



PUC

Monografias em Ciência da Computação
nº 27/92

A Formal Specification for a Hierarchy of Collections

Roberto Ierusalimschy

Departamento de Informática

**PONTIFÍCIA UNIVERSIDADE CATÓLICA DO RIO DE JANEIRO
RUA MARQUÊS DE SÃO VICENTE, 225 - CEP 22454-970
RIO DE JANEIRO - BRASIL**

PUC RIO - DEPARTAMENTO DE INFORMÁTICA

Monografias em Ciência da Computação, Nº 27/92

Editor: Carlos J. P. Lucena

Julho, 1992

A Formal Specification for a Hierarchy of Collections *

Roberto Ierusalimschy

* This work has been sponsored by the Secretaria de Ciência e Tecnologia da Presidência da República Federativa do Brasil.

Abstract

This paper presents a formal specification for a hierarchy of types similar to the Collection hierarchy presented by the Smalltalk language. The specification method is an extension of VDM that supports inheritance of specifications, with the property that subtypes are behavior compatible with their parents. This formalism gives us a clear concept of behavior compatibility, that is used to justify our hierarchy structure and to compare inheritance of specifications, adopted here, with inheritance of implementations, adopted in Smalltalk.

Keywords: Formal Specifications, VDM, Object Oriented Programming, Inheritance.

Resumo

Este artigo apresenta uma especificação formal para uma hierarquia de tipos semelhante à hierarquia "Collection" apresentada pela linguagem Smalltalk. O método de especificação é uma extensão de VDM com suporte a herança de especificações, com a propriedade de subtipos terem sempre comportamento compatível com seus supertipos. Este formalismo nos dá um conceito preciso de compatibilidade de comportamento, que é usado para justificar nossa estrutura hierárquica e para comparar herança de especificações, empregada no artigo, com herança de implementações, adotada em Smalltalk.

Palavras-chave: Especificações Formais, VDM, Programação Orientada a Objetos, Herança.

A Formal Specification for a Hierarchy of Collections

Roberto Ierusalimschy
Departamento de Informática
Pontifícia Universidade Católica
Rio de Janeiro – Brazil*

July 14, 1992

Abstract

This paper presents a formal specification for a hierarchy of types similar to the Collection hierarchy presented by the Smalltalk language. The specification method is an extension of VDM that supports inheritance of specifications, with the property that subtypes are behavior compatible with their parents. This formalism gives us a clear concept of behavior compatibility, that is used to justify our hierarchy structure and to compare inheritance of specifications, adopted here, with inheritance of implementations, adopted in Smalltalk.

Keywords: Formal Specifications, VDM, Object Oriented Programming, Inheritance.

1 Introduction

Among the new concepts introduced by object-oriented languages, maybe the most important one is the concept of inheritance, and the associate notion of subclass (or subtype). Many authors consider inheritance as the distinction mark of object-oriented programming, and is common the “equation” *Objects = AbstractDataTypes + Inheritance* (e.g. [1]). On the other hand, inheritance is also a very controversial subject. There is nothing near a consensus about the meaning of inheritance, and is usually impossible to map type hierarchies written in an object-oriented language into a different language.

One of the sources of such diversity is the choice between multiple and single inheritance. In a hierarchy with multiple inheritance, a type can inherit properties from many different supertypes, while in a single hierarchy every type has at most one parent. Although single inheritance puts an artificial constrain in a type hierarchy, it is sometimes preferred because it avoids many conceptual difficulties posed by multiple inheritance, and also allows more efficient implementations.

Another reason for the diversity of inheritance concepts are the different degrees of compatibility between parents and children that a hierarchy may impose. Wegner & Zdonick [2] classifies the following compatibility levels, in order of increasing congruence:

cancel compatibility: A child has no commitment with its parents’ definitions, and is able to redefine or even eliminate any inherited operation.

name compatibility: A child must have at least all operations defined for its parents, or operations with the same names. Apart the names, there is no other commitment between those operations.

*on leave at the University of Waterloo (from Mar, 91 to Feb, 92).

signature compatibility: A child must have at least all operations defined for its parents, and each child operation must have a type compatible with the type specified by the parent.

behavior compatibility: A child must have at least all operations defined for its parents, and the operations must have compatible behaviors.

Of course, the third criterion needs a concept of type, and it only can be applied in typed languages, while the fourth criterion needs a clear definition of behavior in order to define “compatible behaviors”.

In this paper, we present a formal definition for behavior compatibility, in the framework of VDM specifications [3]. We propose an extension to the specification language that allows inheritance of specifications, and we prove that, in our system, every type is behavior compatible with its parents, in the sense that it is a valid representation for them. Using this language, we specify a hierarchy of collections, inspired by the hierarchy *Collection* of the Smalltalk language [4]. We argue that behavior compatibility can be a powerful (and rigorous) mechanism to structure types. Also we show that, with proper design, behavior compatibility can be very flexible, and even allows cancellation of inherited operations.

The rest of the paper is organized as follows. The next section presents our extensions of VDM to support inheritance of specifications, and proves a lemma that relates it with behavior compatibility. Section 3 describes the collection classes in Smalltalk, and discusses some issues related with their formalization. In section 4 we show the formal specification for the hierarchy of collections. Section 5 discusses some problems arisen by multiple inheritance. Finally, section 6 draws some conclusions.

