representation (IR) of code. Since most modern compilers implement efficient algorithms for translating code in SSA IR, focusing on analysis of SSA IR not only allows more precise results, but also favor reuse of compiler technology, and better integration with existing compilers.

In this paper we exploit static single information (SSI), an extension of SSA proposed in literature and adopted by some compilers, to allow assignments of more precise types to variables in conditional branches. In particular, SSI can be exploited rather easily and effectively to assign more precise static types in presence of explicit runtime typechecking, a necessity that occurs rather often in statically typed object-oriented languages, and even more than often in dynamically typed ones.

We show how the use of SSI form can be smoothly integrated with abstract compilation, our approach to static type analysis that consists in solving goals w.r.t. the coinductive semantics of a constraint logic program automatically generated from the source program to be analyzed. In particular, we define abstract compilation based on union and nominal types for a simple dynamically typed Java-like language in SSI form with a runtime typechecking operator, to show how precise the proposed analysis can be.

1 Introduction

In previous work [5] we have shown that more precise type analysis can be achieved by exploiting union types and static single assignment (SSA) [6] intermediate representation (IR) of code. Most modern compilers (among others, GNU's GCC [11], the SUIF compiler system [10], and JIT compilers including Java HotSpot [9], and Java Jikes RVM [8]) implement efficient algorithms for translating code in SSA IR [7], therefore focusing on analysis of SSA IR not only allows more precise results, but also favor reuse of compiler technology, and better integration with existing compilers. In particular, we have studied how *abstract compilation* [4, 3, 5] can naturally take advantage of SSA IR when union types are considered. Abstract compilation aims to reconcile static type analysis and symbolic execution: one can check whether the execution of a certain expression e is type safe, when variable contents range over possibly infinite

sets of values (represented by types), by solving a goal, obtained by abstract compilation of e, w.r.t. the coinductive[1] semantics of the constraint logic program automatically generated from the source program in which the expression is executed.

Abstract compilation fosters *plug-and-play* static type analysis since one can provide several compilation schemes for the same language, each corresponding to a different kind of analysis, without changing the inference engine that implements the coinductive semantics of constraint logic programs [13, 12, 3]. For instance, in previous work we have defined compilation schemes for Java-like languages based on union and structural object types, to support parametric and data polymorphism, (that is, polymorphic methods and fields) that allow quite precise type analysis, and a smooth integration with the nominal type annotations contained in the programs [4]; other proposed compilation schemes aim to detect uncaught exceptions for Java-like languages [3], or to integrate SSA IR in presence of imperative features [5].

In this paper we exploit static single information (SSI), an extension of SSA proposed in literature [1, 15], to allow assignments of more precise types to variables in conditional branches. SSI has been already adopted by compilers as LLVM [16], PyPy [2], and SUIF [14], and proved to be more effective than SSA for performing data flow analysis, program slicing, and interprocedural analysis.

In particular, we show how SSI can be exploited rather easily and effectively by abstract compilation to assign more precise static types in presence of explicit runtime typechecking, a necessity that occurs rather often in statically typed object-oriented languages [17], and even more than often in dynamically typed ones.

To this aim, we formally define the operational semantics of a simple dynamically typed Java-like language in SSI form equipped with a runtime typechecking operator, and then provide an abstract compilation scheme based on union and nominal types supporting more precise type analysis of branches guarded by explicit runtime typechecks.

The paper is structured as follows: Section 2 introduces SSA and SSI IRs and motivates their usefulness for type analysis; Section 3 formally defines the SSI IR of a dynamically typed Java-like language equipped with an operator **instanceof** for runtime typechecking. Section 4 presents a compilation scheme for the defined IR, based on nominal and union types, and Section 5 concludes with some considerations on future work.

2 Type analysis with SSA and SSI

In this section SSA and SSI IRs are introduced and their usefulness for type analysis is motivated.

---

[1] Coinduction allows proper treatment of recursive types and methods [4].

Type analysis with static single assignment form

Method `read()` declared below, in a hypothetical dynamically typed Java-like language, creates and returns a shape which is read through method `nextLine()` that reads the next available string from some input source. The partially omitted methods `readCircle()` and `readSquare()` read the needed data from the input, create, and return a new corresponding instance of `Circle` or `Square`.

```
class ShapeReader {
   ...
   nextLine() {...}
   readCircle() { ... return new Circle(...); }
   readSquare() { ... return new Square(...); }
   read() {
      st = this.nextLine();
      if(st.equals("circle")) {
         sh = this.readCircle();
         this.print("A circle with radius ");
         this.print(sh.getRadius());
      }
      else if(st.equals("square")) {
         sh = this.readSquare();
         this.print("A square with side ");
         this.print(sh.getSide());
      }
      else throw new IOException();
      this.print("Area = ");
      this.print(sh.area());
   }
}
```

Despite method `read()` is type safe, no type can be inferred for `st` to correctly typecheck the method; indeed, when method `area()` in invoked, variable `st` may hold an instance of `Circle` or `Square`, therefore the most precise type that can be correctly assigned to `st` is $Circle \lor Square$. However, if `st` has type $Circle \lor Square$, then both `sh.getRadius()` and `sh.getSide()` do not typecheck.

There are two different kinds of approaches to solve the problem shown above. Either one defines a rather sophisticated flow-sensitive type system able to associate different types with different occurrences of the same variable, or one can typecheck the SSA IR in which the method is compiled.

In an SSA IR the value of each variable is determined by exactly one assignment statement [6]. To obtain this property, a flow graph is built, and a suitable renaming of variables is performed to keep track of the possibly different versions of the same variable; following Singer's terminology [15] we call these versions *virtual registers*. Conventionally, this is achieved by using a different subscript for each virtual register corresponding to the same variable. For instance, in the SSA IR of method `read()` (Figure 1) there are three virtual registers (sh_0, sh_1 and sh_2) for the variable `sh`.
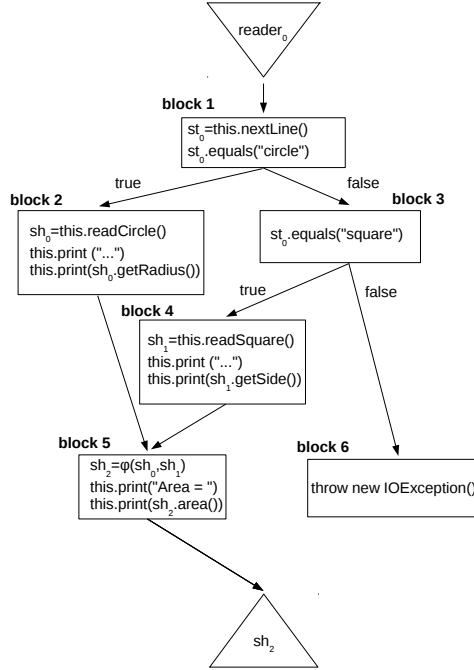
3

Fig. 1. Control flow graph corresponding to the body of method `read`

To transform a program into SSA form, pseudo-functions, conventionally called φ-functions, have to be inserted to correctly deal with merge points. For instance, in block 5 the value of `sh` can be that of either $\mathtt{sh}_0$ or $\mathtt{sh}_1$, therefore a new virtual register $\mathtt{sh}_2$ has to be introduced to preserve the SSA property. The expression $\varphi(\mathtt{sh}_0,\mathtt{sh}_1)$ keeps track of the fact that the value of $\mathtt{sh}_2$ is determined either by $\mathtt{sh}_0$ or $\mathtt{sh}_1$.

