

A Semantics of Multiple Inheritance

Luca Cardelli

AT&T Bell Laboratories
Murray Hill, New Jersey 07974

1. Introduction

There are two major ways of structuring data in programming languages. The first and common one, used for example in Pascal, can be said to derive from standard branches of mathematics. Data is organized as cartesian products (i.e. record types), disjoint sums (i.e. unions or variant types) and function spaces (i.e. functions and procedures).

The second method can be said to derive from biology and taxonomy. Data is organized in a hierarchy of classes and subclasses, and data at any level of the hierarchy *inherits* all the attributes of data higher up in the hierarchy. The top level of this hierarchy is usually called the class of all "objects"; every datum *is an* object and every datum *inherits* the basic properties of objects, like the ability to tell whether two objects are the same or not. Functions and procedures are also considered as local actions of objects, as opposed to global operations.

These different ways of structuring data have generated distinct classes of programming languages, and induced different programming styles. Programming with taxonomically organized data is often called *object-oriented programming*, and has been advocated as an effective way of structuring programming environments, data bases, and large systems in general.

The notions of inheritance and object-oriented programming first appeared in Simula 67 [Dahl 66]. In Simula, objects are grouped into classes and classes can be organized into a subclass hierarchy. Objects are similar to records with functions as components, and elements of a class can appear wherever elements of the respective superclasses are expected. Subclasses inherit all the attributes of their superclasses. In Simula, the issues are somewhat complicated by the use of objects as coroutines, so that communication between objects can be implemented as "message-passing" between processes.

Smalltalk [Goldberg 83] adopts and exploits the idea of inheritance, with some changes. While stressing the message-passing paradigm, a Smalltalk object is not usually a separate process. Message passing is just function call, although the association of message names to functions (called methods) is not straightforward. With respect to Simula, Smalltalk also abandons static scoping, to gain flexibility in interactive use, and strong typing, allowing it to implement system introspection and to introduce the notion of meta-classes.

Inheritance can be single or multiple. In the case of single inheritance, as in Simula or Smalltalk, the subclass hierarchy has the form of a tree, i.e. every class has a unique superclass. A class can be sometime considered a subclass of two incompatible superclasses; then an arbitrary decision has to be made to determine which superclass to use. This problem leads naturally to the idea of multiple inheritance.

Multiple inheritance occurs when an object can belong to several incomparable superclasses: the subclass relation is no longer constrained to form a tree, but can form a dag. Multiple inheritance is more elegant than simple inheritance, but more difficult to implement. So far, it has mostly been considered in the context of type-free dynamically-scoped languages and implemented as Lisp or Smalltalk extensions [Weinreb 81, Borning 82, Steels 83], or as part of knowledge representation languages [Attardi 81]. Exceptions are Galileo [Albano 83] and OBJ [Goguen 84]

where multiple inheritance is typechecked.

The differences between Simula, Smalltalk and other languages suggest that inheritance is the only notion critically associated with object-oriented programming. Coroutines, message-passing, static/dynamic scoping, typechecking and single/multiple superclasses are all fairly independent issues. Hence, a theory of object-oriented programming should first of all focus on the meaning of inheritance.

The aim of this paper is to present a clean semantics of multiple inheritance and to show that, in the context of strongly-typed, statically-scoped languages, a sound typechecking algorithm exists. Multiple inheritance is also interpreted in a broad sense: instead of being limited to objects, it is extended in a natural way to union types and to higher-order functional types.

A clean semantics has the advantage of making clear which issues are fundamental and which are implementation optimizations. The implementation of multiple inheritance suggested by the semantics is very naive, but does not preclude more sophisticated implementation techniques. It should however be emphasized that advanced implementation techniques are absolutely essential to obtain usable systems based on inheritance [Deutsch 84].

The first part of this paper is informal, and presents the basic notations and intuitions by means of examples. The second part is formal: it introduces a language, a semantics, a type inference system and a typechecking algorithm. The algorithm is proved sound with respect to the inference system, and the inference system is proved sound with respect to the semantics [Milner 78].

2. Objects as Records

There are several ways of thinking of what objects *are*. In the pure Smalltalk-like view, objects recall physical entities, like boxes or cars. Physical entities are unfortunately not very useful as semantic models of objects, because they are far too complicated to describe formally.

Two simpler interpretations of objects seem to emerge from the implementations of object-oriented languages. The first interpretation derives from Simula, where objects are essentially records with possibly functional components. Message passing is field selection and inheritance has to do with the number and type of fields possessed by a record.

The second interpretation derives from Lisp. An object is a function which receives a message (a string or an atom) and dispatches on the message to select the appropriate "method". Here message-passing is function application and inheritance has to do with the way messages are dispatched.

In some sense these two interpretations are equivalent because records can be represented as functions from labels (messages) to values. However, to say that objects are functions is misleading, because we must qualify that objects are functions over messages. Instead we can safely assert that objects are records, because labels are an essential part of records.

We also want to regard objects as records for typechecking purposes. While a (character string) message can be the result of an arbitrary computation, a record selection usually requires the selection label to be known at compile-time. In the latter case it is possible to statically determine the set of messages supported by an object, and a compile-time type error can be reported on any attempt to send unsupported messages. This property is true for Simula, but has been lost in all the succeeding languages.

We shall show how records can account for all the basic features of objects, provided that the surrounding language is rich enough. The features we consider are multiple inheritance, message-passing, private instance variables and the concept of "self". The duality between records and functions however remains: in our language objects are records, but in the semantics records are functions.

3. Records

A **record** is a finite association of values to labels, for example:

$$(a = 3, b = \text{true}, c = \text{"abc"})$$

is a record with three fields a , b and c having as values an integer 3, a boolean *true* and a string "abc" respectively. The labels a , b and c belong to a separate domain of labels; they are not identifiers or strings, and cannot be computed as the result of expressions. Records are unordered and cannot contain the same label twice.

The basic operation on records is field selection, denoted by the usual dot notation:

$$(a = 3, b = \text{true}, c = \text{"abc"}) . a \equiv 3$$

An expression can have one or more types; we write

$$e : \tau$$

to indicate that expression e has type τ .

Records have **record types** which are labeled sets of types with distinct labels, for example we have:

$$(a = 3, b = \text{true}) : (a : \text{int}, b : \text{bool})$$

In general, we can write the following informal typing rule for records:

[Rule1] if $e_1 : \tau_1$ and .. and $e_n : \tau_n$ then $(a_1 = e_1, \dots, a_n = e_n) : (a_1 : \tau_1, \dots, a_n : \tau_n)$

This is the first of a series of informal rules which are only meant to capture our initial intuitions about typing. They are not supposed to form a complete set or to be independent of each other.

