

Object-Oriented Type Inference

Jens Palsberg
palsberg@daimi.au.dk

Michael I Schwartzback
mis@daimi.au.dk

Department of Computer Science, Aarhus University
Ny Munkegade, DK-8000 Århus C, Denmark

March 1991

Abstract

We present a new approach to inferring types in untyped object-oriented programs with inheritance, assignments, and late binding. It guarantees that all messages are understood, annotates the program with type information, allows polymorphic methods, and can be used as the basis of an optimizing compiler. Types are finite sets of classes and subtyping is set inclusion. Using a *trace graph*, our algorithm constructs a set of conditional type constraints and computes the least solution by least fixed-point derivation.

1 Introduction

Untyped object-oriented languages with assignments and late binding allow rapid prototyping because classes inherit implementation and not specification. Late binding, however, can cause programs to be unreliable, unreadable, and inefficient [27]. Type inference may help solve these problems, but so far no proposed inference algorithm has been capable of checking most common, completely untyped programs [9].

We present a new type inference algorithm for a basic object-oriented language with inheritance, assignments, and late binding. The algorithm guarantees that all messages are understood, annotates the program with type information, allows polymorphic methods, and can be used as the basis of an optimizing compiler. Types are finite sets of classes and subtyping is set inclusion. Given a concrete program, the algorithm constructs a finite graph of type constraints. The program is *typable* if these constraints are solvable. The algorithm then computes the least solution in worst-case exponential time. The graph contains all type information that can be derived from the program without keeping track of `nil` values or flow analyzing the contents of instance variables. This makes the algorithm capable of checking most common programs; in particular, it allows for polymorphic methods. The algorithm is similar to previous work on type inference [18, 14, 27, 1, 2, 19, 12, 10, 9] in using type constraints, but it differs in handling late binding by *conditional* constraints and in resolving the constraints by least fixed-point derivation rather than unification.

The example language resembles SMALLTALK [8] but avoids metaclasses, blocks, and primitive methods. Instead, it provides explicit `new` and `if-then-else` expressions; classes like `Natural` can be programmed in the language itself.

In the following section we discuss the impacts of late binding on type inference and examine previous work. In later sections we briefly outline the example language, present the type inference algorithm, and show some examples of its capabilities.

2 Late Binding

Late binding means that a message send is dynamically bound to an implementation depending on the class of the receiver. This allows a form of polymorphism which is fundamental in object-oriented programming. It also, however, involves the danger that the class of the receiver does *not* implement a method for the message—the receiver may even be `nil`. Furthermore, late binding can make the control flow of a program hard to follow and may cause a time-consuming run-time search for an implementation.

It would significantly help an optimizing compiler if, for each message send in a program text, it could infer the following information.

- Can the receiver be `nil`?
- Can the receiver be an instance of a class which does not implement a method for the message?
- What are the classes of all possible non-`nil` receivers in any execution of the program?

Note that the available set of classes is induced by the particular program. These observations lead us to the following terminology.

Terminology:

Type: A type is a finite set of classes.

Induced Type: The induced type of an expression in a concrete program is the set of classes of all possible non-`nil` values to which it may evaluate in any execution of that particular program.

Sound approximation: A sound approximation of the induced type of an expression in a concrete program is a *superset* of the induced type.

Note that a sound approximation tells (“the whole truth”) but not always “nothing but the truth” about an induced type. Since induced types are generally uncomputable, a compiler must make do with sound approximations. An induced type is a *subtype* of any sound approximation; subtyping is set inclusion. Note also that our notion of type, which we also investigated in [22], differs from those usually used in theoretical studies of types in object-oriented programming [3, 7]; these theories cannot cope with late binding.

The goals of type inference can now be phrased as follows.

Goals of type inference:

Safety guarantee: A guarantee that any message is sent to either nil or an instance of a class which implements a method for the message; and, given that, also

Type information: A sound approximation of the induced type of any receiver.

Note that we ignore checking whether the receiver is nil; this is a standard data flow analysis problem which can be treated separately.

If a type inference is successful then the program is *typable*; the error `messageNotUnderstood` will not occur. A compiler can use this to avoid inserting some checks in the code. Furthermore, if the type information of a receiver is a *singleton* set, then the compiler can do early binding of the message to the only possible method; it can even do in-line substitution. Similarly, if the type information is an *empty* set, then the receiver is known to always be nil. Finally, type information obtained about variables and arguments may be used to annotate the program for the benefit of the programmer.

Untyped object-oriented languages such as SMALLTALK are traditionally implemented by interpreters. This is ideal for prototyping and exploratory development but often too inefficient and space demanding for real-time applications and embedded systems. What is needed is an optimizing compiler that can be used near the end of the programming phase, to get the required efficiency and a safety guarantee. A compiler which produces good code can be tolerated even it is slow because it will be used much less often than the usual programming environment. Our type inference algorithm can be used as the basis of such an optimizing compiler. Note, though, that both the safety guarantee and the induced types are sensitive to small changes in the program, Hence, separate compilation of classes seems impossible. Typed object-oriented languages such as SIMULA[6] BETA[15], C++[26], and Eiffel [17] allow separate compilation but sacrifice flexibility. The relations between types and implementation are summarized in figure 1.

Graver and Johnson [10, 9], in their type system for SMALLTALK, take an intermediate approach between “untyped” and “typed” in requiring the programmer to specify types for instance variables whereas types of arguments are inferred. Suzuki [27], in his pioneering work on inferring types in

When programs are:	They can be implemented by:
Untyped	Interpretation
Typable	Compilation
Typed	Separate Compilation

Figure 1: Types and implementation.

SMALLTALK, handles late binding by assuming that each message send may invoke all methods for that message. It turned out, however, that this yields an algorithm which is not capable of checking most common programs.