2 Inheritance of Specifications

In order to talk about behavior compatibility, we first need a clear concept of behavior. In this work we adopt a formal specification of a type as the main description of its behavior. The specification language we use is VDM, and we follow the specification method described in [3], with some extensions to support inheritance of specifications¹. In this context, we say that a type S is behavior compatible with a type P if S is a valid representation of P .

In order to support inheritance, we adopted a more object-oriented syntax, that groups together the type declaration and all its operations. So, instead of writing in the usual way:

$$P :: f_1 : T_1$$

...

where

$$inv-P(mk-P(f_1, \dots)) \triangleq invP$$

$$O_1 (\dots)$$

$$ext \ extO_1$$

$$pre \ preO_1$$

$$post \ postO_1$$

¹The rationale for these extensions is presented in [5] and [6].

...
 we are going to use the following notation:

Specification P

$f_1: T_1$

...

$inv-P \triangleq invP$

Operation O_1 (...)

ext $extO_1$

pre $preO_1$

post $postO_1$

...

End P

The main extension to the language is a notation to declare a type as a heir of other types. We do this in a way such that a heir is always behavior compatible with its parents. Suppose we have the following specification:

Specification S

Subtype of P_1

rename $O_{rn_1}^1$ as $NO_{rn_1}^1$, $O_{rn_2}^1$ as $NO_{rn_2}^1$, ...

redefine $O_{rd_1}^1$, $O_{rd_2}^1$, ...

Subtype of P_2

rename $O_{rn_1}^2$ as $NO_{rn_1}^2$, $O_{rn_2}^2$ as $NO_{rn_2}^2$, ...

redefine $O_{rd_1}^2$, $O_{rd_2}^2$, ...

...

$f_1^s: T_1^s$

$f_2^s: T_2^s$

...

$inv-S \triangleq invS$

Operation O_1^s (...)

ext $extO_1^s$

pre $preO_1^s$

post $postO_1^s$

Operation O_2^s (...)

ext $extO_2^s$

pre $preO_2^s$

post $postO_2^s$

...

End S

The above specification declares S as a *subtype* of P_i ; P_i , in turn, is called a *supertype* of S . The following rules give the meaning of the above specification:

1. The actual fields of S are the join of all the fields from P_1, \dots , and S . If fields from different specifications have the same name and the same type (textually equal), then they are merged in one field. If fields from different specifications have the same name but different types, then there is an error condition (i.e. the meaning of the declaration is undefined).
2. The actual invariant of S is $inv-P'_1 \wedge \dots \wedge invS$. The definition of $inv-P'_i$ is as follows:

$$inv-P'_i : S \rightarrow \mathbf{B}$$

$$inv-P'_i(s) \triangleq inv-P_i(proj_{P_i}(s))$$

where:

$$proj_{P_i} : S \rightarrow P_i$$

$$proj_{P_i}(mk-S(\dots, f_1^{P_i}, \dots)) \triangleq mk-P_i(f_1^{P_i}, \dots)$$

that is, $proj_{P_i}$ is the orthogonal projection from S to P_i ; (remember that, by rule 1, S has all the fields from P_i).

Notice that we use $invS$ (without an hyphen) to denote the textual invariant, as written in the specification, while $inv-S$ denotes the final logical function that results from the above operation. The same distinction applies to pre- and post-conditions.

It is easy to verify that $inv-P_i$ and $inv-P'_i$ can be textually identical; the only difference between them is that the latter applies to S , and so it must “throw off” some fields. Notice that $invS$ can refer to the fields inherited from other specifications.

3. The operations of P_i are inserted into S in the following way: first, the operations $O_{rn_1}^i, \dots$, are renamed as $NO_{rn_1}^i, \dots$. Then, from this set, the operations $O_{rd_1}^i, \dots$ are removed. The remaining operations are included in S with unmodified external lists, pre- and post-conditions. If operations inherited from different parents have the same name and the same definition (textually equal), then they are merged. If operations from different parents have the same name but different definitions, then there is an error condition.
4. For each operation in the redefinition list (i.e., each $O_{rd_j}^i$) there must be an operation O_k^s with the same name and same parameter list; this operation *redefines* $O_{rd_j}^i$ (each declaration O_k^s can redefine more than one inherited operation). Suppose that an operation O_k^s redefines the operations O_{k_1}, O_{k_2}, \dots , inherited from the specifications P_{l_1}, P_{l_2}, \dots . Then, its actual specification is a combination of its textual specification and the specifications of all O_{k_i} , according to the following rules:

(a) The external list of O_k^s is the join of its textual external list ($extO_k^s$) with the external lists of O_{k_i} . The list $extO_k^s$ can not include fields inherited from any specification P_{l_i} . Intuitively, this is justified by the fact that an operation must have a behavior compatible with the operations it redefines. If the inherited definition asserts that some fields are not modified (by their absence in the external list), the new operation must keep this assertion.

(b) The actual pre-condition of O_k^s is

$$preO_k^s \vee pre-O'_{k_1} \vee \dots$$

Again, we have that

$$pre-O'_{k_i} : S \rightarrow \mathbf{B}$$

$$pre-O'_{k_i}(s) \triangleq pre-O_{k_i}(proj_{P_{l_i}}(s))$$

where P_i is the specification from where O_{k_i} is inherited.

We assume that when an operation is a redefinition, then the absence of an explicit pre-condition stands for false, instead of the usual true, so that the actual pre-condition simplifies to the conjunction of the inherited conditions.