There is a **subtype** relation on record types which corresponds to the *subclass* relation of Simula and Smalltalk. For example we may define the following types:

<i>type any</i>	=	()
<i>type object</i>	=	(age: int)
<i>type vehicle</i>	=	(age: int, speed: int)
<i>type machine</i>	=	(age: int, fuel: string)
<i>type car</i>	=	(age: int, speed: int, fuel: string)

Intuitively a vehicle *is* an object, a machine *is* an object and a car *is* a vehicle *and* a machine (and therefore an object). We say that *car* is a subtype of *machine* and *vehicle*; *machine* is a subtype of *object*; etc. In general a record type τ is a subtype (written \leq) of a record type τ' if τ has all the fields of τ' , and possibly more, and the common fields of τ and τ' are in the \leq relation. Moreover, all the basic types (like *int* and *bool*) are subtypes of themselves:

[Rule2] • $\iota \leq \iota$ (ι a basic type)
 • $\tau_1 \leq \tau'_1 \dots \tau_n \leq \tau'_n \implies (a_1 : \tau_1, \dots, a_{n+m} : \tau_{n+m}) \leq (a_1 : \tau'_1, \dots, a_n : \tau'_n)$

Let us consider a particular car (value definitions are prefixed by the keyword *val*):

$$\text{val mycar} = (\text{age} = 4, \text{speed} = 140, \text{fuel} = \text{"gasoline"})$$

Of course *mycar*: *car* (*mycar* has type *car*), but we might also want to assert *mycar*: *object*. To obtain this, we say that when a value has a type τ , then it has also all the types τ' such that τ is a subtype of τ' . This leads to our third informal type rule:

[Rule3] if $a : \tau$ and $\tau \leq \tau'$ then $a : \tau'$

If we define the function:

val age(x: object): int = x.age

we can meaningfully compute *age(mycar)* as, by [Rule3] *mycar* has the type required by *age*. Indeed *mycar* has the types *car*, *vehicle*, *machine*, *object*, the empty record type and many other ones.

When is it meaningful to apply a function to an argument? This is determined by the following rules:

[Rule4] if $f: \sigma \rightarrow \tau$ and $a: \sigma$ then $f(a)$ is meaningful, and $f(a): \tau$

[Rule5] if $f: \sigma \rightarrow \tau$ and $a: \sigma'$, where $\sigma' \leq \sigma$ then $f(a)$ is meaningful, and $f(a): \tau$

[Rule5] is just a consequence of [Rule3] and [Rule4]. From [Rule3] we can deduce that $a: \sigma$; than it is certainly meaningful to compute $f(a)$ as $f: \sigma \rightarrow \tau$.

The conventional *subclass* relation is usually only defined on objects or classes. Our *subtype* relation also extends naturally to functional types. Consider the function

serial_number: int → car

We can argue that *serial_number* returns vehicles, as all cars are vehicles. In general, all *car*-valued function are also *vehicle*-valued functions, so that for any domain type t we can say that $t \rightarrow car$ (an appropriate domain of functions from t to *car*) is a subtype of $t \rightarrow vehicle$:

$t \rightarrow car \leq t \rightarrow vehicle$ because $car \leq vehicle$

Now consider the function:

speed: vehicle → int

As all cars are vehicles, we can use this function to compute the speed of a car. Hence *speed* is also a function from *car* to integer. In general every function on vehicles is also a function on cars, and we can say that $vehicle \rightarrow int$ is a subtype of $car \rightarrow int$:

$vehicle \rightarrow t \leq car \rightarrow t$ because $car \leq vehicle$

Something interesting is happening here: note how the subtype relation is inverted on the left hand side of the arrow. This happens because of the particular meaning we are giving to the \rightarrow operator, as explained formally in the following sections. We are assuming a universal value domain V of all computable values. Every function f is a function from V to V , written $f: V \rightarrow V$, where " \rightarrow " is the conventional continuous function space. By $f: \sigma \rightarrow \tau$ we indicate a function $f: V \rightarrow V$ which whenever given an element of $\sigma \subseteq V$ returns an element of $\tau \subseteq V$ (nothing is asserted about the behavior of f outside σ).

Given any function $f: \sigma \rightarrow \tau$ from some domain σ to some codomain τ , we can always consider it as a function from some smaller domain $\sigma' \subseteq \sigma$ to some bigger codomain $\tau' \supseteq \tau$. For example a function $f: vehicle \rightarrow vehicle$ can be used in the context $age(f(mycar))$, where it is used as a function $f: car \rightarrow object$ (the application $f(mycar)$ makes sense because every car is a vehicle; $v = f(mycar)$ is a vehicle; hence it makes sense to compute $age(v)$ as every vehicle is an object).

The general rule of subtyping among functional types can be expressed as follows:

[Rule6] if $\sigma' \leq \sigma$ and $\tau \leq \tau'$ then $\sigma \rightarrow \tau \leq \sigma' \rightarrow \tau'$

As we said, the subtype relation extends to higher types. For example, the following is a definition of a function *mycar_attribute* which takes any integer-valued function on cars and applies it to my car.

val mycar_attribute(f: car → int): int = f(mycar)

We can then apply it to functions of any type which is a subtype of $car \rightarrow int$, e.g., $age: object \rightarrow int$. (Why? Because *car* is a subtype of *object*, hence $object \rightarrow int$ is a subtype of

$car \rightarrow int$, [Rule6] hence $(mycar_attribute: (car \rightarrow int) \rightarrow int)(age: object \rightarrow int)$ makes sense [Rule5]).

$mycar_attribute(age) \equiv 4$

$mycar_attribute(speed) \equiv 140$

Up to now we proceeded by assigning certain types to certain values. However the subtype relation has a very strong intuitive flavor of **inclusion** of types considered as sets of objects, and we want to justify our type assignments on semantic grounds.

Semantically we could regard the type *vehicle* as the set of all the records with a field *age* and a field *speed* having the appropriate types, but then cars would not belong to the set of vehicles as they have three fields while vehicles have two. To obtain the inclusion that we intuitively expect, we must say that the type *vehicle* is the set of all records which have *at least* two fields as above, but may have other fields. In this sense a car is a vehicle, and the set of all cars is included in the set of all vehicles, as we might expect. Some care is however needed to define these "sets", and this will be done formally in the following sections.

Record types can have a large number of fields, hence we need some notation for quickly defining a subtype of some record type, without having to list again all the fields of the record type. The following three sets of definitions are equivalent:

$type\ object = (age: int)$
 $type\ vehicle = (age: int, speed: int)$
 $type\ machine = (age: int, fuel: string)$
 $type\ car = (age: int, speed: int, fuel: string)$

$type\ object = (age: int)$
 $type\ vehicle = object\ and\ (speed: int)$
 $type\ machine = object\ and\ (fuel: string)$
 $type\ car = vehicle\ and\ machine$

$type\ object = (age: int)$
 $type\ car = object\ and\ (speed: int, fuel: string)$
 $type\ vehicle = car\ ignoring\ fuel$
 $type\ machine = car\ ignoring\ speed$

The *and* operator forms the union of the fields of two record types; if two record types have some labels in common (like in *vehicle and machine*), then the corresponding types must match. At this point we do not specify exactly what "match" means, except that in the example above "matching" is equivalent to "being the same". In its full generality, *and* corresponds to a \sqcap operation on type expressions, as explained in a later section.