At the level of types φ-functions naturally correspond to the union type constructor (Figure 2): arrows correspond to data flow and, as usual, to ensure soundness the type at the origin of an arrow must be a subtype of the type the arrow points to. In the figure, $\tau_0, \tau_1 \leq \tau_0 \vee \tau_1 \leq \tau_2$.

The transformation of a source program into its SSA IR is standard [6], and there exists a quite efficient algorithm to perform it [7], therefore one can simply assume that the program to be analyzed is already in SSA IR. While flow graphs are convenient for generating SSA IR, here we adopt a textual language more suitable for defining an abstract compilation scheme. For instance, the SSA IR of method `read()` is the following:

```
read() {
    b1:{st_0 = this.nextLine();
```

4

τ_0 τ_1

$$\varphi(\tau_0, \tau_1) = \tau_0 \vee \tau_1$$

τ_2

Fig. 2. Type theoretic interpretation of φ-functions

```
        if(st0.equals("circle"))
            jump b2;
        else
            jump b3;}
    b2:{sh0=this.readCircle();
        this.print("A circle with radius ");
        this.print(sh0.getRadius());
        jump b5;}
    b3:{if(st0.equals("square"))
            jump b4;
        else
            jump b6;}
    b4:{sh1=this.readSquare();
        this.print("A square with side ");
        this.print(sh1.getSide());
        jump b5;}
    b5:{sh2=φ(sh0,sh1);
        this.print("Area = ");
        this.print(sh2.area());
        jump out;}
    b6:{throw new IOException();}
    out:{return sh2;}
}
```

The body of a method in IR is a sequence of uniquely labeled blocks; each block ends with either a conditional or unconditional `jump`, a **return** or a **throws**[2]. For simplicity, we require that only the last block[3] contains the **return** statement.

---

[2] In the formal treatment that follows we omit exceptions for simplicity.

[3] This can be always obtained by introducing new virtual registers and inserting a φ-function in case of multiple returned values.

Type analysis with static single information form

Let us consider method `largerThan(sh)` of class `Square`, where **instanceof** is exploited to make the method more efficient in case the parameter `sh` contains an instance of (a subclass) of `Square`.

```
class Square {
...
   largerThan(sh) {
     if(sh instanceof Square)
         return this.side > sh.side;
     else
         return this.area() > sh.area();
   }
}
```

The method is transformed in the following SSA IR:

```
   largerThan(sh₀) {
     b1:{if(sh₀ instanceof Square)
            jump b2;
         else
            jump b3;}
     b2:{r₀=this.side > sh₀.side;
         jump out;}
     b3:{r₁=this.area() > sh₀.area();
         jump out;}
     out:{r₂=φ(r₀,r₁);
         return r₂;}
   }
```

Since variable `sh` is not updated, both blocks `b2` and `b3` refer to the same virtual register sh_0. As a consequence, the only possible type that can be correctly associated with sh_0 is `Square`, thus making the method of little use. However, this problem can be encompassed if one considers the SSI IR of the method [1, 15].

```
   largerThan(sh₀) {
     b1:{if(sh₀ instanceof Square) with (sh₁,sh₂) = σ(sh₀)
                                        (this₁,this₂) = σ(this₀)
            jump b2;
         else
            jump b3;}
     b2:{r₀=this₁.side > sh₁.side;
         jump out;}
     b3:{r₁=this₂.area() > sh₂.area();
         jump out;}
     out:{r₂=φ(r₀,r₁);
         return r₂;}
   }
```

SSI is an extension of SSA enforcing the additional constraint that all variables must have different virtual registers in the branches of conditional expressions;

such a property is obtained by a suitable renaming and by the insertion of σ-functions at split points. As a consequence, suitable virtual registers and σ-functions have to be introduced also for the read-only pseudo-variable **this**.


Fig. 3. Type theoretic interpretation of σ-functions

The notion of σ-function is the dual of φ-function (Figure 3); the type theoretic interpretation of σ depends on the specific kind of conditional context. If such a context is of the form (`sh`$_0$ **instanceof** `Square`) as in the example, then σ splits the type τ_0 of `sh`$_0$ in the type $\tau_0 \wedge$`Square`, assigned to `sh`$_1$, and in the type $\tau_0 \setminus$`Square`, assigned to `sh`$_2$, where the intersection and the complement operators have to be properly defined (see Section 4). For instance, if `sh`$_0$ has type `Square`$\vee$`Circle`, then `sh`$_1$ has type (`Square`$\vee$`Circle`)$\wedge$`Square`=`Square`, and `sh`$_2$ has type (`Square`$\vee$`Circle`)$\setminus$`Square`=`Circle`, therefore `Square`$\vee$`Circle` turns out to be a valid type for the parameter `sh`$_0$ of the method `largerThan`. For what concerns **this**, in this particular example no real split is performed: if we assume that **this**$_0$ has type `Square`, then `Square` is split in (`Square`,`Square`), and both **this**$_1$ and **this**$_2$ have type `Square`.

3   Language definition

In this section we formally define an SSI IR for a simple dynamically typed Java-like language equipped with an **instanceof** operator for performing runtime typechecking.

A program is a collection of class declarations followed by a main expression e with no free variables. The notation $\overline{cd}^n$ is a shortcut for $cd_1, \ldots, cd_n$. A class declares its direct superclass (only single inheritance is supported), its fields, and its methods. Two predefined classes are available: *Object*, the usual root class of the inheritance tree, and *Bool*, the class of the two literals `false` and `true`; such a class cannot be extended, and does not provide any constructor.

7

$$prog ::= \overline{cd}^n \ e$$

$$cd ::= \texttt{class } c_1 \texttt{ extends } c_2 \ \{ \ \overline{f}^n \ \overline{md}^k \ \} \quad (c_1 \neq Object, Bool, c_2 \neq Bool)$$

$$md ::= m(\overline{x}^n) \ \{\overline{b}^n\}$$

$$b ::= l{:}e$$

$$r ::= x_i$$

$$e ::= r \mid \texttt{false} \mid \texttt{true} \mid \texttt{new } c(\overline{e}^n) \mid e.f \mid e_0.m(\overline{e}^n) \mid e_1; e_2 \mid r = e \mid e_1.f = e_2 \mid \texttt{jump } l \mid r = \varphi(\overline{r}^n)$$

$$\mid \texttt{if } (r \texttt{ instanceof } c) \texttt{ with } \overline{(r', r'') = \sigma(r''')}^n \texttt{ jump } l_1 \texttt{ else jump } l_2 \mid \texttt{return } r$$

Syntactic assumptions: inheritance is not cyclic, method bodies are in correct SSI form and are terminated with a unique **return** statement, method and class names are disjoint, no name conflicts in class, field, method and parameter declarations, **new** $c(\overline{e}^n)$ allowed only if $c \neq Bool$, and declared parameters cannot be **this**.