(c) The post-condition of O_k^s is

$$postO_k^s \wedge (pre-O_{k_1}^s \Rightarrow post-O_{k_1}^s) \wedge \dots$$

As expected, the definition of $post-O_{k_i}^s$ is as follows:

$$post-O_{k_i}^s : S \times S \rightarrow \mathbf{B}$$

$$post-O_{k_i}^s(\overline{s}, s) \triangleq post-O_{k_i}(proj_{P_i}(\overline{s}), proj_{P_i}(s))$$

Notice that, to avoid inconsistencies between $postO_k^s$ and the inherited behavior we need the condition over external lists (rule 4a). Otherwise, the implicit requirement about unchanged variables could be contradicted by $postO_k^s$.

The other operations of S , which are not redefining any inherited operation, are left unchanged.

5. All operations in S must fulfill the *satisfiability* proof obligation. This must be checked even for the inherited operations, because the new invariant can nullify this property.

Now, we can state the important property of the above definition.

Lemma: Apart from renames, a subtype S of a type P is a valid representation for the type P , that is, it satisfies the following properties:

1. there is a retrieve function from S to P ,
2. S has all operations that P has, and
3. the operations in S model the correspondent operations in P .

Proof:

1. The obvious choice for a retrieve function from S to P is the orthogonal projection function $proj_P$. This function is total, as by rule 2 we can deduce that, for all $s \in S$, $inv-S(s) \vdash inv-P(proj_P(s))$.

A subtle point here is the *adequacy criterion*². Technically, our retrieve function can be not onto, because the invariant of S can put stronger restrictions over the fields of P . This fact corresponds to situations where a type has a level of indeterminacy that is reduced by a subtype. In a typical hierarchy this is very common, as many specifications are declared as generalizations of existing types, with a high degree of indeterminacy³. So, in our framework it is not worth pursuing adequacy. Instead, we will use the more general proof rules for operation modeling, that relies on a relation between abstraction and representation. The obvious choice for such relation is:

$$rel_P : P \times S \rightarrow \mathbf{B}$$

$$rel_P(p, s) \triangleq p = proj_P(s)$$

²This criterion states that a representation function must be onto, i.e., for each P value there must be at least one S value representing it.

³A good example is the type *Collection*, the root of the hierarchy presented in the next section.

	from $rel_P(\overline{p}, \overline{s}) \wedge pre-O_P(\overline{p}) \wedge post-O_S(\overline{s}, s)$	
1	$proj_P(s) \in P$	total- <i>proj</i> _P
2	$rel_P(proj_P(s), s)$	<i>rel</i> _P -defn(\wedge -E(h))
3	$\dots \wedge (pre-O'_P(\overline{s}) \Rightarrow post-O'_P(\overline{s}, s)) \wedge \dots$	<i>post</i> - <i>O</i> _S -defn(\wedge -E(h))
4	$\overline{p} = proj_P(\overline{s})$	<i>rel</i> _P -defn(\wedge -E(h))
5	$pre-O_P(proj_P(\overline{s}))$	\wedge -E(h),4
6	$pre-O'_P(\overline{s})$	<i>pre</i> - <i>O</i> ' _P -defn(5)
7	$post-O'_P(\overline{s}, s)$	\Rightarrow -E(\wedge -E(3),6)
8	$post-O_P(proj_P(\overline{s}), proj_P(s))$	<i>post</i> - <i>O</i> ' _P -defn(7)
9	$post-O_P(\overline{p}, proj_P(s)) \wedge rel_P(proj_P(s), s)$	\wedge -I(8,2),4
	infer $\exists p \in P \cdot post-O_P(\overline{p}, p) \wedge rel_P(p, s)$	\exists -I(1,9)

Figure 1: proof of the result rule

2. The operations of P that do not appear in the redefinition list are straightly included in S (rule 3); the operations in the redefinition list are explicitly declared in S (rule 4).
3. The operations of P that do not appear in the redefinition list are textually identical in S and P , and so are valid models. Let us see the redefined operations. Suppose the operation O_S redefines an operation O_P . The general proof obligations for operation modeling are:

- $\forall s \in S, p \in P \cdot rel_P(p, s) \wedge pre-O_P(p) \Rightarrow pre-O_S(s)$
- $\forall \overline{s}, s \in S, \overline{p} \in P \cdot$
 $rel_P(\overline{p}, \overline{s}) \wedge pre-O_P(\overline{p}) \wedge post-O_S(\overline{s}, s) \Rightarrow \exists p \in P \cdot post-O_P(\overline{p}, p) \wedge rel_P(p, s)$

The first proof obligation is trivial, if we remember that the retrieve function is the orthogonal projection and that $pre-O_S$ has the form $\dots \vee pre-O'_P \vee \dots$ (rule 4). A proof of the second one is presented in figure 1.

This lemma assures that a subtype, in our specification language, is always behavior compatible with its supertypes. More specifically, we have a “is-a” relationship between them. Every object of a given type (that is, that satisfies a given specification) also belongs to all supertypes of that type.

It is important to note the difference between a subtype and an implementation. In the definition of subtypes we said nothing about initial states. The lemma tells us about the behavior of existing objects, but not about the creation of objects. In the specification of a subtype, one can stretch the invariant, narrowing the range of initial states. For instance, one can define a subtype *Square* of a type *Rectangle*, stating in the invariant the equality of width and height. On the other hand, it is obvious that *Square* is not a good implementation for generic rectangles.