The *ignoring* operator simply eliminates a component from a record type; it is undefined on other types.

4. Variants

The two basic non-functional data type constructions in denotational semantics are cartesian products and disjoint sums. We have seen that inheritance can be expressed as a subtype relation on record types, which then extends to higher types. Record types are just labeled cartesian products, and by analogy we can ask whether there is some similar notion deriving from labeled disjoint sums.

A labeled disjoint sum is called here a *variant*. A variant type looks very much like a record type: it is an unordered set of label-type pairs, enclosed in brackets instead of parentheses:

$type\ int_or_bool = [a: int, b: bool]$

An element of a variant type is a labeled value, where the label is one of the labels in the variant type, and the value has a type matching the type associated with that label. A element of *int_or_bool* is either an integer labeled *a* or a boolean labeled *b*.

$[a = 3] : \text{int_or_bool}$

$[b = \text{true}] : \text{int_or_bool}$

The basic operations on variants are *is*, which tests whether a variant object has a particular label, and *as*, which extracts the contents of a variant object having a particular label:

$[a = 3] \text{ is } a \quad \equiv \quad \text{true}$

$[a = 3] \text{ is } b \quad \equiv \quad \text{false}$

$[a = 3] \text{ as } a \quad \equiv \quad 3$

$[a = 3] \text{ as } b \quad \text{fails}$

A variant type σ is a subtype of a variant type τ (written $\sigma \leq \tau$) if τ has all the labels of σ and correspondingly matching types. Hence *int_or_bool* is a subtype of $[a: \text{int}, b: \text{bool}, c: \text{string}]$.

When the type associated to a label is *unit* (the trivial type, whose only defined element is *nil*), we can omit the type altogether; a variant type where all fields have *unit* type is also called an enumeration type. The following examples deal with enumeration types.

$\text{type precious_metal} = [\text{gold}, \text{silver}] \quad (\text{i.e. } [\text{gold}: \text{unit}, \text{silver}: \text{unit}])$

$\text{type metal} = [\text{gold}, \text{silver}, \text{steel}]$

A value of an enumeration type, e.g. $[\text{gold} = \text{nil}]$, can similarly be abbreviated by omitting the "*= nil*" part, e.g. $[\text{gold}]$.

A function returning a precious metal is also a function returning a metal, hence:

$t \rightarrow \text{precious_metal} \leq t \rightarrow \text{metal} \quad \text{because} \quad \text{precious_metal} \leq \text{metal}$

A function working on metals will also work on precious metals, hence:

$\text{metal} \rightarrow t \leq \text{precious_metal} \rightarrow t \quad \text{because} \quad \text{precious_metal} \leq \text{metal}$

It is evident that [Rule6] holds unchanged for variant types. This justifies the use of the symbol \leq for both record and variant subtyping. Semantically the subtype relation on variants is mapped to set inclusion, just as in the case of records: *metal* is a set with three defined elements $[\text{gold}]$, $[\text{silver}]$ and $[\text{steel}]$, and *precious_metal* is a set with two defined elements $[\text{gold}]$ and $[\text{silver}]$.

There are two ways of deriving variant types from previously defined variant types. We could have defined *metal* and *precious_metal* as:

$\text{type precious_metal} = [\text{gold}, \text{silver}]$

$\text{type metal} = \text{precious_metal} \text{ or } [\text{steel}]$

or as:

$\text{type metal} = [\text{gold}, \text{silver}, \text{steel}]$

$\text{type precious_metal} = \text{metal dropping steel}$

The *or* operator makes a union of the cases of two variant types, and the *dropping* operator removes a case from a variant type. The precise definition of these operators is contained in a later section.

5. Multiple Inheritance

In the framework described so far, we can recognize some of the features of what is called *multiple inheritance* between objects, e.g. a car has (inherits) all the attributes of *vehicle* and of *machine*. Some aspects are however unusual; for example the inheritance relation only depends on the structure of objects and need not be declared explicitly. In general, we are not trying to *explain* existing inheritance schemes (e.g. Smalltalk) in detail, but rather trying to present a new perspective on the problem.

We are not aware of other languages where typechecking coexists with multiple inheritance and higher order functions, with the exception of Galileo [Albano 84] which was developed in conjunction with this work. The OBJ language [Goguen 84] comes close with its strongly typed multiple inheritance and first order functions.

Typechecking provides compile-time protection against obvious bugs (like applying the *speed* function to a machine which is not a vehicle), and other less obvious mistakes. Complex type hierarchies can be built where "everything is also something else", and it can be difficult to remember which objects support which messages.

The subtype relation only holds on types, and there is no similar relation on objects. Thus we cannot model directly the *subobject* relation used by, for example, Omega [Attardi 81], where we could define the class of gasoline cars as the cars with fuel equal to "gasoline".

However, in simple cases we can achieve the same effect by turning certain sets of values into variant types. For example, instead of having the fuel field of a machine be a string, we could redefine:

```
type fueltype      = [coal, gasoline, electricity]
type machine       = (age: int, fuel: fueltype)
type car           = (age: int, speed: int, fuel: fueltype)
```

Now we can have:

```
type gasoline_car  = (age: int, speed: int, fuel: [gasoline])
type combustion_car = (age: int, speed: int, fuel: [gasoline, coal])
```

and we have $\text{gasoline_car} \leq \text{combustion_car} \leq \text{car}$. Hence a function over combustion cars, for example, will accept a gasoline car as a parameter, but will give a compile-time type error when applied to electrical cars.

It is often the case that a function which is a field of a record has to refer to other components of the same record. In Smalltalk this is done by referring to the whole record (i.e. object) as *self*, and then selecting the desired components out of that. In Simula there is a similar concept called *this*.

This self-referential capability can be obtained as a special case of the *rec* operator which we are about to introduce. *rec* is used to define recursive functions and data. For example, the recursive factorial function can be written as:

```
rec fact: int → int. λn: int. if n=0 then 1 else n*fact(n-1)
```

(This is an expression, not a declaration.)

In order to prevent looping in case of call-by-value evaluations, the body of *rec* is restricted to be a constant, a record, a variant or a function (or, in general, any data constructor present in the language) [Morris 80].

Examples of circular data definitions are extremely common in object-oriented programming. In the following example, a functional component of a record refers to "its" other components. The functional component *d*, below, is supposed to compute the distance of "this" *active_point* from any other *point* (or any other *active_point*, etc.).


```

type point = (x: real, y: real)
type active_point = point and (d: point → real)
val make_active_point(px: real, py: real): active_point =
  rec self: active_point.
    (x = px, y = py,
     d = λp: point. sqrt((p.x-self.x)**2 + (p.y-self.y)**2))

```

Objects often have **private** variables, which are useful to maintain and update the local state of an object while preventing arbitrary external interference. Here is a counter object which starts from some fixed number and can only be incremented one step at a time. *cell n* is an updatable cell whose initial contents is *n*; a cell can be updated by *:=* and its contents can be extracted by *get*.

```

type counter = (increment: int → unit, fetch: unit → int)
val make_counter(n: int) =
  let count = cell n
  in (increment = λn: int. count := (get count)+1,
      fetch = λnil: unit. get count)

```

Private variables are obtained in full generality by the above well known static scoping technique.