Fig. 4. SSI intermediate language

Every class, except *Bool*, comes equipped with the implicit primary constructor with parameters corresponding to all fields, in the same order as they are inherited and declared, but for simplicity no user declared constructors can be added.

Method bodies are sequences of uniquely labeled blocks that contains sequences of expressions. We assume that all blocks contain exactly one jump, necessarily placed at the end of the block. Three different kinds of jumps are considered: local unconditional and conditional jumps, and returns from methods. Method bodies are implicitly assumed to be in correct SSI IR: each virtual register is determined by exactly one assignment statement, and all variables must have different virtual registers in the branches of conditional expressions. Finally, all method bodies contain exactly one return expression, which is always placed at the end of the body.

The receiver object can be referred inside method bodies with the special implicit parameter **this**; besides usual statements and expressions, we consider φ and σ pseudo-function assignments.

In $r = \varphi(\overline{r}^n)$, $n \geq 2$ and all virtual registers $\overline{r}^n$ occurring in φ are assumed to refer to the same variable, denoted by $var(r_1) = \ldots = var(r_n)$.

All σ-functions are used in association with conditional jumps; each virtual register r occurring in either branches has to be split into two new distinct versions used in the blocks labeled by $\mathtt{l}_1$, and $\mathtt{l}_2$, respectively. The conditions we consider are only of the form $(r \texttt{ instanceof } c)$ since our aim is studying how type analysis can be enhanced by SSI in the presence of explicit runtime typechecks; however, more elaborated forms of conditions can be expressed in terms of the more primitive form $(r \texttt{ instanceof } c)$ by simple transformations performed by the compiler front-end. For instance, the statement

```
if (x.m1() instanceof A && y.m2() instanceof B) e₁ else e₂
```

can be transformed in the equivalent SSI IR

```
b0  : {z₀=xₖ.m1();
        if (z₀ instanceof A) with ... jump b1 else jump b3;}
```

```
b1 : {w₀=yₙ.m2();
      if (w₀ instanceof B) with ... jump b2 else jump b3;}
b2 : {e'₁}
b3 : {e'₂}
```

where z_0 and w_0 are fresh, e'_1 and e'_2 are the SSI IRs of e_1 and e_2, respectively, and σ-functions assignments (that depend on e_1 and e_2) have been omitted. Depending on the types and abstract compilation scheme under consideration, there could be other kinds of conditions for which SSI would improve type analysis; for instance, if an abstract compilation scheme allows analysis of null references, such an analysis could be enhanced by SSI in the case of conditional expressions with conditions of the form (x == **null**). On the other hand, for conditions of the form (x_1 < x_2) SSI does not help refine type analysis as long as the abstract compilation scheme maps numeric values to the standard primitive types **int**, **float**, and **double**; in this case the type theoretic interpretation of σ-functions is the loosest one: $\sigma(\tau)$=(τ,τ) (hence, no split is actually performed).

Semantics: To define the small step semantics of the language we first need to specify values v (see Figure 5), which are either the literals `false` and `true` (recall that they are predefined instances of the *Bool* class) or identities o of dynamically created objects. Furthermore, we add *frame expressions* $ec\{e\}$, where ec is an execution context; frame expressions are *runtime expressions*, that is, expressions that represent intermediate evaluation steps and that are needed for defining the small step semantics of method calls. An execution context ec is a pair consisting of a stack frame fr and a code address a. A frame expression $\langle fr, a \rangle \{e\}$ corresponds to the execution of a call to a method m declared in class c, where e is the residual expression (yet to be evaluated) of the currently executed block, fr is the stack frame of the method call, $a = \mu.l$ is the address of the current block, where $\mu = c.m$ is the fully qualified name of the method, and l is the label of the current block.

Stack frames fr map variables and virtual registers to their corresponding values. These frames are represented by a pair of lists of associations, $x \mapsto v$ and $r \mapsto v$, where variables and virtual registers are all distinct. Keeping track of the values of both virtual registers and their associated variable allows for a simple semantics, as discussed below.

Heaps $\mathcal{H}$ map object identifiers o to objects, that is, pairs consisting of a class name c and the set of field names f with their corresponding value v.

Execution is modeled by the judgment $\mathcal{H} \vdash e \rightarrow \mathcal{H}', e'$, to be read: "runtime expression e, in a heap $\mathcal{H}$, is reduced in one step into runtime expression e', producing the heap $\mathcal{H}'$". The reduction arrow, and all auxiliary functions, should be parametrized by the whole executing program, $\overline{cd}^n$, which is kept implicit.

The execution inside a stack frame, that is, the execution of the body of a method, is modeled by the judgement $\mathcal{H}, ec \vdash e \rightarrow \mathcal{H}', ec', e'$ to be read: "runtime expression e, in an execution context ec and heap $\mathcal{H}$, is reduced in one step into runtime expression e', producing an execution context ec' and heap $\mathcal{H}'$".

Figure 6 show the execution rules.

Rule (main) is the one that allows to start the execution of a program, that is, a sequence of class declarations $\overline{cd}^n$ and a main expression e. The premise of the rule exploits the small-step reduction, starting from an empty heap $\epsilon_{\mathcal{H}}$ and an empty stack frame ϵ, annotated by the pair $\langle \bot, \bot \rangle$, which corresponds to the fact that the main expression is not actually contained in any block/method.

The relations $\rightarrow^*$ and $\rightarrow^*_{max}$ denote the transitive and the maximal transitive closure, respectively, of $\rightarrow$: $\mathcal{H} \vdash e \rightarrow^* \mathcal{H}', e'$ iff $\mathcal{H} \vdash e \rightarrow \mathcal{H}'', e''$ and $\mathcal{H}'' \vdash e'' \rightarrow^* \mathcal{H}', e'$, and $\mathcal{H} \vdash e \rightarrow^*_{max} \mathcal{H}', e'$ iff $\mathcal{H} \vdash e \rightarrow^* \mathcal{H}', e'$ and there exist no $\mathcal{H}''$, e'' s.t. $\mathcal{H}' \vdash e' \rightarrow \mathcal{H}'', e''$.

Rule (ctx-closure) performs a single step inside the most nested execution context, extracted and updated by the two auxiliary functions *currExecContext* and *updCurrExecContext*, respectively.[4]

Rule (return) models the return to the caller, by destroying (popping), the current execution context and substituting the whole frame expression with the value of the returned virtual register r.

Rule (x-acc) models the access to a virtual register (`this` is considered a special read-only local variable), by simply extracting the value of the corresponding variable from the stack frame fr; this works because of the way the assignment is handled in Rule (var-asn).

Variable and field assignments evaluate to their right values; rule (var-asn) models virtual register assignments, and has the side effect of updating, in the current stack frame fr, the values of both the virtual register r and its associated variable x. This dual assignment resembles a cache write-through strategy, and makes evaluation of φ-function simpler (see comments below).

Rule (fld-asn) models field assignments; in that case, the object referenced by o is retrieved from the heap, and its value updated.

Rule (meth-call), describing method invocations, is a bit more involved; first of all, the object referenced by o is retrieved in order to find its class, c. Then, using the two auxiliary functions *firstBlock* and *params*, we obtain the first block of the method and its parameter names. The result of the evaluation is a frame expression, where the stack frame is built by assigning the actual parameters to the formal ones, and the reference o to `this`, and the code address is the fully qualified name of the invoked method, $c.m$, plus the label of its first block, l. Finally, the resulting runtime expression is the body of the first block.

