BETA Language Introduction

Mjølner Informatics Report
MIA 94−26
March 2004

Table of Contents

Table of Contents

Table of Contents

1 Introduction

This report is a an introduction to the BETA language. The BETA language is presented to someone who is familiar with one or more object−oriented language such as C++ or Eiffel.

The overall aspects of the BETA language is presented. The presentation focuses on the concepts and ideas behind the design of BETA, and includes examples on the use of most constructs. The tutorial contains sections on basic constructs, patterns and objects, singular objects, subprocedure, control patterns, nested patterns, virtual patterns, coroutines, concurrency, and inheritance.

For more details about the BETA language than presented in this tutorial please see [MMN 93]. For a tutorial on the Mjølner System, please see [MIA 94−24]

This introduction is based on a chapter in the book [Knudsen 94] written by Ole Lehrmann Madsen.

1.1 Language Concepts

BETA is a modern object−oriented language from the Scandinavian school of object−orientation where the first object−oriented language Simula was developed. BETA supports the object−oriented perspective on programming and contains comprehensive facilities for procedural and functional programming. BETA has powerful abstraction mechanisms for supporting identification of objects, classification and composition. BETA is a strongly typed language like Simula, Eiffel and C++ with most type checking being carried out at compile−time. It is well known that it is not possible to obtain all type checking at compile time without sacrificing the expressiveness of the language. BETA has an optimum balance between compile−time type checking and run−time type checking.

1.2 Powerful Abstraction Mechanisms

BETA has powerful abstraction mechanisms that provide excellent support for design and implementation, including data definition for persistent data. The powerful abstraction mechanisms greatly enhance reusability of designs and implementations.

1.2.1 The pattern

The abstraction mechanisms include class, procedure, function, coroutine, process, exception and many more, all unified into the ultimate abstraction mechanism: the pattern. In addition to the pattern, BETA has subpattern, virtual pattern and pattern variable. This unification gives a uniform treatment of abstraction mechanisms and a number of new ones. Most object−oriented languages have classes, subclasses and virtual procedures, and some have procedure variables. Since a pattern is a generalization of abstraction mechanisms like class, procedure, function, etc., the notions of subpattern, virtual pattern and pattern variable also apply to these abstraction mechanisms. In addition to the above mentioned abstraction mechanisms, the pattern subsumes notions such as generic package and task type as known from Ada.

1.2.2 Subpattern

The subpattern covers subclasses as in most other object−oriented languages. In addition, procedures may be organized in a subprocedure hierarchy in the same way as classes may be organized in a subclass hierarchy. Since patterns may also be used to describe functions, coroutines, concurrent processes, and exceptions, these may also be organized in a pattern hierarchy.

1.2.3 Virtual pattern

The notion of virtual pattern covers virtual procedures as in C++. In addition, virtual patterns cover virtual classes, virtual coroutines, virtual concurrent processes, and virtual exceptions. Virtual classes provide a more general alternative to generic classes as in Eiffel or templates as in C++.

1.3 Pattern variable

BETA includes the notion of pattern variable. This implies that patterns are first class values, that may be passed around as parameters to other patterns. By using pattern variables instead of virtual patterns, it is possible dynamically to change the behavior of an object after its generation. Pattern variables cover procedure variables (i.e. a variable that may be assigned different procedures). Since patterns may be used as classes, it is also possible to have variables that can be assigned classes, etc.

1.4 Coroutines and concurrency

BETA does not only allow for passive objects as in C++ and Eiffel. BETA objects may also act as coroutines, making it possible to model alternating sequential processes and quasi−parallel processes. BETA coroutines may be executed concurrent (non pre−emptive scheduling in current implementation). The basic mechanism for synchronization is semaphores, but high−level abstractions for synchronization and communication, hiding all details about semaphores, are easy to implement, and the standard library includes monitors, and rendezvous. The user may easily define new concurrency abstractions including schedulers for processes.

BETA supports the three main subfunctions of abstraction: identification, classification, and composition as described in the following.

1.5 Identification of Objects

1.5.1 Class−less objects

It is possible to describe objects that are not generated as instances of a class pattern, so−called 'class−less objects'. This is in many cases useful when there is only one object of a kind. In most object−oriented languages, it is necessary to define superfluous classes for such objects. In analysis and design, it is absolutely necessary to be able to describe singular objects without having them as instances of classes.