Both these approaches include a notion of *method type*. Our new type inference algorithm abandons this idea and uses instead the concept of *conditional constraints*, derived from a finite graph. Recently, Hense [11] addressed type inference for a language O'SMALL which is almost identical to our example language. He uses a radically different technique, with type schemes and unification based on work of Rémy [24] and Wand [29]. His paper lists four programs of which his algorithm can type-check only the first three. Our algorithm can type-check all four, in particular the fourth which is shown in figure 13 in appendix B. Hense uses record types which can be extendible and recursive. This seems to produce less precise typings than our approach, and it is not clear whether the typings would be useful in an optimizing compiler. One problem is that type schemes always correspond to either singletons or infinite sets of monotypes; our finite sets can be more precise. Hense's and ours approaches are similar in neither keeping track of nil values nor flow analyzing the contents of variables. We are currently investigating other possible relations.

Before going into the details of our type inference algorithm we first outline an example language on which to apply it.

3 The Language

Our example language resembles SMALLTALK, see figure 2.

A *program* is a set of classes followed by an expression whose value is the result of executing the program. A *class* can be defined using inheritance

(Program)	P	::=	C ₁ . . . C _n E
(Class)	C	::=	class ClassId [inherits ClassId] var Id ₁ . . . Id _k M ₁ . . . M _n end ClassId
(Method)	M	::=	method m ₁ Id ₁ . . . m _n Id _n E
(Expression)	E	::=	Id := E Em ₁ E ₁ . . . m _n E _n E; E if E then E else E ClassId new self class new E instanceof ClassId self super Id nil

Figure 2: Syntax of the example language.

and contains instance variables and methods; a *method* is a message selector ($m_1 \dots m_n$) with formal parameters and an expression. The language avoids metaclasses) blocks) and primitive methods. Instead) it provides explicit **new** and **if-then-else** expressions (the latter tests if the condition is non-nil). The result of a sequence is the result of the last expression in that sequence. The expression “**self class new**” yields an instance of the class of **self**. The expression “ E instanceof ClassId” yields a run-time check for class membership. If the check fails) then the expression evaluates to **nil**.

The SMALLTALK system is based on some primitive methods, written in assembly language. This dependency on primitives is not necessary, at least not in this theoretical study, because classes such as **True**, **False**, **Natural**, and **List** can be programmed in the language itself, as shown in appendix A.

4 Type Inference

Our type inference algorithm is based on three fundamental observations.

Observations:

Inheritance: Class inherit implementation and not specification.

Classes: There are finitely many classes in a program.

Message sends: There are finitely many syntactic message sends in a program.

The first observation leads to separate type inference for a class and its subclasses. Notionally, this is achieved by expanding all classes before doing type inference. This expansion means removing all inheritance by

- Copying the text of a class to its subclasses; and
- Replacing each message send to **super** by a message send to a renamed version of the inherited method; and finally
- Replacing each “**self class new**” expression by a “**ClassId new**” expression where **ClassId** is the enclosing class in the expanded program.

This idea of expansion is inspired by Graver and Johnson [10, 9]; note that the size of the expanded program is at most quadratic in the size of the original.

Example program in:	Illustrates:	Can we type it?
Figure 10	Basic type inference	Yes
Figure 11	A polymorphic method	Yes
Figure 12	A recursive method	Yes
Figure 13	Hense’s program	Yes
Figure 14	A realistic program	Yes
Figure 15	Lack of flow analysis	Yes
Figure 16	Lack of nil detection	No

Figure 3: An overview of the example programs.

The second and third observation lead to a finite representation of type information about all executions of the expanded program; this representation is called the *trace graph*. From this graph a finite set of type constraints will be generated. Typability of the program is then solvability of these constraints. Appendix B contains seven example programs which illustrate different aspects of the type inference algorithm, see the overview in figure 3. The program texts are listed together with the corresponding constraints and their least solution, if it exists. Hense’s program in figure 13 is the one he gives as a typical example of what he cannot type-check [11]. We invite the reader to consult the appendix while reading this section.

The graph will contain the following three kinds of type information.

Three kinds of type information:

Local constraints: Generated from method bodies.

Connecting constraints: Generated from message sends.

Conditions: Discriminate receivers.

Trace Graph Nodes

The *nodes* of the trace graph are obtained from the various methods implemented in the program. Each method yields a number of different nodes: one for each syntactic message send with the corresponding selector. The situation is illustrated in figure 4, where we see the nodes for a method m that is implemented in each of the classes C_1, C_2, \dots, C_n . Thus, the number of nodes in the trace graph will at most be quadratic in the size of the program. There is also a single node for the main expression of the program, which we may think of as a special method without

Methods do not have types, but they can be provided with type annotations, based on the types of their formal parameters and result. A particular method implementation may be represented by several nodes in the trace graph. This enables it to be assigned several different type annotations—one for each syntactic call. This allows us effectively to obtain method polymorphism through a finite set of method “monotypes”.