Rule (seq) models the fact that a value is discarded when followed by another.

Rule (phi) models the assignment of a phi-function, which accesses to a (set $\overline{r}^n$ of different versions of the same) local variable x, by assigning the value contained in x (this is correct because of the write-through semantics of the Rule (var-asn)) to the virtual register r'.

Rule (new) models object creations; a new object, identified by a fresh reference o, is added to the heap $\mathcal{H}$. The fields $\overline{f}^n$ of the newly created object are initialized by the value passed to the constructor.

---

[4] The trivial definition of these and the following auxiliary functions have been omitted.

Rule (fld-acc) models field accesses; its evaluation is quite trivial: the object is retrieved from the heap, and the resulting expression is the value of the selected field.

Rules (jump) and (if) model unconditional and conditional jumps, respectively. These are the only rules that modify the label-part of the execution context. The evaluation of a jump, which, by construction, is known to be the last expression of a sequence, corresponds to replacing the jump expression itself with the expression e contained in the block labeled l' and updating the stack frame annotation accordingly.

The conditional jump, modeled by (if), depending on whether the value of the register r (contained in $fr(r)$) is a reference to an object of a subclass of c, has to both choose which branch to execute and which virtual registers have to be updated. If the referenced object is indeed an instance of c, then the target label is l_1 and the virtual registers $\overline{r'}^n$ are updated; otherwise, the target label is l_2 and the virtual registers $\overline{r''}^n$ are updated.

$$
\begin{aligned}
v &::= \texttt{false} \mid \texttt{true} \mid o \quad \text{(values)} \\
e &::= ec\{e\} \mid \ldots \quad \text{(runtime expressions)} \\
ec &::= \langle fr, a \rangle \quad \text{(execution context)} \\
fr &::= \overline{x \mapsto v}^j \, \overline{r \mapsto v}^k \quad \text{(stack frames)} \\
a &::= \mu.l \quad \text{(block addresses, where } \mu \text{ are full method names)} \\
\mathcal{H} &::= \overline{o \mapsto \langle c, \overline{f \mapsto v}^j \rangle}^k \quad \text{(heaps)} \\
\mathcal{C}[\cdot] &::= [\cdot] \mid ec\{\mathcal{C}[\cdot]\} \mid \texttt{new } c(\overline{v}^n, \mathcal{C}[\cdot], \overline{e}^j) \mid \mathcal{C}[\cdot].f \mid \mathcal{C}[\cdot].m(\overline{e}^k) \mid v_0.m(\overline{v}^j, \mathcal{C}[\cdot], \overline{e}^k) \\
&\quad \mid \mathcal{C}[\cdot]; e \mid v; \mathcal{C}[\cdot] \mid x = \mathcal{C}[\cdot] \mid \mathcal{C}[\cdot].f = e \mid v.f = \mathcal{C}[\cdot] \\
&\quad \mid \texttt{if } (\mathcal{C}[\cdot]) \texttt{ with } \overline{(x', x'') = \sigma(x''')}^n \texttt{ jump } l_1 \texttt{ else jump } l_2
\end{aligned}
$$

Fig. 5. Syntactic definitions instrumental to the operational semantics

4 Abstract compilation

In this section we define an abstract compilation scheme for programs in the SSI IR presented in Section 3. Programs are translated into a Horn formula *Hf* (that is, a logic program) and a goal B; type analysis amounts to resolve B w.r.t. the coinductive semantics (that is the greatest Herbrand model) of *Hf* [4].

In previous work we have defined compilation schemes based on union and structural object types, to support parametric and data polymorphism, (that is, polymorphic methods and fields) that allow quite precise type analysis, but pose problems in terms of termination of the resolved goals. In this paper we present a simpler compilation scheme based on union and nominal object types, that allows a less precise type analysis on the one hand (but still more precise than that obtained with the standard type systems of mainstream languages as Java and C#) but, on the other hand, avoids non termination problems with goal resolution, since the space of all possible valid types for a given program is

$$(\text{main})\ \frac{\epsilon_{\mathcal{H}} \vdash \langle \epsilon_{fr}, \bot.\bot \rangle \{e\} \rightarrow^{*}_{max} \mathcal{H}, e'}{\overline{cd}^{n}\ e \Rightarrow e'}$$

$$(\text{ctx-closure})\ \frac{\begin{array}{c} currExecContext(\mathcal{C}[\cdot]) = ec \\ \mathcal{H}, ec \vdash e \rightarrow \mathcal{H}', ec', e' \\ \mathcal{C}'[\cdot] = updCurrExecContext(\mathcal{C}[\cdot], ec') \end{array}}{\mathcal{H} \vdash \mathcal{C}[e] \rightarrow \mathcal{H}', \mathcal{C}'[e']}$$

$$(\text{return})\ \frac{}{\mathcal{H} \vdash \mathcal{C}[\langle fr, a \rangle \{\texttt{return}\ r\}] \rightarrow \mathcal{H}, \mathcal{C}[fr(r)]}$$

$$(\text{x-acc})\ \frac{}{\mathcal{H}, \langle fr, a \rangle \vdash r \rightarrow \mathcal{H}, \langle fr, a \rangle, fr(r)}$$

$$(\text{new})\ \frac{\begin{array}{c} o\ \text{fresh in}\ \mathcal{H} \\ fieldNames(c) = \overline{f}^{n} \end{array}}{\mathcal{H}, ec \vdash \texttt{new}\ c(\overline{v}^{n}) \rightarrow \mathcal{H}[\langle c, \overline{f \mapsto v}^{n} \rangle / o], ec, o}$$

$$(\text{fld-acc})\ \frac{\mathcal{H}(o) = \langle c, \overline{f \mapsto v}^{n} \rangle \quad f = f_{j}}{\mathcal{H}, ec \vdash o.f \rightarrow \mathcal{H}, ec, v_{j}}$$

$$(\text{meth-call})\ \frac{\begin{array}{c} \mathcal{H}(o) = \langle c, \_ \rangle \\ firstBlock(c.m) = l : e \\ params(c.m) = \overline{r}^{n} \end{array}}{\begin{array}{c} \mathcal{H}, ec \vdash o.m(\overline{v}^{n}) \rightarrow \\ \mathcal{H}, ec, \langle (\overline{r \mapsto v}^{n}, \texttt{this}_{0} \mapsto o), c.m.l \rangle \{e\} \end{array}}$$

$$(\text{seq})\ \frac{}{\mathcal{H}, ec \vdash v_{1}; v_{2} \rightarrow \mathcal{H}, ec, v_{2}}$$

$$(\text{var-asn})\ \frac{x = var(r) \quad x \neq \texttt{this}}{\mathcal{H}, \langle fr, a \rangle \vdash r = v \rightarrow \mathcal{H}, \langle fr[v/r, v/x], a \rangle, v}$$