1.6 Classification

Classification is supported by patterns, subpatterns, and virtual patterns that make it possible to describe classification hierarchies of objects and patterns (objects, classes, procedures, functions, coroutines, processes, exceptions, etc.).

1.7 Composition (Aggregation)

Objects and patterns may be defined as a composition of other objects and patterns. The support for composition includes:

- Whole−part composition: an attribute of an object may be a part−object. This makes it possible to describe objects in terms of their physical parts.
- Reference composition: an attribute may be a reference to an object. Reference composition is the basis for modeling arbitrary relations between objects.
- Localization: an attribute of an object may be a (nested) patternÑalso known as block−structure. The block−structure makes it easy to create arbitrary nested patterns. This makes it possible for objects to have local patterns used as classes, procedures, etc. Local patterns greatly enhance the modeling capabilities of an object−oriented language.

1.8 Inheritance

In BETA, inheritance is not only restricted to inheritance from superpatterns. It is also possible to inherit from a part−object. Virtual patterns in the part−object may be redefined to influence the enclosing object. Multiple inheritance is supported through inheritance from multiple part−objects. This gives a much cleaner structure than inheritance from multiple superpatterns.

1.9 Conceptual Framework

BETA is intended for modeling and design as well as implementation. During the design of BETA the development of the underlying conceptual framework has been just as important as the language itself.

1.9.1 Modeling

BETA is a language for representing/modeling concepts and phenomena from the application domain and for implementing such concepts and phenomena on a computer system. Part of a BETA program describes objects and patterns that represent phenomena and concepts from the application model. This part is said to be representative since BETA elements at this level are meaningful with respect to the application domain. Other parts of a BETA program are non−representative, since they do not correspond to elements of the application domain, but are intended for realizing the model as a computer system.

The BETA language as presented in this introduction is for describing objects and patterns. The objects and patterns constitute the logical structure of a program execution. The physical structure of a program execution is handled by other components of the Mjølner System. A tutorial on using the this system is given in [MIA 94−24].

2 Basic Constructs

The most fundamental elements of BETA are objects and patterns. This section describes the basic patterns and values, simple assignments, control structures, variable declarations, repetitions and patterns used as composite types.

2.1 Simple Types and Values

The simple types (or also called basic patterns) are integer, boolean, char, and, real. The following table shows the simple types with examples of values, including text constant. Notice, that text is not a simple type in BETA, but a pattern defined in the basic BETA environment called `betaenv`.

Type Value integer7, −4, 0x4FFC, 2x101101 booleantrue, false char'c' real3.141, −1.234E3 text constant'abc'

2.2 Simple Static Variables

In BETA, a static variable (also called a static reference) is declared like:

```
i: @integer;
r: @real;
```

Variables of the simple types can only be declared static, see below for dynamic references.

2.3 Simple Assignments

2.3.1 Value assignment

Simple value assignments in BETA goes left to right:

```
2 -> i          (* assign the value 2 to i *)
i -> j          (* assign the value of i to j *)
i*j -> k        (* assign the value of i*j to k *)
(i,j) -> (x,y) (* assign the value of i to x and
                 *        the value of j to y
                 *)
```

2.4 Control Structures

BETA has two build−in control structures: if and for, both having two forms. The simple if imperative with one boolean expression:

2.4.1 if

```
(if <expression> then
    <imperatives>
```

```
 else
    <imperatives>
if)
```

and the if with several alternatives:

```
(if <expression>
 // <expression> then <imperatives>
 // <expression> then <imperatives>
 else
   <imperatives>
if)
```

where // means equals.

The simple for imperative just iterates a given number of times:

2.4.2 for

```
(for <expression> repeat <imperatives> for)
```

but the for imperative may implicitly declare an iteration variable, only available inside the for loop, by:

```
(for <variable>: <expression> repeat <imperatives> for)
```

The for loop always starts in 1 and stops at <expression>. The loop can be terminated or restarted using labels, see below.

The following BETA code is a general object−descriptor (or descriptor for short):

2.4.3 descriptor

```
<declarations>
enter <enter−list>
do <imperatives>
exit <exit−list>
```

A descriptor consists of type and variable declarations, an enter part for parameters (enter <enter−list>), a do−part for the action (do <imperatives>), and finally an exit part for the results (exit <exit−list>). All elements are optional.