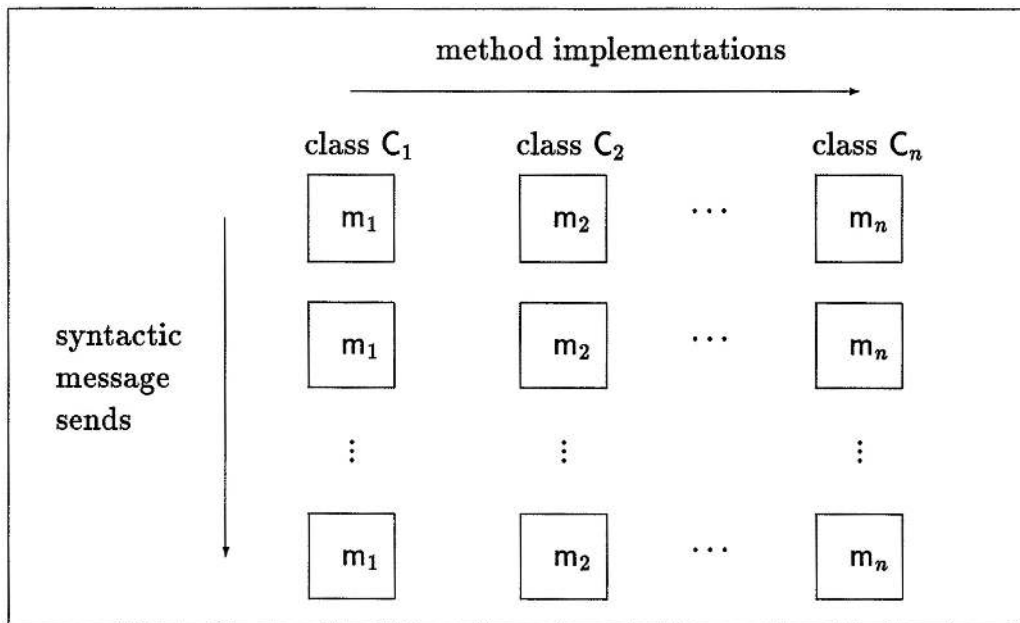


Figure 4: Trace graph nodes.

Local Constraints

Each node contains a collection of *local constraints* that the types of expressions must satisfy. For each syntactic occurrence of an expression E in the implementation of the method, we regard its type as an unknown variable $\llbracket E \rrbracket$. Exact type information is, of course, uncomputable. In our approach, we will ignore the following two aspects of program executions.

Approximations:

Nil values: It does not keep track of nil values.

Instance variables: It does not flow analyze the contents of instance variables.

The first approximation stems from our discussion of the goals of type inference; the second corresponds to viewing an instance variable as having a single possibly large type, thus leading us to identify the type variables of different occurrences of the same instance variable. In section 5 we present

two program fragments that are typical for what we cannot type because of these approximations. In both cases the constraints demand the false inclusion $\{\text{True}\} \subseteq \{\text{Natural}\}$. Suzuki [27] and Hense [11] make the same approximations.

For an expression E , the local constraints are generated from all the phrases in its derivation, according to the rules in figure 5. The idea of generating constraints on type variables from the program syntax is also exploited in [28, 25].

	<u>Expression:</u>	<u>Constraint:</u>
1)	$\text{Id} := E$	$\llbracket \text{Id} \rrbracket \supseteq \llbracket E \rrbracket \wedge \llbracket \text{Id} := E \rrbracket = \llbracket E \rrbracket$
2)	$E \ m_1 \ E_1 \ \dots \ m_n \ E_n$	$\llbracket E \rrbracket \subseteq \{C \mid C \text{ implements } m_1 \ \dots \ m_n\}$
3)	$E_1; E_2$	$\llbracket E_1; E_2 \rrbracket = \llbracket E_2 \rrbracket$
4)	$\text{if } E_1 \text{ then } E_2 \text{ else } E_3$	$\llbracket \text{if } E_1 \text{ then } E_2 \text{ else } E_3 \rrbracket = \llbracket E_2 \rrbracket \cup \llbracket E_3 \rrbracket$
5)	$C \ \text{new}$	$\llbracket C \ \text{new} \rrbracket = \{C\}$
6)	$E \ \text{instanceof } C$	$\llbracket E \ \text{instanceof } C \rrbracket = \{C\}$
7)	self	$\llbracket \text{self} \rrbracket = \{\text{the enclosing class}\}$
8)	Id	$\llbracket \text{Id} \rrbracket = \llbracket \text{Id} \rrbracket$
9)	nil	$\llbracket \text{nil} \rrbracket = \{ \}$

Figure 5: The local constraints.

The constraints guarantee safety; only in the cases 4) and 8) do the approximations manifest themselves. Notice that the constraints can all be expressed as inequalities of one of the three forms: “constant \subseteq variable”, “variable \subseteq constant”, or “variable \subseteq variable”; this will be exploited later.

Each different node employs unique type variables, except that the types of instance variables are common to all nodes corresponding to methods implemented in the same class. A similar idea is used by Graver and Johnson [10, 9].

Trace Graph Edges

The *edges* of the trace graph will reflect the possible connections between a message send and a method that may implement it. The situation is illustrated in figure 6.

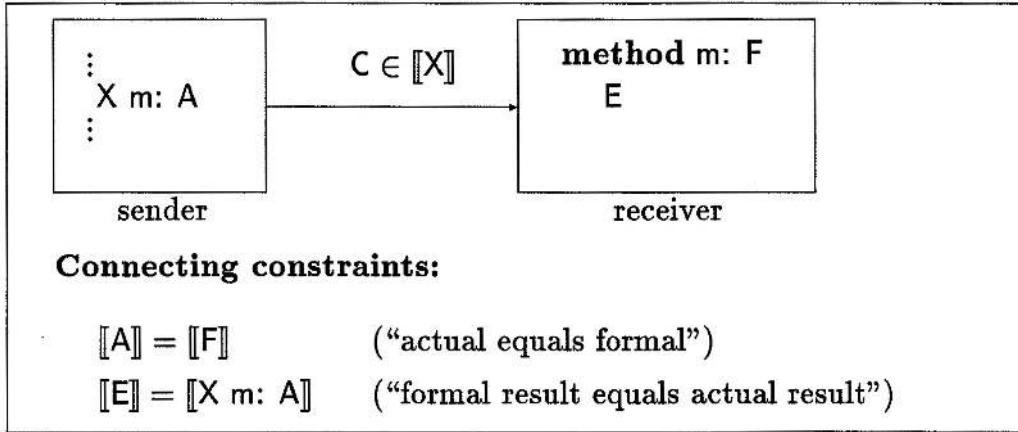


Figure 6: Trace graph edges.