$$(\text{fld-asn})\ \frac{\begin{array}{c} \mathcal{H}(o) = \langle c, \overline{f \mapsto v}^{n} \rangle \\ f = f_{j} \quad \text{if}\ i = j\ \text{then}\ v'_{i} = v\ \text{else}\ v'_{i} = v_{i} \end{array}}{\mathcal{H}, ec \vdash o.f = v \rightarrow \mathcal{H}[\langle c, \overline{f \mapsto v'}^{n} \rangle / o], ec, v}$$

$$(\text{jump})\ \frac{block(\mu.l') = l' : e}{\mathcal{H}, \langle fr, \mu.l \rangle \vdash \texttt{jump}\ l' \rightarrow \mathcal{H}, \langle fr, \mu.l' \rangle, e}$$

$$(\text{phi})\ \frac{v = fr(var(r_{1}))}{\mathcal{H}, \langle fr, a \rangle \vdash r' = \varphi(\overline{r}^{n}) \rightarrow \mathcal{H}, \langle fr[v/r'], a \rangle, v}$$

$$(\text{if})\ \frac{\begin{array}{c} \mathcal{H}(fr(r)) = \langle c', \_ \rangle \\ \text{if}\ c' \leq c\ \text{then}\ \ l' = l_{1}, fr' = fr[\overline{fr(r''')/r'}^{n}] \\ \text{else}\ \ \ l' = l_{2}, fr' = fr[\overline{fr(r''')/r''}^{n}] \\ block(\mu.l') = l' : e \end{array}}{\mathcal{H}, \langle fr, \mu.l \rangle \vdash \texttt{if}\ (r\ \texttt{instanceof}\ c)\ \texttt{with}\ \overline{(r', r'') = \sigma(r''')}^{n}\ \texttt{jump}\ l_{1}\ \texttt{else}\ \texttt{jump}\ l_{2} \rightarrow \mathcal{H}, \langle fr', \mu.l' \rangle, e}$$

Fig. 6. Small-step semantics

12

always bounded. In this way there is no need to resort to coinductive constraint logic programming [3] (where constraints are expressed in terms of the subtyping relation) as required for structured types to ensure termination of goal resolution. Therefore, to make the presentation simpler, here subtyping is treated as an ordinary predicate, rather than as a constraint, even though switching from ordinary to constraint logic programs would allow more efficient goal resolution.

Summarizing, here we prefer using nominal rather than structured types for the following reasons: SSI allows more precise type analysis even in presence of less expressive types, the presentation is simpler, and, last but not least, we take advantage of the *plug-and-play* facility offered by the abstract compilation approach by providing yet another compilation scheme; in practice, more advanced compilations schemes could be adopted, including structural [4, 5] and exception types [3] (see further comments in the conclusion).

The compilation of programs, class, and method declarations is defined in Figure 7. We follow the usual syntactic conventions for logic programs: logic variable names begin with upper case, whereas predicate and functor names begin with lower case letters. Underscore denotes anonymous logic variables that occur only once in a clause; $[\,]$ and $[e|l]$ respectively represent the empty list, and the list where e is the first element, and l is the rest of the list.

$$(\text{prog})\frac{\forall\, i = 1..n\ cd_i \rightsquigarrow Hf_i \qquad e \rightsquigarrow (t\mid B)}{\overline{cd}^n\ e \rightsquigarrow (Hf^d \cup \overline{Hf}^n \mid B)}$$

$$(\text{class})\frac{\forall\, i = 1..k\ md_i\ \textbf{in}\ c_1 \rightsquigarrow Hf_i \qquad inhFields(c_1) = \overline{f'}^h}{\texttt{class}\ c_1\ \texttt{extends}\ c_2\ \{\ \overline{f}^n\ \overline{md}^k\ \} \rightsquigarrow}$$

$$\overline{Hf}^k \cup \left\{ \begin{array}{l} class(c_1) \leftarrow true. \\ extends(c_1, c_2) \leftarrow true. \\ \overline{dec\_field(c_1, f)}^n \leftarrow true. \\ new(CE, c_1, [\overline{T'}^h, \overline{T}^n]) \leftarrow new(CE, c_2, [\overline{T'}^h]), \overline{field(CE, c_1, f, T)}^n. \end{array} \right\}$$

$$(\text{meth})\frac{\overline{b}^n \rightsquigarrow (t\mid B)}{m(\overline{r}^n)\{\overline{b}^n\}\ \textbf{in}\ c \rightsquigarrow}$$

$$\left\{ \begin{array}{l} dec\_meth(c, m) \leftarrow true. \\ has\_meth(CE, c, m, [This_0, \overline{r}^n], t) \leftarrow subclass(This_0, c), B. \end{array} \right\}$$

$$(\text{body})\frac{\forall\, i = 1..n\ b_i \rightsquigarrow B_i}{\overline{b}^n\ l\text{:}\texttt{return}\ r \rightsquigarrow (r\mid \overline{B}^n)}$$

Fig. 7. Compilation of programs, class, and method declarations and bodies.

Each rule defines a different compilation judgment. The judgment $\overline{cd}^n\ e \rightsquigarrow (Hf^d \cup \overline{Hf}^n \mid B)$ means that the program $\overline{cd}^n\ e$ is compiled into the pair $(Hf^d \cup \overline{Hf}^n \mid B)$, where $Hf^d \cup \overline{Hf}^n$ is a Horn formula (that is, a set of Horn clauses), and

B is a goal[5]. Static analysis of the program corresponds to resolving the goal B w.r.t. the coinductive semantics of $Hf^d \cup \overline{Hf}^n$. The Horn formula Hf^d contains all generated clauses that are invariant w.r.t. the program (see Figure 9), whereas each Hf_i is obtained by compiling the class declaration cd_i (see below); the goal B is obtained from the compilation of the main expression e (see below); the term t corresponding to the returned type of e is simply ignored here, but it is necessary for compiling expressions (see comments on Figure 8).

The compilation of a class declaration class c_1 extends c_2 $\{ \overline{f}^n \ \overline{md}^k \}$ is a set of clauses, including each clause Hf_i obtained by compiling the method md_i (see below), a clause asserting that class c_1 declares field f_i, for all $i = 1..n$, and three specific clauses for predicates *class*, *extends*, and *new*. The clause for *new* deserves some explanations: the atom $new(ce, c, [\overline{t}^n])$ succeeds iff the invocation of the implicit constructor of c with n arguments of type $\overline{t}^n$ is type safe in the global class type environment ce. The class environment ce is required for compiling field access and update expressions (see Figure 8): it is a finite map (simply represented by a list) associating class names with field records, which are finite maps (again simply represented by lists) associating all fields of a class with their corresponding types. Class environments are required because of nominal types: abstract compilation with structural types allows data polymorphism on a per-object basis, whereas here we obtain only a very limited form of data polymorphism on a per-class basis. Type safety of object creation is checked by ensuring that object creation for the direct superclass c_2 is correct, where only the first h arguments corresponding to the inherited fields (computed by the auxiliary function *inhFields* whose straightforward definition has been omitted) are passed; then, predicate *field* defined in Figure 9 checks that all remaining n arguments, corresponding to the new fields declared in c_1, have types that are compatible with those specified in the class environment. The clause dealing with the base case for the root class *Object* is defined in Figure 9.