6. Expressions

We now begin the formal treatment of multiple inheritance. First, we define a simple applicative language supporting inheritance (side effects could be treated without introducing any new concept, but they make the formal treatment more complicated). Then a denotational semantics is presented, in a domain of values *V*. Certain subsets of *V* are regarded as types, and inheritance corresponds directly to set inclusion among types. A type inference system and a typechecking algorithm are then presented. The soundness of the algorithm is proved by showing that the algorithm is consistent with the inference system, and that the inference system is in turn consistent with the semantics.

Our language is typed lambda calculus with records and variants. The following notation is often used for records (and similarly for record and variant types):

$$(a_1=e_1, \dots, a_n=e_n) \equiv (a_i=e_i) \quad i \in 1..n$$

$$(a_1=e_1, \dots, a_n=e_n, a'_1=e'_1, \dots, a'_m=e'_m) \equiv (a_i=e_i, a'_j=e'_j) \quad i \in 1..n, j \in 1..m$$

Here is the syntax of expressions and type expressions:

<i>e</i> ::=	expressions
<i>x</i>	identifiers
<i>b</i>	constants
if <i>e</i> then <i>e</i> else <i>e</i>	conditionals
(<i>a_i</i> = <i>e_i</i>) <i>e.a</i>	records (<i>i</i> ∈ 1..<i>n</i>, <i>n</i> ≥ 0)
[<i>a</i> = <i>e</i>] <i>e is a</i> <i>e as a</i>	variants
λ <i>x</i> : τ. <i>e</i> <i>e e</i>	functions
rec <i>x</i> : τ. <i>e</i>	recursive data
<i>e</i> : τ	type specs
(<i>e</i>)	
τ ::=	type expressions
ι	type constants
(<i>a_i</i> : τ _{<i>i</i>})	record types (<i>i</i> ∈ 1..<i>n</i>, <i>n</i> ≥ 0)
[<i>a_i</i> : τ _{<i>i</i>}]	variant types (<i>i</i> ∈ 1..<i>n</i>, <i>n</i> ≥ 0)
τ → τ	function types
(τ)	

where $i \neq j \Rightarrow a_i \neq a_j$
 take $v_0 = \text{unit}, v_1 = \text{bool}, v_2 = \text{int}, \text{ etc.}$

Syntactic restriction: the body e of $\text{rec } x: \tau. e$ can only be a constant, a record, a variant, a lambda expression, or another rec obeying this restriction.

Labels a , and identifiers x have the same syntax, but are distinguishable by the syntactic context. Among the type constants we have unit (the domain with one defined element) bool and int . Among the constants we have nil (of type unit), booleans ($\text{true}, \text{false}$) and numbers ($0, 1, \dots$).

Global definitions of values and types are introduced by the syntax:

$$d ::=$$

$$\text{val } x = e \mid$$

$$\text{type } x = \tau$$

where the type definitions are meant as simple abbreviations.

Standard abbreviations are:

$$\text{let } x: \tau = e \text{ in } e' \quad \text{for} \quad (\lambda x: \tau. e') e$$

$$f(x: \tau): \tau' = e \quad \text{for} \quad f = \lambda x: \tau. (e: \tau')$$

$$\text{rec } f(x: \tau): \tau' = e \quad \text{for} \quad f = \text{rec } f: \tau \rightarrow \tau'. \lambda x: \tau. e$$

(the last two abbreviations can only appear after a let or a val).

Record and variant type expressions are unordered, so for any permutation π_n of $1..n$, we identify:

$$(a_i: \tau_i) \quad \equiv \quad (a_{\pi_n(i)}: \tau_{\pi_n(i)}) \quad i \in 1..n$$

$$[a_i: \tau_i] \quad \equiv \quad [a_{\pi_n(i)}: \tau_{\pi_n(i)}] \quad i \in 1..n$$

7. The Semantic Domain

The semantics of expressions is given in the recursively defined domain V of values. The domain operators used below are coalesced sum (+), cartesian product (\times), continuous function space (\rightarrow) and finite functions (\rightarrow_{fn} , explained later).

$$V = B_0 + B_1 + \dots + R + U + F + W$$

$$R = L \rightarrow_{fn} V$$

$$U = L \times V$$

$$F = V \rightarrow V$$

$$W = \{\perp, w\}$$

where L is a countable flat domain of character strings, called *labels*, and B_i are flat domains of basic values. We take:

$$B_0 = O = \{\perp, \text{nil}\}$$

$$B_1 = T = \{\perp, \text{true}, \text{false}\}$$

$$B_2 = N = \{\perp, 0, 1, \dots\}$$

W is a domain which contains a single defined element w , the *wrong* value. The value w is used to model run-time **type errors** (e.g. trying to apply an integer as if it were a function) which we want a compiler to trap before execution. It is not used to model run-time **exceptions** (like trying to extract the head of an empty list); in our context these can only be generated by the as operator. Run-time exceptions should be modeled by an extra summand of V , but for simplicity we shall instead use the undefined element \perp . The name *wrong* is used to denote w as a member of V (instead of simply a member of W).

$R = L \rightarrow_{fin} V$ is the domain of *records*, which are associations of values to labels. We are only interested in finite associations, so we define $L \rightarrow_{fin} V = \{r \in L \rightarrow V \mid \{a \mid r(a) \neq \text{wrong}\} \text{ is finite}\}$.

$U = L \times V$ is the domain of *variants* which are pairs $\langle l, v \rangle$ with a label l and a value v .

$F = V \rightarrow V$ are the continuous functions from V to V , used to give semantics to lambda expressions.