If a node corresponds to a method which contains a message send of the form $X \ m: \ A$, then we have an edge from that sender node to any other receiver node which corresponds to an implementation of a method m . We label this edge with the condition that the message send may be executed, namely $C \in \llbracket X \rrbracket$ where C is the class in which the particular method m is implemented. With the edge we associate the *connecting constraints*, which reflect the relationship between formal and actual parameters and results. This situation generalizes trivially to methods with several parameters. Note that the number of edges is again quadratic in the size of the program.

Global Constraints

To obtain the *global constraints* for the entire program we combine local and connecting constraints in the manner illustrated in figure 7. This produces *conditional constraints*, where the inequalities need only hold if all the conditions hold. The global constraints are simply the union of the conditional constraints generated by all paths in the graph, originating in the node corresponding to the main expression of the program. This is a finite set, because the graph is finite; as shown later in this section, the size of the constraint set may in (extreme) worst-cases become exponential.

If the set of global constraints has a solution, then this provides approximate

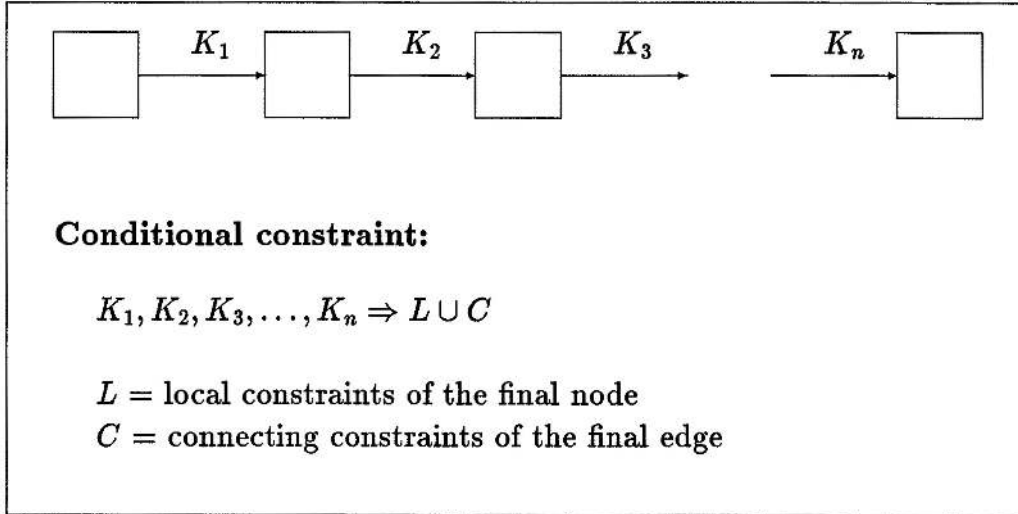


Figure 7: Conditional constraints from a path.

information about the dynamic behavior of the program.

Consider any execution of the program. While observing this, we can trace the pattern of method executions in the trace graph. Let E be some expression that is evaluated at some point, let $\text{VAL}(E)$ be its value, and let $\text{CLASS}(b)$ be the class of an object b . If L is some solution to the global constraints, then the following result holds.

Soundness Theorem:

If $\text{VAL}(E) \neq \text{nil}$, then $\text{CLASS}(\text{VAL}(E)) \in L(\llbracket E \rrbracket)$.

It is quite easy to see that this must be true. We sketch a proof by induction in the number of message sends performed during the trace. If this is zero, then we rely on the local constraints alone; given a dynamic semantics [5, 4, 23, 13] one can easily verify that their satisfaction implies the above property. If we extend a trace with a message send $X \text{ m: A}$ implemented by a method in a class C , then we can inductively assume that $C \in L(\llbracket X \rrbracket)$. But this implies that the local constraints in the node corresponding to the invoked method must hold, since all their conditions now hold and L is a solution. Since the relationship between actual and formal parameters and results is soundly represented by the connecting constraints, which also must hold, the result follows.

Note that an expression E occurring in a method that appears k times in the trace graph has k type variables $\llbracket E \rrbracket_1, \llbracket E \rrbracket_2, \dots, \llbracket E \rrbracket_k$ in the global constraints. A sound approximation to the induced type of E is obtained as

$$\bigcup_i \mathcal{L}(\llbracket E \rrbracket_i)$$

Appendix C gives an efficient algorithm to compute the smallest solution of the extracted constraints, or to decide that none exists. The algorithm is at worst quadratic in the size of the constraint set.

The complete type inference algorithm is summarized in figure 8.

Input:	A program in the example language.
Output:	Either: a safety guarantee and type information about all expressions; or: “unable to type the program”.
1)	Expand all classes.
2)	Construct the trace graph of the expanded program.
3)	Extract a set of type constraints from the trace graph.
4)	Compute the least solution of the set of type constraints. If such a solution exists, then output it as the wanted type information, together with a safety guarantee; otherwise, output “unable to type the program”.

Figure 8: Summary of the type inference algorithm.

Type Annotations

Finally, we will consider how a solution L of the type constraints can produce a *type annotation* of the program. Such annotations could be provided for the benefit of the programmer.

An instance variable x has only a single associated type variable. The type annotation is simply $L(\llbracket x \rrbracket)$. The programmer then knows an upper bound of the set of classes whose instances may reside in x .

A method has finitely many type annotations, each of which is obtained from a corresponding node in the trace graph. If the method, implemented in the class C , is

$$\text{method } m_1 : F_1 \ m_2 : F_2 \dots m_n : F_n \\ E$$

then each type annotation is of the form

$$\{C\} \times L(\llbracket F_1 \rrbracket) \times \dots \times L(\llbracket F_n \rrbracket) \rightarrow L(\llbracket E \rrbracket)$$

The programmer then knows the various manners in which this method may be used.