The judgment $m(\overline{r}^n)\{\overline{b}^n\}$ **in** $c \rightsquigarrow Hf$ means that the method declaration $m(\overline{r}^n)\{\overline{b}^n\}$ contained in class c compiles to Horn clauses Hf. Just two clauses are generated per method declaration: the first simply states that method m is declared in class c (and is needed to deal with inherited methods, see Figure 9), whereas the second is obtained by compiling the body of the method. The atom $has\_meth(ce, c, m, [t_0, \overline{t}^n], t)$ succeeds iff in class environment ce method m of class c can be safely invoked on target object of type t_0, with n arguments of type $\overline{t}^n$ and returned value of type t. The predicate *subclass* (defined in Figure 9) ensures that the method can be invoked only on objects that are instances of c or of one of its subclasses. For simplicity we assume that all names (including this) are translated to themselves, even though, in practice appropriate injective renaming should be applied [4].

---

[5] For simplicity we use the same meta-variable B to denote conjunctions of atoms (that is, clause bodies), and goals, even though more formally goals are special clauses of the form $false \leftarrow B$.

The compilation of a method body $\overline{b}^{\,n}$ $l{:}\texttt{return}$ r consists of the type of the returned virtual register r, and the conjunction of all the atoms generated by the compilation of blocks $\overline{b}^{\,n}$.

Figure 8 defines abstract compilation for blocks, and expressions.

$$(\text{block})\frac{e \rightsquigarrow (t \mid B)}{l{:}e \rightsquigarrow B} \qquad (\text{seq})\frac{e_1 \rightsquigarrow (t_1 \mid B_1) \qquad e_2 \rightsquigarrow (t_2 \mid B_2)}{e_1; e_2 \rightsquigarrow (t_2 \mid B_1, B_2)} \qquad (\text{jmp})\frac{}{\texttt{jump } l \rightsquigarrow (void \mid true)}$$

$$(\text{c-jmp})\frac{\begin{array}{c}\text{if } var(r_i''') = var(r)\\ \text{then } t_i' = T, t_i'' = F\\ \text{else } t_i' = r_i''', t_i'' = r_i''' \qquad T, F \text{ fresh}\end{array}}{\texttt{if } (r \texttt{ instanceof } c) \texttt{ with } \overline{(r', r'') = \sigma(r''')}^{\,n} \texttt{ jump } l_1 \texttt{ else jump } l_2 \texttt{ in } \rightsquigarrow \\ (void \mid inter(r, c, T), diff(r, c, F), \overline{var\_upd(r', t'), var\_upd(r'', t'')}^{\,n})}$$

$$(\text{var-upd})\frac{e \rightsquigarrow (t \mid B)}{r = e \rightsquigarrow (t \mid B, var\_upd(r, t))}$$

$$(\text{field-upd})\frac{e_1 \rightsquigarrow (t_1 \mid B_1) \qquad e_2 \rightsquigarrow (t_2 \mid B_2)}{e_1.f = e_2 \rightsquigarrow (t_2 \mid B_1, B_2, field\_upd(CE, t_1, f, t_2))}$$

$$(\text{phi})\frac{}{r = \varphi(\overline{r}^{\,n}) \rightsquigarrow (\overline{\vee r}^{\,n} \mid var\_upd(r, \overline{\vee r}^{\,n}))}$$

$$(\text{new})\frac{\forall i = 1..n \ e_i \rightsquigarrow (t_i \mid B_i)}{\texttt{new } c(\overline{e}^{\,n}) \rightsquigarrow (c \mid \overline{B}^{\,n}, new(CE, c, [\overline{t}^{\,n}]))}$$

$$(\text{field-acc})\frac{e \rightsquigarrow (t \mid B) \qquad R \text{ fresh}}{e.f \rightsquigarrow (R \mid B, field(CE, t, f, R))}$$

$$(\text{invk})\frac{\forall i = 0..n \ e_i \rightsquigarrow (t_i \mid B_i) \qquad R \text{ fresh}}{e_0.m(\overline{e}^{\,n}) \rightsquigarrow (R \mid B_0, \overline{B}^{\,n}, invoke(CE, t_0, m, [\overline{t}^{\,n}], R))}$$

$$(\text{var})\frac{}{r \rightsquigarrow (r \mid true)} \qquad (\text{bool})\frac{v \in \{\texttt{true}, \texttt{false}\}}{v \rightsquigarrow (bool \mid true)}$$

Fig. 8. Compilation of blocks and expressions

Compiling a block $l{:}e$ returns the conjunction of atoms obtained by compiling e; the type t of e is discarded.

The compilation of $e_1; e_2$ returns the type of e_2 and the conjunction of atoms generated from the compilation of e_1 and e_2.

The compilation of an unconditional jump generates the type $void$ and the empty conjunction of atoms $true$. A conditional jump with condition $(r \texttt{ instanceof } c)$ has type $void$ as well, but a non empty sequence of predicates is generated to deal with the splitting performed by σ-functions; predicates $inter$ and $diff$ (defined in Figure 10) correspond to a sort of intersection T and difference F between the type of r and c, respectively, and predicate var_upd (defined in Figure 9) ensures that the type of virtual registers r_i' and r_i'' are compatible with the pairs

of types returned by the σ-functions. In case r_i''' refers to the same variable of r the types of such a pair are the computed intersection T and difference F, respectively, otherwise the pair (r_i''', r_i''') is returned (hence, no split is actually performed).

Compilation of assignments to virtual registers and fields yields the conjunction of the atoms generated from the corresponding sub-expressions, together with the atoms that ensure that the assignment is type compatible (with predicates *var_upd* and *field_upd* defined in Figure 9). The returned type is the type of the right-hand side expression.

Compilation of φ-function assignments to virtual registers is just an instantiation of rule (var-upd) where the type of the expression is the union of the types of the virtual registers passed as arguments to φ.

Compilation rules for object creation, field selection, and method invocation follow the same pattern: the type of the expression is a fresh logical variable corresponding to the type returned by the specific predicate (*new*, *field*, and *invoke* defined in Figure 9). The generated atoms are those obtained from the compilation of the sub-expressions, together with the atom specific of the expression.

Rules (var) and (bool) are straightforward.

Figure 9 and Figure 10 define the set Hf^d used in compilation rule (prog) corresponding to all generated clauses that are invariant w.r.t. the compiled program.

Clauses in Figure 9 deserve some comments for what concerns the subtyping relation (predicate *subtype*); as expected, classes c_1 and c_2 are both subtypes of $c_1 \vee c_2$, but no subclass of c_1 or c_2 is a subtype of $c_1 \vee c_2$, because subclassing is not subtyping, since no rules are imposed on method overriding. Consider for instance the following source code snippet:

```
class Square {
...
equals(s) {return this.side==s.side;}
...
}
class ColoredSquare extends Square {
...
equals(cs) {return this.side==cs.side && this.color==cs.color;}
...
}
```

According to our compilation scheme, the expression `s1.equals(s2)` has type `Bool` if `s1` and `s2` have type `Square` and `Square∨ColoredSquare`, respectively, but the same expression is not well-typed if `s1` has type `ColoredSquare` (hence, `ColoredSquare⊀Square`), since `s2` may contain an instance of `Square` for which field `color` is not defined. Subtyping is required for defining the predicates *var_upd* and *field_upd* for virtual register and field updates: the type of the source must be a subtype of the type of the destination.