8. Semantics of Expressions

The semantic function is $\mathcal{E} \in Exp \rightarrow Env \rightarrow V$, where Exp are syntactic expressions according to our grammar, and $Env = Id \rightarrow V$ are environments for identifiers. The semantics of basic values is given by $\mathcal{B} \in Exp \rightarrow V$, whose obvious definition is omitted; b_{ij} is the j -th element of the basic domain B_i .

$$\begin{aligned}
\mathcal{E}[x]v &= v[x] \\
\mathcal{E}[b_{ij}]v &= \mathcal{B}[b_{ij}] \\
\mathcal{E}[\text{if } e \text{ then } e' \text{ else } e'']v &= \\
&\quad \text{if } \mathcal{E}[e]v \in T \text{ then } (\text{if } (\mathcal{E}[e]v \mid T) \text{ then } \mathcal{E}[e']v \text{ else } \mathcal{E}[e'']v) \text{ else wrong} \\
\mathcal{E}[(a_1 = e_1, \dots, a_n = e_n)]v &= \\
&\quad \text{if } \mathcal{E}[e_1]v \in W \text{ or } \dots \text{ or } \mathcal{E}[e_n]v \in W \text{ then wrong} \\
&\quad \text{else } (\lambda l. \text{ if } l=a_1 \text{ then } \mathcal{E}[e_1]v \text{ else } \dots \text{ if } l=a_n \text{ then } \mathcal{E}[e_n]v \text{ else wrong}) \text{ in } V \\
\mathcal{E}[e.a]v &= \text{if } \mathcal{E}[e]v \in R \text{ then } (\mathcal{E}[e]v \mid R)(a) \text{ else wrong} \\
\mathcal{E}[[a=e]]v &= \text{if } \mathcal{E}[e]v \in W \text{ then wrong else } \langle a, \mathcal{E}[e]v \rangle \text{ in } V \\
\mathcal{E}[e \text{ is } a]v &= \text{if } \mathcal{E}[e]v \in U \text{ then } \text{fst}(\mathcal{E}[e]v \mid U) = a \text{ else wrong} \\
\mathcal{E}[e \text{ as } a]v &= \\
&\quad \text{if } \mathcal{E}[e]v \in U \text{ then } (\text{let } \langle b, v \rangle \text{ be } (\mathcal{E}[e]v \mid U) \text{ in if } b=a \text{ then } v \text{ else } \perp) \text{ else wrong} \\
\mathcal{E}[\lambda x: \tau. e]v &= (\lambda v. \mathcal{E}[e]v\{v/\llbracket x \rrbracket\}) \text{ in } V \\
\mathcal{E}[e \ e']v &= \\
&\quad \text{if } \mathcal{E}[e]v \in F \text{ then } (\text{if } \mathcal{E}[e']v \in W \text{ then wrong else } (\mathcal{E}[e]v \mid F)(\mathcal{E}[e']v)) \text{ else wrong} \\
\mathcal{E}[\text{rec } x: \tau. e]v &= Y((\lambda v. \mathcal{E}[e]v\{v/\llbracket x \rrbracket\}) \text{ in } V) \\
\mathcal{E}[e: \tau]v &= \mathcal{E}[e]v
\end{aligned}$$

Comments on the equations:

- $d \text{ in } V$ (where $d \in D$ and D is a summand of V) is the injection of d in the appropriate summand of V . Hence $d \text{ in } V \in V$ and $\perp \text{ in } V = \perp$. This is not to be confused with the *let...be...in...* notation for local variables.
- $v \in D$ (where $v \in V$ and D is a summand of V) is a function yielding: \perp if $v = \perp$; *true* if $v = d \text{ in } V$ for some $d \in D$; *false* otherwise.
- $v \mid D$ (where D is a summand of V) is a function yielding: d if $v = d \text{ in } V$ for some $d \in D$; \perp otherwise.
- *fst* extracts the first element of a pair, *snd* extracts the second one.
- \mathcal{E} defines a call by value semantics.

Intuitively, a well-typed program will never return the *wrong* value at run-time. For example, consider the second occurrence of *wrong* in the semantics of records. The typechecker will make sure that any record selection will operate on records having the appropriate field, hence that instance of *wrong* will never be returned. A similar reasoning applies to all the instances of *wrong* in the semantics: *wrong* is a run-time type error which can be detected at compile-time. Run-time exceptions which cannot be detected are represented as \perp ; the only instance of this in the above

semantics is in the equation for e as a .

Formally, we proceed by defining \mathcal{E} (so that it satisfies the above intuitions about run-time errors), then we define " e is semantically well-typed" to mean " $\mathcal{E}[e]v \neq \text{wrong}$ ", and later we give an algorithm which statically checks well-typing.

9. Semantics of Type Expressions

The semantics of types is given in the *weak ideal model* [MacQueen 84] $\mathcal{J}(V)$ (the set of non-empty weak ideals which are subset of V and do not contain *wrong*). $\mathcal{J}(V)$ is a lattice of domains, where the ordering is set inclusion. $\mathcal{J}(V)$ is closed under union and intersection, as well as the usual domain operations.

$$\mathcal{D}[\iota_i] = B_i \text{ in } V$$

$$\mathcal{D}[(a_i: \tau_i)] = \bigcap_i \{r \in R \mid r(a_i) \in \mathcal{D}[\tau_i]\} \text{ in } V \quad (\text{where we take } \mathcal{D}[] = R \text{ in } V)$$

$$\mathcal{D}[[a_i: \tau_i]] = \bigcup_i \{<a_i, v> \in U \mid v \in \mathcal{D}[\tau_i]\} \text{ in } V \quad (\text{where we take } \mathcal{D}[[]] = \{\perp\})$$

$$\mathcal{D}[\sigma \rightarrow \tau] = \{f \in F \mid v \in \mathcal{D}[\sigma] \Rightarrow f(v) \in \mathcal{D}[\tau]\} \text{ in } V$$

where $\mathcal{D} \text{ in } V = \{d \text{ in } V \mid d \in D\}$

THEOREM (\mathcal{D} properties)

$$\forall \tau. \perp \in \mathcal{D}[\tau]$$

$$\forall \tau, v. v \in \mathcal{D}[\tau] \Rightarrow v \neq \text{wrong}$$

The *wrong* value is deliberately left out of the type domains so that if a value has a type, then that value is not a run-time type error. Another way of saying this is that *wrong* has no type.

10. Type Inclusion

A subtyping relation can be defined syntactically on the structure of type expressions. This definition formalizes our initial discussion of subtyping for records, variants and functions.

$$\iota_i \leq \iota_i$$

$$(a_i: \sigma_i, a_j: \sigma_j) \leq (a_i: \sigma'_i) \iff \sigma_i \leq \sigma'_i \quad (i \in 1..n, n \geq 0; j \in 1..m, m \geq 0)$$

$$[a_i: \sigma_i] \leq [a_i: \sigma'_i, a_j: \sigma'_j] \iff \sigma_i \leq \sigma'_i \quad (i \in 1..n, n \geq 0; j \in 1..m, m \geq 0)$$

$$\sigma \rightarrow \tau \leq \sigma' \rightarrow \tau' \iff \sigma' \leq \sigma \text{ and } \tau \leq \tau'$$

no other type expressions are in the \leq relation

PROPOSITION:

\leq is a partial order

It is possible to extend type expressions by two constants *anything* and *nothing*, such that *nothing* $\leq \tau \leq \text{anything}$ for any τ . Then, \leq defines a lattice structure on type expressions, which is a sublattice of $\mathcal{J}(V)$. Although this is mathematically appealing, we have chosen not to do it in view of our intended application. For example, the expression *if x then 3 else true*, should produce a type error because of a conflict between *int* and *bool* in the two branches of the conditional. If we have the full lattice of type expression, it is conceivable to return *anything* as the type of the expression above, and carry on typechecking. This is bad for two reasons. First, no use can be made of objects of type *anything* (at least in the present framework). Second, type errors are difficult to localize as their presence is only made manifest by the eventual occurrence of *anything* or *nothing* in the resulting type.