A descriptor can be labeled, and the descriptor can be restarted and/or left using the label:

2.4.4 labeled descriptor

```
L: (#  leave L; restart L #)
```

In general any imperative can have a label:

2.4.5 labeled imperative

```
L: <imperative>
L: (if   leave L   if)
L: (for  leave L  for)
```

leave L implies that control is transferred to immediately after the labeled imperative/descriptor.
restart L implies that control is transferred to immediately before the labeled imperative/descriptor.

2.5 Static and Dynamic Variables

2.5.1 Reference attributes

In BETA variables are two examples of reference attributes − static references that constantly denote the same object, and dynamic references that may denote different objects.

2.5.2 Static Reference

Examples of static reference variables are:

```
i: @integer (* i refers to a simple type: integer *)
p: @A       (* an instance of A is automatically generated and
             * p always refers to this object *)
s: @(#  #) (* an instance of (#  #) is automatically generated
             * and s always refers to this singular object *)
```

2.5.3 Dynamic reference

Examples of dynamic reference variables are:

```
i: ^integerObject
p: ^A
```

Assignments between dynamic references can be done using the reference operator '[]' (read box):

2.5.4 Reference assignment

```
p1[] -> p2[]  (* reference assignment *)
```

Dynamic reference variables are initially NONE i.e. refers to nothing. Objects can be created using the new operator '&':

```
&A[] -> p[]  (* create an instance of A and assign the reference
              * to p *)
```

It is illegal to declare dynamic references to simple types:

```
i: ^integer (* ILLEGAL *)
r: ^real    (* ILLEGAL *)
```

Instead use integerObject, charObject, booleanObject, or realObject defined in the Mjølner System basic betaenv environment.

2.6 Repetitions

In BETA it is possible to declare a repetition of static (simple types) or dynamic references. A repetition is declared like:

```
R: [10] @integer  (* repetition of 10 static references *)
P: [10] ^A        (* repetition of 10 dynamic references *)
R[1] -> i         (* value assignment *)
P[1][] -> x[]     (* reference assignment *)
RR: [1] @integer  (* repetition of 1 static reference *)
R -> RR           (* repetition assignment:
                   * all values from R is copied into RR
                   * RR is automatically extended if needed
                   *)
R.range           (* the size of the repetition *)
n -> R.extend     (* extends the repetition with n elements *)
n -> R.new        (* allocates a new repetition with n elements *)
```

The range of a repetition is 1 to R.range, thus repetitions always start with 1.

2.7 Composite Types (Records)

Using the object−descriptor it is possible to declare composite types:

```
point:  (# x,y: @integer #) (* point is a composite type
                             * consisting of two integers *)
p: @point    (* static declaration of a point *)
p.x          (* remote access to x *)
circle:      (* composite type using simple and composite types *)
  (# center: @point;
     radius: @integer;
  #)
```

2.7.1 Pattern

The declaration of point and circle above is in general called a pattern declaration. The pattern will be described in details in the following sections.

3 Patterns and Objects

Most object−oriented languages supporting the object−oriented perspective have constructs such as class, subclass, virtual procedure, and qualified reference variable. These constructs all originated with Simula. Eiffel and C++ include these constructs although a different terminology is used. In addition to virtual procedures BETA also has non−virtual procedures.

In this introduction, the BETA version of the above constructs will be described and compared to other object−oriented languages. The example used in the following is a company with different kinds of employees, including salesmen and workers. employee is an abstract superpattern describing the common properties of all employees.