Predicate *field* looks up the type of a field in the global class environment, and is defined in terms of the auxiliary predicates *has_field*, *class_fields* and *field_type*.

$class(object) \leftarrow true.$
$class(bool) \leftarrow true.$
$extends(bool, object) \leftarrow true.$
$subclass(X, X) \leftarrow class(X).$
$subclass(X, Y) \leftarrow extends(X, Z), subclass(Z, Y).$
$subtype(T, T) \leftarrow true.$
$subtype(T1 \vee T2, T) \leftarrow subtype(T1, T), subtype(T2, T).$
$subtype(T, T1 \vee T2) \leftarrow subclass(T, T1).$
$subtype(T, T1 \vee T2) \leftarrow subclass(T, T2).$
$field(CE, C, F, T) \leftarrow has\_field(C, F), class\_fields(CE, C, R), field\_type(R, F, T).$
$field(CE, T1 \vee T2, F, FT1 \vee FT2) \leftarrow field(CE, T1, F, FT1), field(CE, T2, F, FT2).$
$class\_fields([C : R|\_], C, R) \leftarrow true.$
$class\_fields([C1 : \_|CE], C2, R) \leftarrow class\_fields(CE, C2, R), C1 \neq C2.$
$field\_type([F{:}T|\_], F, T) \leftarrow true.$
$field\_type([F1 : \_|R], F2, T) \leftarrow field\_type(R, F2, T), F1 \neq F2.$
$invoke(CE, C, M, A, RT) \leftarrow has\_meth(CE, C, M, [C|A], RT).$
$invoke(CE, T1 \vee T2, M, A, RT) \leftarrow invoke(CE, T1, M, A, RT1),$
$\qquad\qquad\qquad\qquad\qquad\qquad invoke(CE, T2, M, A, RT2), subtype(RT1 \vee RT2, RT).$
$new(\_, object, [\,]) \leftarrow true.$
$has\_field(C, F) \leftarrow dec\_field(C, F).$
$has\_field(C, F) \leftarrow extends(C, P), has\_field(P, F), \neg dec\_field(C, F).$
$has\_meth(CE, C, M, A, R) \leftarrow extends(C, P), has\_meth(CE, P, M, A, R), \neg dec\_meth(C, M).$
$var\_upd(T1, T2) \leftarrow subtype(T2, T1).$
$field\_upd(CE, C, F, T2) \leftarrow field(CE, C, F, T1), subtype(T2, T1).$
$field\_upd(CE, T1 \vee T2, F, T) \leftarrow field\_upd(CE, T1, F, T), field\_upd(CE, T2, F, T).$

Fig. 9. Clauses defining the predicates used by the abstract compilation

In particular, predicate *class_fields* checks that a class has actually a certain field, either declared or inherited. The definitions of *class_fields* and *field_type* are straightforward, whereas the clause for *has_field* dealing with inherited fields is similar to the corresponding one for *invoke* (see below).

If the target object has a union type, predicate *invoke* checks that method invocation is correct for both types of the union, and then merges the types of the results into a single union type. If the target object has a class type c, then the correctness of method invocation is checked with predicate *has_meth* applied to class c and to the same list of arguments where, however, the type c of **this** is added at the beginning.

Finally, the clause for *has_meth* deals with the inherited methods: if class c does not declare method m, then *has_meth* must hold on the direct superclass of c.

$inter(C1, C2, C1) \leftarrow subclass(C1, C2).$
$inter(T1 \vee T2, C, IT1 \vee IT2) \leftarrow inter(T1, C, IT1), inter(T2, C, IT2).$
$inter(T1 \vee T2, C, IT1) \leftarrow inter(T1, C, IT1), \neg inter(T2, C, \_).$
$inter(T1 \vee T2, C, IT2) \leftarrow inter(T2, C, IT2), \neg inter(T1, C, \_).$
$diff(C1, C2, C1) \leftarrow class(C1), \neg subclass(C1, C2).$
$diff(T1 \vee T2, C, IT1 \vee IT2) \leftarrow diff(T1, C, IT1), diff(T2, C, IT2).$
$diff(T1 \vee T2, C, IT1) \leftarrow diff(T1, C, IT1), \neg diff(T2, C, \_).$
$diff(T1 \vee T2, C, IT2) \leftarrow diff(T2, C, IT2), \neg diff(T1, C, \_).$

Fig. 10. Clauses defining type intersection and difference

Predicates *inter* and *diff* define type splitting for σ-functions; the asymmetric definition for *inter* is due to the fact that subclass is not subtyping: if r has type c_1, then it means that it contains an object that is an instance of c_1, therefore the condition (r **instanceof** c_2) is false when c_2 is a proper subclass of c_1. Both predicates fail if the returned type is empty, therefore a conditional jump is not considered correct if either branches are dead (that is, not reachable). In practice, it would be better to avoid this kind of failures by introducing an explicit empty type constant to be able to detect dead code without any failure. Such an alternative option does not pose any technical issue, but since it is more verbose (a new clause dealing with the empty type must be added for most predicates) has not been considered here, just for space limitations.

5 Conclusion

We have defined the small step operational semantics of a simple Java-like language in SSI IR, equipped with an **instanceof** operator for runtime typechecks, and shown how precise type analysis of branches guarded by runtime typechecks can be achieved by abstract compilation in the presence of union and nominal types, and by suitably defining two predicates *inter* and *diff* that provide the type theoretic interpretation of σ-functions.

Despite the use of nominal types, the analysis is more precise than that we would get from the type system of a statically typed language as Java and C#; however, using structural types to trace the type of each field object leads to a more precise analysis, as already shown in previous work [4, 5]. Here we have preferred to keep the presentation simpler, but we envisage no problems in extending the compilation scheme and the definition of the predicates *inter* and *diff* to accommodate structural types.

For what concerns future developments, the approach proposed here seems promising also for precise static analysis of null pointer references when branches are guarded by conditions of the form (r == **null**).