However, it is not possible to build subtypes in unrestricted ways: the ultimate restriction for the range of admissible subtypes is posed by rule 5. It is not possible to stretch a post-condition more than allowed by the satisfiability criterion. Moreover, one can not put a new invariant that conflicts with the inherited post-conditions. So, the limits to modify a type are given not only by its model, but also by its operations. As an example, we can return to the type *Rectangle*. If the

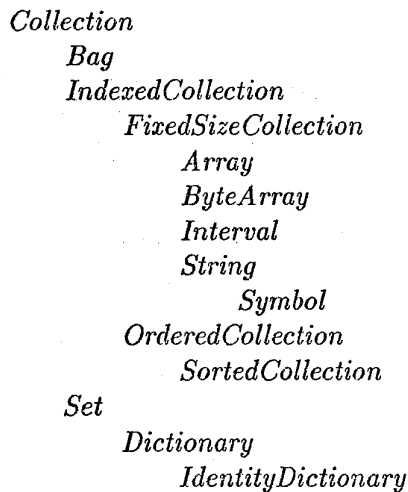


Figure 2: The Smalltalk hierarchy of collections

type is immutable, then a subtype *Square* only restricts the range of initial states, and is correct. However, if *Rectangle* has an operation like *SET-WIDTH*, that changes the width for an arbitrary value and does not modify the height, than this operation would be unsatisfiable under the *Square* invariant. In that case, *Square* would not be a valid subtype of *Rectangle*.

3 Smalltalk Hierarchy \times Formal Hierarchy

The Smalltalk library is organized as a hierarchy of classes, rooted in the class *Object*. All classes representing collections of objects, like sets and bags, are united in a subtree under the class *Collection*.

Before going on, it is important to notice that Smalltalk implements what we call *inheritance of implementations*. This means that the inheritance mechanism has no relationship with a concept of subtype or with the polymorphism of the language; instead, it is solely a mechanism for code reuse. As a consequence, Smalltalk presents only name compatibility between subclasses; moreover, as an operation can be redefined to rise an error if called, the language actually allows cancel compatibility. Implementation aspects, together with the restriction of single inheritance, guides the organization of the Smalltalk library hierarchy. The position of a class inside the hierarchy is mainly dictated by where it can reuse more code.

On the other hand, the formal hierarchy we are going to present is based on *inheritance of specifications*, that is, behavior compatibility. Moreover, our formal system allows multiple inheritance. So, it is expected that we have to modify the original hierarchy to fit the formal specification. Nevertheless, the final hierarchy encompasses all interesting classes under a root *Collection*, and it is rather similar to the Smalltalk structure.

The primary members of the *Collection* class hierarchy in Smalltalk⁴ are shown in figure 2. The *Collection* class defines the basic functionality for all collections. *Bag*, *Set*, and *Array* have the expected meanings. *IndexedCollection* introduces operations for indexing elements by their positions inside the collection, and so introduces an order among the elements. *FixedSizeCollection* comprises the collections that have their size specified at creation time, while an *OrderedCollection*

⁴We are following the library presented in [7].

allows insertion of elements at both ends, growing as needed. The classes *ByteArray*, *String*, and *Symbol* are all special implementations for arrays, optimized to handle integers and characters; from a specification point of view they are similar to *Array*, and so we will not consider them here. Class *Interval* implements immutable sequence of integers in arithmetic progression. *Dictionary* corresponds to generic symbol tables, while *IdentityDictionary* differs from its parent only in the way it does key comparison, and will not be considered here, either.

In order to match the *Collection* tree with a formal hierarchy, and also to make use of multiple inheritance, we made the following modifications. First, as already explained, we removed the classes *ByteArray*, *String*, *Symbol*, and *IdentityDictionary*. On the other hand, we introduced the types *Stack*, *Queue* and *DoubleQueue*, mainly because we think no paper about abstract data types is complete without a stack⁵.

The operation to remove a generic element poses some difficulties. Not all collections allow the deletion of a generic element: for instance, in a stack one can only remove the element at the top. We solved this point with multiple inheritance. We declared the remove operation in a subtype of *Collection*, called *RemovableCollection*. Any type that needs the operation can inherit it, while keeping all other parents.

The other changes were all in the *IndexedCollection* subtree. From a behavior perspective, many ordered collections are not compatible with an indexing specification. For instance, a sorted collection does not have an operation to put an arbitrary element in an fixed position — Smalltalk uses cancel compatibility to achieve that. Stacks and queues do not allow access to arbitrary elements, either. So, we declared *SortedCollection* as a sibling of *IndexedCollection*, and created a new type above them. As this new type encompasses all kinds of collections with an internal order, we took the name *OrderedCollection* for it, renaming the old *OrderedCollection* as *List* in our hierarchy. Finally, we took advantage of multiple inheritance to put *FixedSizeCollection* as a direct child of *Collection*, making it a sibling of *OrderedCollection*. Then, an array can be declared as an indexed collection with a fixed size, inheriting both behaviors.

4 The Specification of The Collection Hierarchy

In this section we present a formal specification for a hierarchy of collections, using the method presented in section 2.

The root of the hierarchy is the type *Collection* (figure 3). That specification uses two auxiliary functions:

$$\begin{aligned}
 mpc &: (T \xrightarrow{m} N_1) \times T \rightarrow N \\
 mpc(m, e) &\triangleq \text{if } e \in \text{dom } m \text{ then } m(e) \text{ else } 0 \\
 \\
 bag\text{-}size &: T \xrightarrow{m} N_1 \rightarrow N \\
 bag\text{-}size(m) &\triangleq \sum_{i \in \text{dom } m} m(i)
 \end{aligned}$$

Both operate over bags, implemented as a map that gives, for each element in the bag, its number of occurrences. *mpc* gives the number of occurrences of a given element inside a bag, and *bag-size* returns the total number of elements in a bag.