A constraint solution contains more type information about methods than the method types used by Suzuki. Consider for example the polymorphic identity function in figure 11. Our technique yields both of the method type annotations

$$\text{id} : \{C\} \times \{\text{True}\} \rightarrow \{\text{True}\}, \quad \text{id} : \{C\} \times \{\text{Natural}\} \rightarrow \{\text{Natural}\}$$

whereas the method type using Suzuki's framework is

$$\text{id} : \{C\} \times \{\text{True}, \text{Natural}\} \rightarrow \{\text{True}, \text{Natural}\}$$

which would allow neither the `succ` nor the `isTrue` message send, and, hence, would lead to rejection of the program.

An Exponential Worst-case

The examples in appendix B show several cases where the constraint set is quite small, in fact linear in the size of the program. While this will often be the situation, the theoretical worst-case allows the constraint set to become exponential in the size of the program. The running time of the inference algorithm depends primarily on the topology of the trace graph.

In figure 9 is shown a program and a sketch of its trace graph. The induced constraint set will be exponential since the graph has exponentially many different paths. Among the constraints will be a family whose conditions are similar to the words of the regular language

$$(\text{CCC} + \text{DCC})^{\frac{n}{3}}$$




Figure 9: A worst-case program.

the size of which is clearly exponential in n .

Note that this situation is similar to that of type inference in ML, which is also worst-case exponential but very useful in practice. The above scenario is in fact not unlike the one presented in [16] to illustrate exponential running times in ML. Another similarity is that both algorithms generate a potentially exponential constraint set that is always solved in polynomial time.

5 Conclusion

Our type inference algorithm is sound and can handle most common programs. It is also conceptually simple: a set of uniform type constraints is constructed and solved by fixed-point derivation. It can be further improved by an orthogonal effort in data flow analysis.

The underlying type system is simple: types are finite sets of classes and subtyping is set inclusion.

Future work includes implementation of an optimizing compiler based on the type inference algorithm. The algorithm should be easy to extend to work for full SMALLTALK, because metaclasses are simply classes, blocks can be treated as objects with a single method, and primitive methods can be handled by stating the constraints that the machine code must satisfy. Another challenge is to extend the algorithm to produce type annotations together with type substitution, see [21, 20, 22].

A Basic classes

```
class Object
end Object

class True
  method isTrue
    Object new % Returns something non-nil.
end True

class False
  method isTrue
    nil
end False

% Henceforth, we abbreviate “True new” as “true”, and “False new” as “false”.

class Natural
  var rep % The predecessor of non-zero receivers; otherwise nil.
  method isZero
    if rep then false else true
  method succ
    (Natural new) update: self
  method update: x
    rep := x; self
  method pred
    if (self isZero) isTrue then self else rep
  method less: i
    if (i isZero) isTrue
    then false
    else if (self isZero) isTrue then true else (self pred) less: (i pred)
end Natural

% Henceforth, we abbreviate “Natural new” as “0” and, recursively,
% “n succ” as “n + 1”.
% For example, “2” is an abbreviation for “((Natural new) succ) succ”.
```

```

class List
  var head, tail
  method setHead: h setTail: t
    head := h; tail := t
  method cons: x
    (self class new) setHead: x setTail: self
  method isEmpty
    if head then false else true
  method car
    head
  method cdr
    tail
  method append: aList
    if (self isEmpty) isTrue
      then aList
    else (tail append: aList) cons: head
  method insert: x
    if (self isEmpty) isTrue
      then self cons: x
    else
      if (head less: x) isTrue
        then self cons: x
      else (tail insert: x) cons: head
  method sort
    if (self isEmpty) isTrue then self else (tail sort) insert: head
  method merge: aList
    if (self isEmpty) isTrue
      then aList
    else
      if (head less: (aList car)) isTrue
        then (tail merge: aList) cons: head
      else else (self merge: (aList cdr)) cons: (aList car)
end List

```

```
class Comparable
  var key
  method getKey
    key
  method setKey: k
    key := k
  method less: c
    key less: (c getKey)
end Comparable
```

B Example Programs

```
class A
  method f
    7
end A
```

```
class B
  method f
    true
end B
```

```
x := A new; (x f) succ
```

Constraints:

```
[[A new]] = {A}
[[x]] ⊇ [[A new]]
[[x := A new]] = [[A new]]
[[x]] ⊇ {A, B}
A ∈ [[x]] ⇒ [[x f]] = [[7]]
A ∈ [[x]] ⇒ [[7]] = {Natural}
B ∈ [[x]] ⇒ [[x f]] = [[true]]
B ∈ [[x]] ⇒ [[true]] = {True}
[[x f]] ⊆ {Natural}
Natural ∈ [[x f]] ⇒ [[(x f) succ]] = {Natural}
[[x := A new; (x f) succ]] = [[(x f) succ]]
```

Smallest Solution:

```
[[x]] = [[A new]] = [[x := A new]] = {A}
[[x f]] = [[(x f) succ]] = [[x := A new; (x f) succ]] = [[7]]{Natural}
[[true]] = {True}
```