References

1. C. S. Ananian. The static single information form. Technical Report MITLCS-TR-801, MIT, 1999.
2. D. Ancona, M. Ancona, A Cuni, and N. Matsakis. RPython: a Step Towards Reconciling Dynamically and Statically Typed OO Languages. In *OOPSLA 2007 Proceedings and Companion, DLS'07: Proceedings of the 2007 Symposium on Dynamic Languages*, pages 53–64. ACM, 2007.
3. D. Ancona, A. Corradi, G. Lagorio, and F. Damiani. Abstract compilation of object-oriented languages into coinductive CLP(X): can type inference meet verification? In B. Beckert and C. Marché, editors, *Formal Verification of Object-Oriented Software International Conference, FoVeOOS 2010, Paris, France, June 28-30, 2010, **Revised Selected Papers***, volume 6528 of *Lecture Notes in Computer Science*. Springer Verlag, 2011.
4. D. Ancona and G. Lagorio. Coinductive type systems for object-oriented languages. In S. Drossopoulou, editor, *ECOOP 2009 - Object-Oriented Programming*, volume 5653 of *Lecture Notes in Computer Science*, pages 2–26. Springer Verlag, 2009. **Best paper prize**.
5. D. Ancona and G. Lagorio. Idealized coinductive type systems for imperative object-oriented programs. *RAIRO - Theoretical Informatics and Applications*, 45(1):3–33, 2011.
6. R. Cytron, J. Ferrante, B. K. Rosen, M. N. Wegman, and F. K. Zadeck. Efficiently computing static single assignment form and the control dependence graph. *ACM Transactions on Programming Languages and Systems*, 13:451–490, 1991.
7. D. Das and U. Ramakrishna. A practical and fast iterative algorithm for phi-function computation using DJ graphs. *ACM Transactions on Programming Languages and Systems*, 27(3):426–440, 2005.
8. Alpern B. e al. The jalapeño virtual machine. *IBM Systems Journal*, 39(1), 2000.
9. R. Griesemer and S. Mitrovic. A compiler for the java hotspottm virtual machine. In *The School of Niklaus Wirth, "The Art of Simplicity"*, pages 133–152, 2000.
10. G. Holloway. The machine-suif static single assignment library. Technical report, Harvard School of Engineering and Applied Sciences, 2001.
11. D. Novillo. Tree ssa - a new optimization infrastructure for gcc. In *GCC Developers' Summit*, pages 181–193, 2003.
12. L. Simon, A. Bansal, A. Mallya, and G. Gupta. Co-logic programming: Extending logic programming with coinduction. In *Automata, Languages and Programming, 34th International Colloquium, ICALP 2007*, pages 472–483, 2007.

13. L. Simon, A. Mallya, A. Bansal, and G. Gupta. Coinductive logic programming. In *Logic Programming, 22nd International Conference, ICLP 2006*, pages 330–345, 2006.
14. J. Singer. Static single information form in machine suif. Technical report, University of Cambridge Computer Laboratory, UK, 2004.
15. J. Singer. *Static Program Analysis based on Virtual Register Renaming*. PhD thesis, Christs College, 2005.
16. A. Tavares, F.M. Pereira, M. Bigonha, and R. Bigonha. Efficient ssi conversion. In *SBLP 2010*, 2010.
17. J. Winther. Guarded type promotion (eliminating redundant casts in java). In *FTfJP 2011*. ACM, 2011.

A   Generated code

```
class(object).
class(bool).
class(int).
class(string).
class(shape_reader).
class(circle).
class(square).
class(col_square).

extends(bool,object).
extends(int,object).
extends(string,object).
extends(shape_reader,object).
extends(circle,object).
extends(square,object).
extends(col_square,square).

subclass(X,X) :- class(X).
subclass(X,Y) :- extends(X,Z),subclass(Z,Y).

subclasstype([X],Y) :- subclass(X,Y).

subtype(T1,T2) :- (var(T1);var(T2)) -> T1=T2;ord_subset(T1,T2).

field(CE,[C],F,FT) :- has_field(C,F),memberchk(C:R,CE),memberchk(F:FT1,R),subtype(FT1
field(CE,O,F,FT) :- O\=[_],O=[C|T],has_field(C,F),memberchk(C:R,CE),memberchk(F:FT1,F
field(_,[],_,[]).

invoke(CE,[C],M,A,RT) :- has_meth(CE,C,M,[[C]|A],RT1),subtype(RT1,RT).
invoke(CE,O,M,A,RT) :- O\=[_],O=[C|T],has_meth(CE,C,M,[[C]|A],RT1),invoke(CE,T,M,A,RT
invoke(_,[],_,_,[]).

or(T1,T2,T3) :- (T1==[]-> T3=T2,!).
or(T1,T2,T3) :- (T2==[]-> T3=T1,!).
```

```prolog
or(T1,T2,T3) :- ((var(T1);var(T2))) -> T3=T2,T1=T2;ord_union(T1,T2,T3).

new(_,object,[]).
new(CE,shape_reader,[]) :- new(CE,object,[]).
new(CE,circle,[R0]) :- new(CE,object,[]),field(CE,[circle],radius,R0).
new(CE,square,[R0]) :- new(CE,object,[]),field(CE,[square],side,R0).

dec_meth(object,equals).
dec_meth(object,toString).
dec_meth(int,mult).
dec_meth(int,gth).
dec_meth(shape_reader,nextLine).
dec_meth(shape_reader,readCircle).
dec_meth(shape_reader,readSquare).
dec_meth(shape_reader,print).
dec_meth(shape_reader,read).
dec_meth(circle,getRadius).
dec_meth(circle,area).
dec_meth(circle,largerThan).
dec_meth(square,getSide).
dec_meth(square,area).
dec_meth(square,largerThan).

dec_field(circle,radius).
dec_field(square,side).
has_field(C,F) :- dec_field(C,F).
has_field(C,F) :- extends(C,P),has_field(P,F),\+ dec_field(C,F).

has_meth(CE,object,equals,[This0,_],R) :- subclasstype(This0,object),subtype([bool],R
has_meth(CE,object,toString,[This0],R) :- subclasstype(This0,object),subtype([string]
has_meth(CE,int,mult,[This0,T],R) :- subclasstype(This0,int),subtype(T,[int]),subtype
has_meth(CE,int,gth,[This0,T],R) :- subclasstype(This0,int),subtype(T,[int]),subtype(
has_meth(CE,shape_reader,nextLine,[This0],R) :- subclasstype(This0,shape_reader),subt
has_meth(CE,shape_reader,readCircle,[This0],R) :- subclasstype(This0,shape_reader),ne
has_meth(CE,shape_reader,readSquare,[This0],R) :- subclasstype(This0,shape_reader),ne
has_meth(CE,shape_reader,print,[This0,[string]],[void]) :- subclasstype(This0,shape_r
has_meth(CE,shape_reader,read,[This0],SH2) :- subclasstype(This0,shape_reader), invok
    subtype(R1,[bool]),var_upd(This1,This0),var_upd(This2,This0),var_upd(ST1,ST0),var_upd
has_meth(CE,circle,getRadius,[This0],R0) :- subclasstype(This0,circle),field(CE,This0
has_meth(CE,circle,area,[This0],R1) :- subclasstype(This0,circle),field(CE,This0,radi
has_meth(CE,circle,largerThan,[This0,SH0],RT2) :- subclasstype(This0,circle),inter(SH
has_meth(CE,square,getSide,[This0],R0) :- subclasstype(This0,square),field(CE,This0,s
has_meth(CE,square,area,[This0],R1) :- subclasstype(This0,square),field(CE,This0,side
has_meth(CE,square,largerThan,[This0,SH0],RT2) :- subclasstype(This0,square),inter(SH
has_meth(CE,C,M,A,R):- extends(C,P),has_meth(CE,P,M,A,R), \+ dec_meth(C,M).

var_upd(T1,T2) :- subtype(T2,T1).
field_upd(CE,T1,F,T2) :- field(CE,T1,F,FT),subtype(T2,FT).

inter([C1|T],C2,T2) :- inter(T,C2,T1),(subclass(C1,C2) -> ord_add_element(T1,C1,T2);T
```

```
inter([],_,[]).
diff([C1|T],C2,T2) :- diff(T,C2,T1),(subclass(C1,C2) -> T2=T1;ord_add_element(T1,C1,T
diff([],_,[]).
```