The basic structure of a collection is a bag, stored in the field *contents*. Other options would be a set or a sequence. A set would not allow us to define a generic *INSERT* operation, as one would

⁵It is amazing that the Smalltalk library does not include a class *Stack*. Maybe more amazing is the fact that it does not need stacks.

Specification *Collection*

contents: $T \xrightarrow{m} \mathbf{N}_1$
maxSize: \mathbf{N}

Operation *INSERT* ($e: T$)

ext wr *contents*: $T \xrightarrow{m} \mathbf{N}_1$
rd *maxSize*: \mathbf{N}

pre $\text{bag-size}(\text{contents}) < \text{maxSize}$

post $e \in \text{dom } \text{contents} \wedge \text{mpc}(\text{contents}, e) \geq \text{mpc}(\overline{\text{contents}}, e) \wedge$
 $(\forall x \in T \cdot x \neq e \Rightarrow (\text{mpc}(\text{contents}, x) = \text{mpc}(\overline{\text{contents}}, x)))$

Operation *CONTAINS* ($e: T$) $b: \mathbf{B}$

ext rd *contents*: $T \xrightarrow{m} \mathbf{N}_1$
post $b = (e \in \text{dom } \text{contents})$

Operation *IS_EMPTY* () $b: \mathbf{B}$

ext rd *contents*: $T \xrightarrow{m} \mathbf{N}_1$
post $b = (\text{bag-size}(\text{contents}) = 0)$

Operation *GET_ELEMENT* () $e: T$

ext rd *contents*: $T \xrightarrow{m} \mathbf{N}_1$
pre $\text{bag-size}(\text{contents}) > 0$
post $e \in \text{dom } \text{contents}$

Operation *SIZE* () $s: \mathbf{N}$

ext rd *contents*: $T \xrightarrow{m} \mathbf{N}_1$
post $s = \text{bag-size}(\text{contents})$

End *Collection*

Figure 3: Collection

```

Specification RemovableCollection
  Subtype of Collection

  Operation REMOVE ( $e: T$ )
  ext wr contents:  $T \xrightarrow{m} \mathbf{N}_1$ 
  pre  $\text{mpc}(\text{contents}, e) > 0$ 
  post  $(\forall x \in T \cdot x \neq e \Rightarrow (\text{mpc}(\text{contents}, x) = \overline{\text{mpc}(\text{contents}, x)}) \wedge$ 
       $\text{mpc}(\text{contents}, e) = \overline{\text{mpc}(\text{contents}, e)} - 1$ )
End RemovableCollection

```

Figure 4: *RemovableCollection*

be able to deduce that the size of a collection does not change after the insertion of an already present element, an obviously false statement for generic collections. A sequence would be too biased for non-ordered collections, like sets. The other field of the type *Collection* is its maximum size.

Notice that, although we define a maximum size, we do not set an invariant to assure it. The check is made only in the pre-condition of the insertion operation. In this way, we can define unbounded collections just weakening this pre-condition to true. We do not do that in this hierarchy, because ultimately all real collections are bounded, and any program that uses a collection must be able to treat an overflow in a sensible way.

From the operations of *Collection*, the only one a little more complex is *INSERT*. We must make sure that its specification is flexible enough to allow redefinitions for all kinds of collections, but at the same time captures all the commonality of insert operations. The pre-condition only checks the available space. The post-condition has three terms: the first one asserts the presence of the inserted element in the final collection; the second asserts that the number of such elements does not decrease; and the third part assures that the operation does not affect other elements.

Not all collections support a generic remove operation, that is, one that can remove an arbitrary element. So we define this operation in a separate type, called *RemovableCollection* (figure 4). *RemovableCollection* inherits all operations from the basic collection, and adds the operation *REMOVE*. The pre-condition asserts the presence of the element to be removed, and the post-condition states that the number of such elements is decremented by 1 after the operation.

Both *Collection* and *RemovableCollection* are examples of what is called an *abstract class* in the Smalltalk terminology. They are not intended to have implementations, but rather to be used as supertypes for other types. Using them we can define the first two *concrete classes* of our hierarchy: *Set* and *Bag* (figure 5). Both are subtypes of *RemovableCollection*, as they allow the deletion of any contained element. For the type *Set*, the only modification is a stronger invariant, disallowing repetitions. For bags, the invariant is correct, but we must redefine the operation *INSERT* to make sure it always inserts one and only one element. Notice that the redefinition only specifies this new property, that is joined with the inherited properties according to the rules of section 2.

Our next specification is another abstract class, that encompasses all collections with an internal order for its elements (figure 6). To keep track of this order, we use a sequence, called *order*; the invariant assures the sequence and the bag (*contents*) have the same elements. The auxiliary function *seq-count* counts the number of occurrences of a given element inside a sequence.