Figure 10: Conditions at work.

```

class C
  method id: x
    x
  end C

((C new ) id: 7) succ;
((C new ) id: true) isTrue

Constraints:
[[C new]]1 = {C}
[[C new]]1 ⊆ {C}
C ∈ [[C new]]1 ⇒ [[7]] = [[x]]1
C ∈ [[C new]]1 ⇒ [[x]]1 = [(C new) id: 7]
[[7]] = {Natural}
[(C new) id: 7] ⊆ {Natural}
Natural ∈ [(C new) id: 7] ⇒ { Natural } = [((C new) id: 7) succ]
[[C new]]2 = {C}
[[C new]]2 ⊆ {C}
C ∈ [[C new]]2 ⇒ [[true]] = [[x]]2
C ∈ [[C new]]2 ⇒ [[x]]2 = [(C new) id: true]
[[true]] = {True}
[(C new) id: true] ⊆ {True,False}
True ∈ [(C new) id: true] ⇒ {Object} = [((C new) id: true) isTrue]
False ∈ [(C new) id: true] ⇒ { } = [((C new) id: true) isTrue]

Smallest Solution:
[[C new]]1 = [[C new]]2 = {C}
[[7]] = [[x]]1 = [(C new) id: 7] = [((C new) id: 7) succ]{Natural}
[[true]] = [[x]]2 = [(C new) id: true] = {True}
[((C new) id: true) isTrue] = {Object}

```

Figure 11: A polymorphic method

```

class D
  method f: x
    if x then self f: x else nil
end D

(D new) f: nil

Constraints:
[[D new]] = {D}
[[D new]] ⊆ {D}
D ∈ [[D new]] ⇒ [[nil]] = [[x]]1
D ∈ [[D new]] ⇒ [[if x then self f: x else nil]]1 = [[(D new) f: nil]]
D ∈ [[D new]] ⇒ [[if x then self f: x else nil]]1 = [[self f: x]]1 ∪ [[nil]]1
D ∈ [[D new]] ⇒ [[nil]]1 = {}
D ∈ [[D new]] ⇒ [[self]]1 = {D}
D ∈ [[D new]] ⇒ [[self]]1 ⊆ {D}
D ∈ [[D new]], D ∈ [[self]]1 ⇒ [[x]]1 = [[x]]2
D ∈ [[D new]], D ∈ [[self]]1 ⇒ [[if x then self f: x else nil]]2 = [[self f: x]]1
D ∈ [[D new]], D ∈ [[self]]1 ⇒ [[if x then self f: x else nil]]2 = [[self f: x]]2 ∪ [[nil]]2
D ∈ [[D new]], D ∈ [[self]]1 ⇒ [[nil]]2 = {}
D ∈ [[D new]], D ∈ [[self]]1 ⇒ [[self]]2 = {D}
D ∈ [[D new]], D ∈ [[self]]1 ⇒ [[self]]2 ⊆ {D}
D ∈ [[D new]], D ∈ [[self]]1, D ∈ [[self]]2 ⇒ [[x]]2 = [[x]]2
D ∈ [[D new]], D ∈ [[self]]1, D ∈ [[self]]2 ⇒ [[if x then self f: x else nil]]2
    = [[self f: x]]2

[[nil]] = {}

Smallest Solution:
[[D new]]1 = [[self]]1 = [[self]]2 = {D}
[[nil]] = [[x]]1 = [[nil]]1 = [[if x then self f: x else nil]]1 = [[self f: x]]1 =
[[ (D new) f: nil ]]1 = [[x]]2 = [[nil]]2 = [[if x then self f: x else nil]]2 =
[[self f: x]]2 = [[ (D new) f: nil ]]2 = {}

```

Figure 12: A recursive method

```

class A
  method m
    0
end A

class B inherits A
  method n
    0
end B

a := A new
b := B new
a := b
a m

Constraints:
[[A new]] = {A}
[[a]] ⊇ [[A new]]
[[B new]] = {B}
[[b]] ⊇ [[B new]]
[[a]] ⊇ [[b]]
[[a]] ⊆ {A, B}
A ∈ [[a]] ⇒ [[a m]] = [[0]]
B ∈ [[a]] ⇒ [[a m]] = [[0]]
[[0]] = {Natural}
⋮

Smallest Solution:
[[a]] = {A, B}
[[b]] = {B}
[[a m]] = {Natural}
[[A new]] = {A}
[[B new]] = {B}
⋮

```

Figure 13: Hense's program

```

class Student inherits Comparable
  ...
end Student
class ComparableList inherits List
  method studentcount % Returns the number of Students instances.
    if (self isEmpty) isTrue
      then 0
    else
      if (self car) instanceof Student
        then ((self cdr) studentcount) succ
      else (self cdr) studentcount
    end
  end ComparableList

```

Figure 14: An example program.

```

x := 7;
x succ;
x := true;
x isTrue

Constraints:
[[x]] ⊇ [[7]]
[[7]] = {Natural}
[[x]] ⊆ {Natural}
[[x]] ⊇ [[true]]
[[true]] = {True}
[[x]] ⊆ {True, False}
⋮

```

Figure 15: A safe program rejected.

(if nil then true else 7) succ

Constraints:

$\llbracket \text{if nil then true else 7} \rrbracket \subseteq \{\text{Natural}\}$
 $\llbracket \text{if nil then true else 7} \rrbracket = \llbracket \text{true} \rrbracket \cup \llbracket 7 \rrbracket$
 $\llbracket \text{true} \rrbracket = \{\text{True}\}$
 $\llbracket 7 \rrbracket = \{\text{Natural}\}$
 \vdots

Figure 16: Another safe program rejected.

C Solving Systems of Conditional Inequalities

This appendix shows how to solve a finite system of conditional inequalities in quadratic time.

Definition C.1: A *CI-system* consists of

- a finite set \mathcal{A} of *atoms*.
- a finite set $\{\alpha_i\}$ of *variables*.
- a finite set of *conditional inequalities* of the form

$$C_1, C_2, \dots, C_k \Rightarrow Q$$

Each C_i is a *condition* of the form $a \in \alpha_j$, where $a \in \mathcal{A}$ is an atom, and Q is an *inequality* of one of the following forms

$$\begin{aligned} A &\subseteq \alpha_i \\ \alpha_i &\subseteq A \\ \alpha_i &\subseteq \alpha_j \end{aligned}$$

where $A \subseteq \mathcal{A}$ is a set of atoms.