As we said, the ordering of domains in the $\mathcal{J}(V)$ model is set inclusion. This allows us to give a very direct semantics to subtyping, as simple set inclusion of domains.

THEOREM (Semantic Subtyping)

$$\tau \leq \tau' \Rightarrow \mathcal{D}[\tau] \subseteq \mathcal{D}[\tau']$$

The proof is by induction on the structure of τ and τ' .

11. Type Inference Rules

In this section we formally define the notion of a **syntactically well-typed** expression. An expression is well-typed when a type can be deduced for it, according to a set of type rules forming an **inference system**. If no type can be deduced, then the expression is said to contain type errors.

In general, many types can be deduced for the same expression. Provided that the inference system is consistent, all those types are in some sense compatible. A typechecking algorithm can then choose any of the admissible types as *the* type of an expression, with respect to that algorithm (in some type systems there may be a *best*, or *most general*, or *principal* type). Inference systems can be shown to be consistent with respect to the semantics of the language, as we shall see at the end of this section.

Hence, here is the inference system for our language. It is designed so that (1) it contains exactly one type rule for each syntactic construct; (2) it satisfies the intuitive subtyping property expressed by the syntactic subtyping theorem below; and (3) it satisfies a semantic soundness theorem, relating it to the semantics of the language.

The use of the subtyping predicate \leq is critical in many type rules. However it should be noted that subtyping does not affect the fundamental λ -calculus typing rules [ABS] and [COMB]. This indicates that this style of subtyping merges naturally with functional types.

[VAR]	$A.x: \tau \vdash x: \tau'$	where $\tau \leq \tau'$
[BAS]	$A \vdash b_{ij}: \iota_i$	
[COND]	$\frac{A \vdash e: \text{bool} \quad A \vdash e': \tau \quad A \vdash e'': \tau}{A \vdash (\text{if } e \text{ then } e' \text{ else } e''): \tau}$	
[RECORD]	$\frac{A \vdash e_1: \tau_1 \quad \dots \quad A \vdash e_n: \tau_n}{A \vdash (a_1 = e_1, \dots, a_n = e_n): (a_i: \tau_i)}$	where $i \in I \subseteq 1..n$
[DOT]	$\frac{A \vdash e: (\dots a: \tau \dots)}{A \vdash e.a: \tau}$	
[VARIANT]	$\frac{A \vdash e: \tau}{A \vdash [a = e]: [\dots a: \tau \dots]}$	
[IS]	$\frac{A \vdash e: [\dots a: \sigma \dots]}{A \vdash (e \text{ is } a): \text{bool}}$	
[AS]	$\frac{A \vdash e: [\dots a: \tau \dots]}{A \vdash (e \text{ as } a): \tau}$	
[ABS]	$\frac{A.x: \sigma \vdash e: \tau}{A \vdash (\lambda x: \sigma. e): \sigma \rightarrow \tau}$	
[COMB]	$\frac{A \vdash e: \sigma \rightarrow \tau \quad A \vdash e': \sigma}{A \vdash (e \ e'): \tau}$	
[REC]	$\frac{A.x: \sigma \vdash e: \rho}{A \vdash (\text{rec } x: \sigma. e): \tau}$	where $\rho \leq \sigma$ and $\rho \leq \tau$

$$[\text{SPEC}] \quad \frac{A \vdash e: \sigma}{A \vdash (e: \sigma): \tau} \quad \text{where } \sigma \leq \tau$$

Some comments on the rules:

- A is a list of assumptions for variables, of the form $x: \tau$. We use the notation $A \equiv A'.x: \tau$ to single out the assumption at the head of A . The handling of nested scopes for variables requires some care: a list of assumptions can be permuted as long as assumptions involving the same variable are not swapped.
- If there are some nontrivial inclusions in the basic types (e.g. $\text{int} \leq \text{real}$) then [BAS] must be changed to $A \vdash b_{ij}: \tau$ where $\iota_j \leq \tau$.
- In [RECORD], the derived record type can have fewer fields than the corresponding record object.
- In [VARIANT], the derived variant type can have any number of fields, as long as it includes a field corresponding to the variant object.
- In [IS], the antecedent could safely be changed to have an arbitrary variant type. It would however seem strange to be able to test the existence of a label but fail to typecheck when trying to extract out of that same label (see [AS]).
- The [IS] rule assumes that the set of basic types does not contain a supertype of *bool*, otherwise a more refined rule is needed. Similarly, [COND] assumes that there are no subtypes of *bool*.

The basic syntactic property of this inference system is expressed in the syntactic subtyping theorem below: if an expression has a type τ , and τ is a subtype of τ' , then the expression has also type τ' . The lemma is required to prove the [ABS] case of the theorem. Both the lemma and the theorem are proved by induction on the structure of the derivations.

LEMMA (Syntactic Subtyping):

$$A.x: \sigma \vdash e: \tau \text{ and } \sigma' \leq \sigma \implies A.x: \sigma' \vdash e: \tau$$

THEOREM (Syntactic Subtyping):

$$A \vdash e: \tau \text{ and } \tau \leq \tau' \implies A \vdash e: \tau'$$

The next theorem states the soundness of the type system with respect to the semantics: if it is possible to deduce that e has type τ , then the value denoted by e belongs to the domain denoted by τ . A list of assumptions A agrees with an environment ν if for any x , $\nu[x] \in \mathcal{D}[\tau]$ $\iff A = A'.x: \tau$.

THEOREM (Semantic Soundness):

$$\text{If } \nu \text{ agrees with } A \text{ and } A \vdash e: \tau \text{ then } \mathcal{E}[e]\nu \in \mathcal{D}[\tau]$$

The proof is by induction on the structure of the derivation of $A \vdash e: \tau$, using the semantic subtyping and \mathcal{D} -properties theorems.

Another way of looking at this theorem is that if e is syntactically well-typed (i.e. for some τ , $A \vdash e: \tau$), then it is also semantically well-typed (i.e. for some τ , $\mathcal{E}[e]\nu \in \mathcal{D}[\tau]$), which implies that $\mathcal{E}[e]\nu \neq \text{wrong}$.

12. Join and Meet Types

In the examples at the beginning of the paper we used the *and* and *or* type operators, and we are now going to need them in the definition of the typechecking algorithm. However those operators are not part of the syntax of type expressions, nor are *ignoring* and *dropping*.

This is because the above operators only work on restricted kinds of type expressions. Applied to arbitrary type expressions they either are undefined, or can be eliminated by a normalization process. Hence, if we have a type expression containing the above operators we can process the expression checking that the operators can be indeed used in that context, and in such case we

can normalize them away obtaining a normal type expression.

The *and* operator is interpreted as a meet operation on types (written \downarrow), and *or* is interpreted as join (written \uparrow). Joins and meets are taken in the partial order determined by \leq , when they exist.