Specification Set
 Subtype of *RemovableCollection*
 $inv\text{-Set} \triangleq \forall x \in \text{dom } contents \cdot contents(x) = 1$
 End Set

Specification Bag
 Subtype of *RemovableCollection*
 redefine *INSERT*
 Operation *INSERT* ($e: T$)
 post $mpc(contents, e) = \overline{mpc(contents, e)} + 1$
 End Bag

Figure 5: Sets and Bags

Specification OrderedCollection
 Subtype of *Collection*
 rename *GET_ELEMENT* as *GET_FIRST*
 redefine *GET_FIRST*, *INSERT*
 $order: T^*$
 $inv\text{-OrderedCollection} \triangleq \forall x \in T \cdot seq\text{-count}(order, x) = mpc(contents, x)$
 Operation *GET_FIRST* () $e: T$
 ext rd $order: T^*$
 post $e = \text{hd } order$
 Operation *INSERT* ($e: T$)
 ext wr $order: T^*$
 pre false
 post true
 End *OrderedCollection*

Figure 6: OrderedCollection

Specification *Stack*

Subtype of *OrderedCollection*

rename *GET-FIRST* as *TOP*, *INSERT* as *PUSH*

redefine *PUSH*

Operation *PUSH* ($e: T$)

post $order = [e] \frown \overline{order}$

Operation *POP* () $e: T$

ext wr $order: T^*$

wr $contents: T \xrightarrow{m} N_1$

pre $len\ order > 0$

post $order = \overline{tl\ order} \wedge e = hd\ \overline{order}$

End *Stack*

Figure 7: Stack

$seq\text{-}count : T^* \times T \rightarrow N$

$seq\text{-}count(s, x) \triangleq$ if $s = []$ then 0 else
if $x = hd\ s$ then $seq\text{-}count(tl\ s, x) + 1$ else $seq\text{-}count(tl\ s, x)$

There is a little trick in the redefinition of *INSERT*. To keep the invariant, *INSERT* must modify the field *order* according to *contents*. But where to put the new element depends on the kind of collection. A stack may need to insert the new element at the beginning of the sequence, while a queue would insert at the end. So, the redefinition of *INSERT* declares *order* in its external list to allow the maintenance of the invariant, but does not strengthen the post-condition; this is left for the subtypes. The redefinition of *GET-ELEMENT* to *GET-FIRST* is done because the latter seems to be more useful (but not available in unordered collections).

The simplest subtypes of *OrderedCollection* are *Stack* (figure 7) and *Queue* (figure 8). Each one redefines *INSERT* in order to specify the position of the inserted element, at the beginning (Stacks) or at the end (Queues) of the sequence. Also, each one specifies a particular operation to remove elements (*POP* and *REMOVE*). As they do not allow the deletion of a generic element, they can not inherit this operation from *RemovableCollection*.

Using stacks and queues, it is trivial to declare a double queue (figure 9). We only need to add one new operation, to remove elements from the end. Notice the use of the redefinition and rename facilities to join two inherited operations (*POP* and *REMOVE*) into one (*REMOVE-FIRST*).

Another useful subtype of *OrderedCollection* is *SortedCollection* (figure 10). For this type we assume that the basic type *T* has a total order, given by the infix function " \leq_T ". The specification joins the basic properties from *OrderedCollection* with other properties from *Bag*, adding an invariant stating that the elements inside *order* must be sorted. From *Bag* comes the restriction that *INSERT* adds only one element to the collection. This property and the invariant are enough to completely specify the insertion; the redefinition is used again only to join both definitions (in this case we have not needed to rename them, as they already had the same name). *SortedCollection* also has a *REMOVE* operation, inherited from *Bag*. The redeclaration allows the operation to modify the field *order*. Again, the invariant is enough to ensure that this field is modified in the

```

Specification Queue
  Subtype of OrderedCollection
    redefine INSERT
  Operation INSERT (e: T)
    post order =  $\overleftarrow{\text{order}} \frown [e]$ 
  Operation REMOVE-FIRST () e: T
    ext wr order: T*
      wr contents: T  $\xrightarrow{m}$   $\mathbb{N}_1$ 
    pre len order > 0
    post order =  $\text{tl } \overleftarrow{\text{order}} \wedge e = \text{hd } \overleftarrow{\text{order}}$ 
End Queue

```

Figure 8: Queue

```

Specification DoubleQueue
  Subtype of Queue
    rename INSERT as INSERT-LAST
    redefine REMOVE-FIRST
  Subtype of Stack
    rename PUSH as INSERT-FIRST, POP as REMOVE-FIRST
    redefine REMOVE-FIRST
  Operation REMOVE-FIRST () e: T
    pre false
    post true
  Operation REMOVE-LAST () e: T
    ext wr contents: T  $\xrightarrow{m}$   $\mathbb{N}_1$ 
      wr order: T*
    pre len order > 0
    post  $\overleftarrow{\text{order}} = \text{order} \frown [e]$ 
End DoubleQueue

```

Figure 9: DoubleQueue


```

Specification SortedCollection
  Subtype of OrderedCollection
    redefine INSERT
  Subtype of Bag
    redefine INSERT, REMOVE
  inv-SortedCollection  $\triangleq \forall i, j \in \text{inds } \text{order} \cdot i < j \Rightarrow c(i) \leq_T c(j)$ 
  Operation INSERT (e: T)
    pre false
    post true
  Operation REMOVE (e: T)
    ext wr order: T*
    pre false
    post true
End SortedCollection

```

Figure 10: SortedCollection

proper way.

Indexed collections are specified at figure 11. The main operations are *AT*, to access an element by its index (i.e., its position within the sequence), and *PUT*, to modify the element at a given position. Both operations check at the pre-condition that the index is inside the bounds of the sequence. The post-condition of *PUT* must assure that only the specified position is modified, without disturbing the other elements in the sequence.