A *solution* L of the system assigns to each variable α_i a set $L(\alpha_i) \subseteq \mathcal{A}$ such that all the conditional inequalities are satisfied. \square

In our application, \mathcal{A} models the set of classes occurring in a concrete program.

Lemma C.2: Solutions are closed under intersection. Hence, if a CI-system has solutions, then it has a unique minimal one.

Proof: Consider any conditional inequality of the form $C_1, C_2, \dots, C_k \Rightarrow Q$, and let $\{L_i\}$ be all solutions. We shall show that $\cap_i L_i$ is a solution. If a condition $a \in \cap_i L_i(\alpha_j)$ is true, then so is all of $a \in L_i(\alpha_j)$. Hence, if all the conditions of Q are true in $\cap_i L_i$, then they are true in each L_i ; furthermore, since they are solutions, Q is also true in each L_i . Since, in general, $A_k \subseteq B_k$ implies $\cap_k A_k \subseteq \cap_k B_k$, it follows that $\cap_i L_i$ is a solution. Hence, if there are any solutions, then $\cap_i L_i$ is the unique smallest one. \square

Definition C.3: Let \mathcal{C} be a CI-system with atoms \mathcal{A} and n distinct variables. An *assignment* is an element of $(2^{\mathcal{A}})^n \cup \{error\}$ ordered as a lattice, see figure 17.

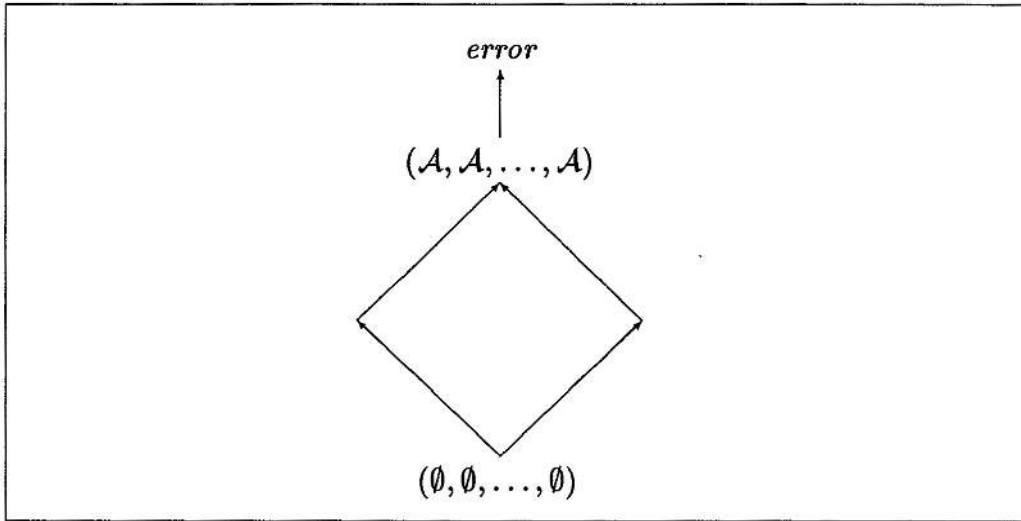


Figure 17: The lattice of assignments.

If different from *error*, then it assigns a set of atoms to each variable. If V is an assignment, then $\tilde{\mathcal{C}}(V)$ is a new assignment, defined as follows. If $V = error$, then $\tilde{\mathcal{C}}[V] = error$. An inequality is *enabled* if all of its conditions are true under V . If for any enabled inequality of the form $\alpha_i \subseteq A$ we do not have $V(\alpha_j) \subseteq A$, then $\tilde{\mathcal{C}}(V) = error$; otherwise, $\tilde{\mathcal{C}}(V)$ is the smallest

pointwise extension of V such that

- for every enabled inequality of the form $A \subseteq \alpha_j$ we have $A \subseteq \tilde{\mathcal{C}}(V)(\alpha_j)$.
- for every enabled inequality of the form $\alpha_i \subseteq \alpha_j$ we have $V(\alpha_i) \subseteq \tilde{\mathcal{C}}(V)(\alpha_j)$.

Clearly, $\tilde{\mathcal{C}}$ is monotonic in the above lattice. \square

Lemma C.4: An assignment $L \neq error$ is a solution of a CI-system \mathcal{C} iff $L = \tilde{\mathcal{C}}(L)$. If \mathcal{C} has no solutions, then $error$ is the smallest fixed-point of $\tilde{\mathcal{C}}$.

Proof: If L is a solution of \mathcal{C} , then clearly $\tilde{\mathcal{C}}$ will not equal $error$ and cannot extend L ; hence, L is a fixed-point. Conversely, if L is a fixed-point of $\tilde{\mathcal{C}}$, then all the enabled inequalities must hold. If there are no solutions, then there can be no fixed-point below $error$. Since $error$ is by definition a fixed-point, the result follows. \square

This means that to find the smallest solution, or to decide that none exists, we need only compute the least fixed-point of $\tilde{\mathcal{C}}$.

Lemma C.5: For any CI-system \mathcal{C} , the least fixed-point of $\tilde{\mathcal{C}}$ is equal to

$$\lim_{k \rightarrow \infty} \tilde{\mathcal{C}}^k(\emptyset, \emptyset, \dots, \emptyset)$$

Proof: This is a standard result about monotonic functions on complete lattices. \square

Lemma C.6: Let n be the number of different conditions in a CI-system \mathcal{C} . Then

$$\lim_{k \rightarrow \infty} \tilde{\mathcal{C}}^k(\emptyset, \emptyset, \dots, \emptyset) = \tilde{\mathcal{C}}^{n+1}(\emptyset, \emptyset, \dots, \emptyset)$$

Proof: When no more conditions are enabled, then the fixed-point is obtained by a single application. Once a condition is enabled in an assignment, it will remain enabled in all larger assignments. It follows that after n iterations no new conditions can be enabled; hence, the fixed-point is obtained in at most $n + 1$ iterations. \square

Lemma C.7: The smallest solution to any CI-system, or the decision that none exists, can be obtained in quadratic time.