The definition of the operators also immediately defines the normalization process which eliminates them:

$$\begin{aligned}
 \iota_i \uparrow \iota_i &= \iota_i \\
 (a_i: \tau_i, b_j: \sigma_j) \uparrow (a_i: \tau'_i, c_k: \rho_k) &= (a_i: \tau_i \uparrow \tau'_i) & (\forall j, k. b_j \neq c_k) \\
 [a_i: \tau_i, b_j: \sigma_j] \uparrow [a_i: \tau'_i, c_k: \rho_k] &= [a_i: \tau_i \uparrow \tau'_i, b_j: \sigma_j, c_k: \rho_k] & (\forall j, k. b_j \neq c_k) \\
 (\sigma \rightarrow \tau) \uparrow (\sigma' \rightarrow \tau') &= (\sigma \downarrow \sigma') \rightarrow (\tau \uparrow \tau') \\
 \tau \uparrow \tau' &\text{ undefined otherwise}
 \end{aligned}$$

$$\begin{aligned}
 \iota_i \downarrow \iota_i &= \iota_i \\
 (a_i: \tau_i, b_j: \sigma_j) \downarrow (a_i: \tau'_i, c_k: \rho_k) &= (a_i: \tau_i \downarrow \tau'_i, b_j: \sigma_j, c_k: \rho_k) & (\forall j, k. b_j \neq c_k) \\
 [a_i: \tau_i, b_j: \sigma_j] \downarrow [a_i: \tau'_i, c_k: \rho_k] &= [a_i: \tau_i \downarrow \tau'_i] & (\forall j, k. b_j \neq c_k) \\
 (\sigma \rightarrow \tau) \downarrow (\sigma' \rightarrow \tau') &= (\sigma \uparrow \sigma') \rightarrow (\tau \downarrow \tau') \\
 \tau \downarrow \tau' &\text{ undefined otherwise}
 \end{aligned}$$

$$\begin{aligned}
 (a_i: \tau_i) \text{ ignoring } a &= (a_j: \tau_j) & (i \in 1..n, j \in 1..n - \{k \mid a_k = a\}) \\
 \tau \text{ ignoring } a &\text{ undefined otherwise} \\
 [a_i: \tau_i] \text{ dropping } a &= [a_j: \tau_j] & (i \in 1..n, j \in 1..n - \{k \mid a_k = a\}) \\
 \tau \text{ dropping } a &\text{ undefined otherwise}
 \end{aligned}$$

PROPOSITION (\uparrow and \downarrow properties):

$$\begin{aligned}
 \uparrow \text{ and } \downarrow &\text{ are join and meet for } \leq, \text{ when defined} \\
 \mathcal{D}[\sigma \text{ and } \tau] &= \mathcal{D}[\sigma] \cap \mathcal{D}[\tau] & (\text{when defined}) \\
 \mathcal{D}[\tau \text{ ignoring } a] &= \{r \in R \mid ((r/a) \text{ in } V) \in \mathcal{D}[\tau]\} \text{ in } V & (\text{when defined}) \\
 \mathcal{D}[\sigma \text{ or } \tau] &= \mathcal{D}[\sigma] \cup \mathcal{D}[\tau] & (\text{when defined}) \\
 \mathcal{D}[\tau \text{ dropping } a] &= \mathcal{D}[\tau] - \{<a, v> \in U\} \text{ in } V & (\text{when defined})
 \end{aligned}$$

where $r/a = (\lambda b. \text{ if } b=a \text{ then } \perp \text{ else } r(b))$.

13. Typechecking

The typechecking function is $\mathcal{T} \in \text{Exp} \rightarrow \text{TypeEnv} \rightarrow \text{TypeExp}$, where *Exp* and *TypeExp* are respectively expressions and type expressions according to our grammar, and $\text{TypeEnv} = \text{Id} \rightarrow \text{TypeExp}$ are type environments for identifiers.

The following description is to be intended as a scheme for a program that returns a type expression denoting the type of a term, or fails in case of type errors. The *fail* word is a global jump-out: when a type error is detected the program stops. Similarly, typechecking fails when the \uparrow and \downarrow operations are undefined. When we assert that $\mathcal{T}[e]\mu = \tau$, we imply that the typechecking of e does not fail.

$$\mathcal{T}[\llbracket x \rrbracket]_{\mu} = \mu[\llbracket x \rrbracket]$$

$$\mathcal{T}[\llbracket b_{ij} \rrbracket]_{\mu} = \iota_i$$

$$\mathcal{T}[\text{if } e \text{ then } e' \text{ else } e'']_{\mu} = \text{if } \mathcal{T}[\llbracket e \rrbracket]_{\mu} = \text{bool} \text{ then } \mathcal{T}[\llbracket e' \rrbracket]_{\mu} \uparrow \mathcal{T}[\llbracket e'' \rrbracket]_{\mu} \text{ else fail}$$

$$\mathcal{T}[(a_1 = e_1, \dots, a_n = e_n)]_{\mu} = (a_1: \mathcal{T}[\llbracket e_1 \rrbracket]_{\mu}, \dots, a_n: \mathcal{T}[\llbracket e_n \rrbracket]_{\mu})$$

$$\mathcal{T}[e.a]_{\mu} = \text{if } \mathcal{T}[\llbracket e \rrbracket]_{\mu} = (\dots a: \tau \dots) \text{ then } \tau \text{ else fail}$$

$$\mathcal{T}[\llbracket a = e \rrbracket]_{\mu} = [a: \mathcal{T}[\llbracket e \rrbracket]_{\mu}]$$

$$\mathcal{T}[e \text{ is } a]_{\mu} = \text{if } \mathcal{T}[\llbracket e \rrbracket]_{\mu} = [\dots a: \tau \dots] \text{ then bool else fail}$$

$$\mathcal{T}[e \text{ as } a]_{\mu} = \text{if } \mathcal{T}[\llbracket e \rrbracket]_{\mu} = [\dots a: \tau \dots] \text{ then } \tau \text{ else fail}$$

$$\mathcal{T}[\lambda x: \tau. e]_{\mu} = \tau \rightarrow \mathcal{T}[\llbracket e \rrbracket]_{\mu\{\tau/\llbracket x \rrbracket\}}$$

$$\mathcal{T}[e \ e']_{\mu} = \text{if } \mathcal{T}[\llbracket e \rrbracket]_{\mu} = (\tau \rightarrow \tau') \text{ and } \mathcal{T}[\llbracket e' \rrbracket]_{\mu} \leq \tau \text{ then } \tau' \text{ else fail}$$

$$\mathcal{T}[\text{rec } x: \sigma. e]_{\mu} = \text{if } \mathcal{T}[\llbracket e \rrbracket]_{\mu\{\sigma/\llbracket x \rrbracket\}} = \tau \text{ and } \tau \leq \sigma \text{ then } \tau \text{ else fail}$$

$$\mathcal{T}[e: \sigma]_{\mu} = \text{if } \mathcal{T}[\llbracket e \rrbracket]_{\mu} = \tau \text{ and } \tau \leq \sigma \text{ then } \tau \text{ else fail}$$

This typechecking algorithm is correct with respect to the type inference system: if the algorithm succeeds and returns a type τ for an expression e , then it is possible to prove that e has type τ . A type environment μ agrees with a list of assumptions A if $\mu[\llbracket x \rrbracket] = \tau \iff A = A'.x: \tau$.