A generic list structure can be defined only combining predefined behaviors — see figure 12. Such structure supports insertion and deletion at both ends (from type *DoubleQueue*), general indexing facilities (from *IndexedCollection*), and a remove operation (from *RemovableCollection*).

In order to specify a common array, we need a trick. At a first glance, arrays could not be a subtype of collection, because there is no way to insert an element into an array, but only to change elements: an array does not have “empty slots”. However, we can disallow the *INSERT* operation and still keep behavior compatibility with collections. This is done in the specification *FixedSizeCollection* (figure 13), simply stating in the invariant that such collections are always full, and so the pre-condition for an insertion can never be satisfied! Now, an array is just an indexed collection with a fixed size, and its specification is shown in figure 14. Notice that a type invariant must be true at the initial state of an object, but our specification does not define how to achieve this (and does not need to). A typical implementation may fill new arrays with an invalid or neutral initial value.

The specification for intervals is shown in figure 15. As a subtype of *FixedSizeCollection*, it has an invalid *INSERT* operation. The invariant states that the sequence must be an arithmetic progression. The type *T* must be **N**, **Z** or **R**, but our language does not have a notation for this restriction. This specification describes a very unusual kind of collection, and it is defined here only because it is present in the Smalltalk hierarchy.

Our last specification is a dictionary, or a symbol-table (figure 16). Following the Smalltalk hierarchy, we declared it as a subtype of *Set*. The key to keep a compatible behavior is to identify the set with the keys in the dictionary. We rename most operations to reflect this concept. The

Specification *IndexedCollection*

Subtype of *OrderedCollection*

Operation *AT* ($i: \mathbb{N}$) $e: T$

ext rd *order*: T^*

pre $i \in \text{inds } \textit{order}$

post $e = \textit{order}(i)$

Operation *PUT* ($i: \mathbb{N}, e: T$)

ext wr *order*: T^*

wr *contents*: $T \xrightarrow{m} \mathbb{N}_1$

pre $i \in \text{inds } \textit{order}$

post $\text{len } \textit{order} = \overline{\text{len } \textit{order}} \wedge \textit{order}(i) = e \wedge \forall j \in \text{inds } \textit{order} \cdot j \neq i \Rightarrow \textit{order}(j) = \overline{\textit{order}(j)}$

End *IndexedCollection*

Figure 11: IndexedCollection

Specification *List*

Subtype of *DoubleQueue*

Subtype of *IndexedCollection*

Subtype of *RemovableCollection*

End *List*

Figure 12: List

Specification *FixedSizeCollection*

Subtype of *Collection*

inv-FixedSizeCollection $\triangleq \text{bag-size}(\textit{contents}) = \textit{maxSize}$

End *FixedSizeCollection*

Figure 13: FixedSizeCollection

Specification *Array*
 Subtype of *FixedSizeCollection*
 Subtype of *IndexedCollection*
End *Array*

Figure 14: Array

Specification *Interval*
 Subtype of *FixedSizeCollection*
 Subtype of *OrderedCollection*
 $inv\text{-}Interval \triangleq \exists a, r \in \mathbf{R} \cdot \forall i \in \text{inds } order \cdot order(i) = i * r + a$
Operation *AT* ($i: \mathbf{N}$) $e: T$
 ext rd *order*: T^*
 pre $i \in \text{inds } order$
 post $e = order(i)$
End *Interval*

Figure 15: Interval

Specification Dictionary

Subtype of Set

rename *INSERT* as *INSERT-KEY*, *CONTAINS* as *CONTAINS-KEY*

redefine *INSERT-KEY*, *REMOVE*

values: $T \xrightarrow{m} V$

inv-Dictionary \triangleq $\text{dom } \textit{values} = \text{dom } \textit{contents}$

Operation *INSERT-KEY* ($e: T$)

ext wr *values*: $T \xrightarrow{m} V$

pre false

post true

Operation *REMOVE* ($e: T$)

ext wr *values*: $T \xrightarrow{m} V$

pre false

post true

Operation *SET-VALUE* ($e: T, v: V$)

ext wr *contents*: $T \xrightarrow{m} \mathbb{N}_1$

wr *values*: $T \xrightarrow{m} V$

rd *maxSize*: \mathbb{N}

pre $\text{bag-size}(\textit{contents}) < \textit{maxSize} \vee e \in \text{dom } \textit{contents}$

post $\textit{values} = \overline{\textit{values}} \uparrow \{e \mapsto v\}$

Operation *GET-VALUE* ($e: T$) $v: V$

ext rd *values*: $T \xrightarrow{m} V$

pre $e \in \text{dom } \textit{values}$

post $v = \textit{values}(e)$

End Dictionary

Figure 16: Dictionary

new field, *values*, tracks the relationship between keys and values, and the invariant assures our identification between sets and keys. Notice that we could get a much simpler specification using only the field *value*, without any inheritance, but the point here is to show that a dictionary can be (formally) defined as a subtype of sets, according to the Smalltalk hierarchy.

5 Some Difficulties with Multiple Inheritance

Whenever a system allows multiple inheritance, it arises the possibility of different kinds of conflicts among the heritage. Object Oriented programming languages have introduced many mechanisms to cope with these conflicts, ranging from single inheritance (e.g. Smalltalk [4]) to some scheme of priorities (e.g. Common Lisp [8]) to facilities for renaming and redefining (Eiffel [9]). Obviously, single inheritance is not a solution to multiple inheritance. The scheme of priorities adopted by some languages can lead to unexpected results in more intricate hierarchies. Moreover, although some systems allow some form of control over the solution (for instance, by means of changing the order of the subtype declarations), such control is very restricted. So we have adopted here a solution based on explicit statements to avoid conflicts, the *renames* and *redefines* declarations. Nevertheless, it is important to keep in mind the drawbacks of those solutions.