Proof: This follows from the previous lemmas. \square

References

- [1] Alan H. Borning and Daniel H. H. Ingalls. A type declaration and inference system for Smalltalk. In *Ninth Symposium on Principles of Programming Languages*, pages 133–141. ACM Press, January 1982.
- [2] L. Cardelli. A semantics of multiple inheritance. In G. Kahn, D. MacQueen, and Gordon Plotkin, editors, *Semantics of Data Types*, pages 51–68. Springer-Verlag (*LNCS* 173), 1984.
- [3] L. Cardelli and P. Wegner. On understanding types, data abstraction, and polymorphism. *ACM Computing Surveys*, 17(4), December 1985.
- [4] W. R. Cook. *A Denotational Semantics of Inheritance*. PhD thesis, Brown University, 1989.
- [5] William Cook and Jens Palsberg. A denotational semantics of inheritance and its correctness. In *Proc. OOPSLA '89, ACM SIGPLAN Fourth Annual Conference on Object-Oriented Programming Systems, Languages and Applications*, 1989. To appear in *Information and Computation*.
- [6] O. J. Dahl, B. Myhrhaug, and K. Nygaard. Simula 67 common base language. Technical report, Norwegian Computing Center, Oslo, Norway, 1968.
- [7] Scott Danforth and Chris Tomlinson. Type theories and object-oriented programming. *ACM Computing Surveys*, 20(1), March 1988.
- [8] A. Goldberg and D. Robson. *Smalltalk-80—The Language and its Implementation*. Addison- Wesley, 1983.
- [9] Justin O. Graver and Ralph E. Johnson. A type system for Smalltalk. In *Seventeenth Symposium on Principles of Programming Languages*, pages 136–150. ACM Press, January 1990.
- [10] Justin Owen Graver. *Type-Checking and Type-Inference for Object-Oriented Programming Languages*. PhD thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, August 1989. UIUCD-R-89-1539.

- [11] Andreas V. Hense. Polymorphic type inference for a simple object oriented programming language with state. Technical Report Tech. Bericht Nr. A 20/90, Universität des Saarlandes, 1990.
- [12] R. E. Johnson. Type-checking Smalltalk. In *Proc. OOPSLA'86, Object-Oriented Programming Systems, Languages and Applications*. Sigplan Notices, 21(11), November 1986.
- [13] S. Kaolin. Inheritance in Smalltalk-80: A denotational definition. In *Fifteenth Symposium on Principles of Programming Languages*, pages 80–87. ACM Press, January 1988.
- [14] Marc A. Kaplan and Jeffrey D. Ullman. A general scheme for the automatic inference of variable types. In *Fifth Symposium on Principles of Programming Languages*, pages 60–75. ACM Press, January 1978.
- [15] B. B. Kristensen, O. L. Madsen, B. Møller-Pedersen, and K. Nygaard. The BETA programming language. In B. Shriver and P. Wegner, editors, *Research Directions in Object-Oriented Programming*, pages 7–48. MIT Press, 1987.
- [16] H. G. Mairson. Decidability of ML typing is complete for deterministic exponential time. In *Seventeenth Symposium on Principles of Programming Languages*. ACM Press, January 1990
- [17] Bertrand Meyer. *Object-Oriented Software Construction*. Prentice-Hall, Englewood Cliffs, NJ, 1988.
- [18] Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17, 1978.
- [19] Prateek Mishra and Uday S. Reddy. Declaration-free type checking. In *Twelfth Symposium on Principles of Programming Languages*, pages 7–21. ACM Press, January 1985.
- [20] Jens Palsberg and Michael I. Schwartzbach. Genericity And Inheritance. Computer Science Department, Aarhus University. PB-318, 1990.
- [21] Jens Palsberg and Michael I. Schwartzbach. Type substitution for object-oriented programming. In *Proc. OOPSLA/ECOOP790, ACM SIGPLAN Fifth Annual Conference on Object-Oriented Programming*

Systems, Languages and Applications; European Conference on Object-Oriented Programming, 1990.

- [22] Jens Palsberg and Michael I. Schwartzbach. A unified type system for object-oriented programming. Computer Science Department, Aarhus University. PB-341, 1990.
- [23] U. S. Reddy. Objects as closures: Abstract semantics of object-oriented languages. In *Proc. ACM Conference on Lisp and Functional Programming*, pages 289–297. ACM, 1988.
- [24] Didier Rémy. Typechecking records and variants in a natural extension of ML. In *Sixteenth Symposium on Principles of Programming Languages*, pages 77–88. ACM Press, January 1989.
- [25] Michael I. Schwartzbach. Type inference with inequalities. In *Proc. TAPSOFT'91*. Springer-Verlag (*LNCS*), 1991.
- [26] B. Stroustrup. *The C++ Programming Language*. Addison:Wesley, 1986.
- [27] Norihisa Suzuki. Inferring types in Smalltalk. In *Eighth Symposium on Principles of Programming Languages*, pages 187–199. ACM Press, January 1981.
- [28] Mitchell Wand. A simple algorithm and proof for type inference. *Fundamentae Informaticae*, X:115–122, 1987.
- [29] Mitchell Wand. Type inference for record concatenation and multiple inheritance. In *LICS'89, Fourth Annual Symposium on Logic in Computer Science*, 1989.