THEOREM (Syntactic Soundness):

if A agrees with μ and $\mathcal{T}[\llbracket e \rrbracket]_{\mu} = \tau$ then $A \vdash e: \tau$

The proof is by induction on the structure of e , using the properties of \uparrow , \downarrow and \leq .

Combining the syntactic soundness, semantic soundness and \mathfrak{D} -properties theorems we immediately obtain:

COROLLARY (Typechecking prevents type errors):

if $\mathcal{T}[\llbracket e \rrbracket]_{\mu} = \tau$ then $\mathcal{E}[\llbracket e \rrbracket]_{\nu} \neq \text{wrong}$ (where μ agrees with some A which agrees with ν)

i.e. if e can be successfully typechecked, then e cannot produce run-time type errors.

14. Conclusions

This work originated as an attempt to justify the multiple inheritance constructs present in the Galileo data base language [Albano 83] and to provide a sound typechecking algorithm for that language. I believe this paper adequately solves the basic problems, although some practical issues may require more work.

Parametric polymorphism has not been treated in this paper. The intention was to study multiple inheritance problems in the cleanest possible framework, without interaction of other features. This restriction also has the advantage of considerably simplifying proofs.

Some confusion may arise from the fact that languages like Smalltalk are often referred to as polymorphic languages. This is correct, if by polymorphism we mean that an object or a function can have many types. However it now appears that there are two subtly different kinds of polymorphism, which I'll call *horizontal* polymorphism (having to do with inheritance) and *vertical* polymorphism (the ordinary parametric kind).

These adjectives come from the following considerations (where α and β are type variables). The function $f \ x = x.a$ where $f: (a: \alpha) \rightarrow \alpha$ is polymorphic in two senses. It can be considered as a function of type $(a: \alpha, b: \beta) \rightarrow \alpha$ by extending its domain *horizontally*, and also as a function of type $(a: \beta \rightarrow \beta) \rightarrow (\beta \rightarrow \beta)$ by instantiating its domain *vertically*.

These two kinds of polymorphism are not incompatible. We have seen here that horizontal polymorphism can be explained in the semantic domains normally used for vertical polymorphism. Moreover the technical explanation of polymorphism is the same in both cases: domain

intersection. Merging these two kinds of polymorphism does not seem to introduce new semantic problems. The interactions of horizontal and vertical polymorphism in typechecking may however be complex, and this is left here as an open problem.

There are now several competing (although not totally independent) styles of parametric polymorphism, noticeably [Milner 78], [Reynolds 74, McCracken 84] and [MacQueen 84]. Inheritance is orthogonal to all of these, so it seems better to study it independently, at least initially. However the interactions will have to be investigated in order to obtain the best of all possible worlds.

15. Related work and acknowledgements

Inheritance is not a totally new idea. Not surprisingly there are several independent and/or related works: my apologies for not referencing some of them in this or earlier versions of this paper. I would like to mention here [Reynolds 80, Oles 84] which expose the same basic semantic ideas in a different formal framework, [Ait-Kaci 83] again very similar ideas in a Prolog-related framework, [Mitchell 84] this time different, but related, ideas in the same formal framework, [Goguen 84] whose OBJ system implements a multiple inheritance typechecker, and [Fairbairn 82] whose algorithms for typechecking quantified types may be relevant to unifying inheritance with polymorphism.

Finally, I would like to thank Dave MacQueen for many discussions, and Antonio Albano and Renzo Orsini for motivating me to carry out this work.

16. References

- [Ait-Kaci 83] H.Ait-Kaci: "Outline of a calculus of type subsumptions", Technical report MS-CIS-83-34, Dept of Computer and Information Science, The Moore School of Electrical Engineering, University of Pennsylvania, August 1983.
- [Albano 83] A.Albano, L.Cardelli, R.Orsini: "Galileo: a strongly typed, interactive conceptual language", Bell Labs Technical Memorandum TM 83-11271-2.
- [Attardi 81] G.Attardi, M.Simi: "Semantics of inheritance and attributions in the description system Omega", M.I.T. A.I. Memo 642, August 81.
- [Dahl 66] O.Dahl, K.Nygaard: "Simula, an Algol-based simulation language", Comm. ACM, Vol 9, pp. 671-678, 1966.
- [Deutsch 84] P.Deutsch: "An efficient implementation of Smalltalk-80", Proc. Popl 84.
- [Goguen 84] J.A.Goguen, J.Meseguer: "Equality, types, modules and generics for logic programming", Second International Logic Programming Conference, Uppsala University, Sweden, July 1984.
- [Goldberg 83] A.Goldberg, D.Robson: "Smalltalk-80. The language and its implementation", Addison-Wesley, 1983.
- [McCracken 84] N.McCracken: "The typechecking of programs with implicit type structure", this conference.
- [MacQueen 84] D.B.MacQueen, R.Seti, G.D.Plotkin: "An ideal model for recursive polymorphic types", Proc. Popl 84.
- [Milner 78] R.Milner: "A theory of type polymorphism in programming", Journal of Computer and System Science 17, pp. 348-375, 1978.
- [Mitchell 84] J.C.Mitchell: "Coercion and type inference", Proc. Popl 84.
- [Oles 84] F.J.Oles: "Type algebras, functor categories, and block structure", to appear in "Algebraic semantics", M.Nivat and J.C.Reynolds ed., Cambridge University Press 1984.
- [Fairbairn 82] J.Fairbairn: "Ponder and its type system", Technical report No 31, Nov 82, University of Cambridge, Computer Laboratory.
- [Reynolds 74] J.C.Reynolds: "Towards a theory of type structure", in "Colloquium sur la

programmation" pp. 408-423, Springer-Verlag Lecture Notes in Computer Science, n.19, 1974.

[Reynolds 80] J.C.Reynolds: "Using category theory to design implicit type conversions and generic operators", in "Semantics-directed compiler generation", Lecture Notes in Computer Science 94, pp. 211-258, Springer-Verlag 1980.

[Morris 80] L.Morris, J.Schwarz: "Computing cyclic list structures", Conference Record of the 1980 Lisp Conference, pp.144-153.

[Steels 83] L.Steels: "Orbit: an applicative view of object-oriented programming", in: Integrated Interactive Computing Systems, pp. 193-205, P.Degano and E.Sandewall editors, North-Holland 1983.

[Weinreb 81] D.Weinreb, D.Moon: "Lisp machine manual", Fourth Edition, Chapter 20: "Objects, Message Passing, and Flavors", Symbolics Inc., 1981.