Redefinition is a very powerful mechanism to enhance the flexibility of inheritance, even in a system with single inheritance. In our system, the semantics of redefinitions is given in such a way that subtypes are always consistent with their parents. But we still can have some problems. Suppose we try to define a new specification, which is a subtype of both *Stack* and *Queue* (pages 12 and 13). This will lead us to an unfeasible specification, as the post-conditions from both *INSERT* operations are incompatible. The strange thing is that these operations were originally the same, the *INSERT* operation from type *Collection*, that being inherited through two different paths became incompatible.

The mechanism for renaming has more problems. The main use for renaming is to avoid name clashes between operations inherited from different parents. This is a very important facility in order to be able to join specifications written by different groups. However, the facility is not compatible with the notion of subtyping (remember that in the lemma of section 2 we did not consider renamings). A partial solution, adopted by the programming language Eiffel, is to keep track of the renames when viewing a type as a supertype. For example, if a variable of type *Collection* contains a *Stack*, when we refer to the operation *INSERT* the system automatically translate it to a reference to the operation *PUSH*. However, such solution does not work sometimes. Referring again to the above example, if the variable contains a *DoubleQueue*, the same reference to the operation *INSERT* can be translated to *INSERT-FIRST* or *INSERT-LAST*, giving different results.

6 Conclusions

We have presented a formal definition for behavior compatibility. We have extended the VDM notation to support inheritance of specifications, in a way that an heir is always behavior compatible with its ancestors. Based on this extension, we have built a hierarchy of collections, resembling the tree of collection classes in the Smalltalk library. Our hierarchy uses multiple inheritance, and follows a strict “is-a” discipline: an implementation always satisfies the formal specifications of all supertypes of its basic type.

We argue that behavior compatibility is not too restrict, mainly when a language supports separate hierarchies for types and implementations (like Duo-Talk [10] and O=M [11]). In the implementation tree, parents can be chosen according to a criterion of code reuse. On the other

hand, the specification hierarchy must be based strictly in behavior considerations; this gives an objective criterion to organize the hierarchy and permits the use of formal methods for reasoning about programs with inheritance.

A key point in our method is the use of a specification language that supports indeterminacy, like VDM. This feature allows definitions of very general types in the first levels of a hierarchy; these types specify only what is known or needed at that level. Then subtypes can be defined narrowing the range of indeterminacy of the inherited behavior, while adding new behaviors (in the form of new fields and new operations). As we have shown in the collection example, this method presents great flexibility, and in some cases a proper design can even cancel inherited operations. Moreover, the method allows reuse of specifications, as shown in the definitions of *Array* and *List* — both types only combine inherited specifications to define their behavior.

A point we have neglected in our work is iteration. Collections in Smalltalk offer a variety of methods to iterate over their elements. However, all those methods are based on *block* parameters, i.e., the iterator method receives as a parameter the operation to be performed over the elements of the collection. Our specification language does not support this kind of second order functions, and so we have not included those operations in the formal specifications. In [5] we propose the use of iterator objects⁶ to cope with iterations without needing second order facilities, and we also show how to include such mechanisms in a formal specification for ordered collections.

References

- [1] S. Danforth and C. Tomlinson. Type theories and object-oriented programming. *ACM Computing Surveys*, 20(1):29–71, 1988.
- [2] P. Wegner and S. Zdonik. Inheritance as an incremental modification mechanism or what like is and isn't like. In *ECOOP'88 Proceedings*, pages 55–77, 1988.
- [3] Cliff B. Jones. *Systematic Software Development using VDM*. International Series in Computer Science. Prentice Hall, second edition, 1990.
- [4] Adele Goldberg and Dave Robson. *Smalltalk-80 : The Language and its Implementation*. Addison-Wesley, 1983.
- [5] Roberto Ierusalimschy. *O=M : Uma Linguagem Orientada a Objetos para Desenvolvimento Rigoroso de Programas*. PhD thesis, Dep. Informática, PUC-Rio, Rio de Janeiro, Brazil, 1990.
- [6] Roberto Ierusalimschy. A method for object-oriented specifications with VDM. *Monografias em Ciência da Computação 2/91*, PUC-Rio, Rio de Janeiro, Brazil, 1991.
- [7] Digitalk, Inc., Los Angeles – CA. *Smalltalk/V Windows – Tutorial and Programming Handbook*, 1991.
- [8] D. Bobrow, L. DeMichiel, R. Gabriel, S. Keene, G. Kiczales, and D. Moon. *Common Lisp Object System Specification*. ANSI Common Lisp, 1988. Doc. 88-003, X3J13 Standards Committee.
- [9] Bertrand Meyer. Eiffel – a language and environment for software engineering. *The Journal of Systems and Software*, 8(3):129–46, 1988.

⁶This concept is somewhat similar to streams in Smalltalk.

- [10] C. Lunau. Separation of hierarchies in Duo-Talk. *Journal of Object-Oriented Programming*, 2(2):20–26, 1989.
- [11] Roberto Ierusalimschy. The O=M programming language. Monografias em Ciência da Computação 3/91, PUC-Rio, Rio de Janeiro, Brazil, 1991.