3.1 Pattern Employee

```
employee:
  (# name: @text;
     birthday: @date;
     dept: ^Department;
     totalHours: @integer;
     registerWork:
       (# noOfHours: @integer
       enter noOfHours
       do noOfHours + totalhours -> totalHours
       #);
     computeSalary:<
       (# salary: @integer
       do inner
       exit salary
       #);
  #);
```

The elements of the employee pattern have the following meaning:

3.2 Elements of Employee

- The attributes name, birthday, dept and totalHours are reference attributes denoting instances of the patterns text, date, department and integer respectively.
- Name, birthday, and totalHours refer to part−objects. A part−object is a fixed part of its enclosed object and is generated together with the enclosing object. Part−objects are also found in Eiffel and C++.
- Dept is a dynamic reference that either has the value NONE or refers to a separate instance of the pattern department.
- The attributes registerWork, and computeSalary are pattern attributes describing actions to be executed. They correspond to procedures in most other languages. The enter−part describes the input parameters of a pattern and the exit−part describes its output parameters. registerWork has one input parameter noOfHours and computeSalary has one output parameter, salary.
- registerWork is a non−virtual pattern attribute. This means that its complete description is given as part of the description of employee. It is similar to non−virtual functions in C++.
- computeSalary is a virtual pattern attribute (specified by using the ':<' symbol). Only part of its description is given since the computation of the salary is different for salesmen and workers. The description of a virtual pattern may be extended in subpatterns of employee. A virtual pattern attribute is similar to a virtual function in C++.

3.3 Class and procedure patterns

- employee, registerWork and computeSalary are all examples of patterns. employee is an example of a pattern used as a class and is therefore called a class pattern. registerWork and computeSalary are examples of patterns used as procedures and are therefore called procedure patterns. Technically there is no difference between class patterns and procedure patterns.

The following patterns are subpatterns of employee corresponding to salesmen and workers.

3.4 Subpatterns of Employee

```
worker: employee
  (# seniority: @integer;
     computeSalary::<
       (# do totalHours*80+seniority*4->salary; 0->totalHours #)
  #);
salesman: employee
  (# noOfSoldUnits: @integer;
     computeSalary::<
       (# do totalHours*80+noOfSoldUnits*6->salary;
             0->noOfSoldUnits->totalHours
       #)
  #)
```

- The class pattern worker adds the attribute seniority and extends the definition of computeSalary. The salary for a worker is a function of the totalHours being worked and the seniority of the worker.
- The class pattern salesman adds the attribute noOfSoldUnits and describes another extension of computeSalary. The salary for a salesman is a function of the totalHours being worked and the noOfSoldUnits.
- The symbol '::<' describe the fact that the definition of computeSalary from the superpattern employee is extended.

3.5 Part object

The above examples have shown instantiation of objects from patterns in the form of part−object attributes (like birthday: @date). An instance of, say worker, may in a similar way be generated by a declaration of the form:

```
mary: @worker
```

3.6 Dynamic reference

The above examples have also shown a dynamic reference (like dept: ^department). Such a reference is initially NONE. A dynamic reference to instances of worker may be declared as follows:

```
theForeman: ^worker
```

theForeman may be assigned a reference to the object referred by mary by execution of the following imperative:

```
mary[] -> theForeman[]
```

Note that the opposite assignment (theForeman[]−>mary[]) is not legal since mary is a static reference. An instance of worker may be generated and its reference assigned to theForeman by executing the following imperative:

```
&worker[] -> theForeman[]
```

A few additional comments about constructs used so far:

- The symbol & means new.
- The symbol −> is used for assignment of state.
- An expression R[] denotes the reference to the object referred by R, whereas an expression R denotes the object itself. The above assignment thus means that the qualified reference theForeman is assigned a reference to the generated instance of worker.
- An assignment of the form mary−>theForeman means that the state of the object referred by mary is enforced upon the state of the object referred by theForeman. This form of assignment is called value assignment. If X and Y are integer objects then X −> Y means that the value of X is assigned to the object Y.

In this section, it was shown how the most common object−oriented constructs may be expressed in BETA. In the following sections, examples of the more unique constructs will be given.

3.7 Singular Objects

Often there is only one object of a given type. In most languages it is necessary to make a class and generate a single instance. In BETA it is possible to describe a singular object directly. There is only one president of our company and he may be described as the following singular object:

```
president: @employee(# computeSalary::< (# do BIG -> salary #) #)
```

The declaration president is similar to the declaration of mary. The difference is that in the declaration of mary, a pattern name (worker) describes the objects whereas a complete object description is used to describe the president.

The president object is an example of a singular data object corresponding to an instance of a class pattern. It is also possible to describe singular action objects corresponding to an instance of a procedure pattern. Examples of singular action objects are given below.

3.8 Subprocedure

The previous sections has shown examples of patterns used as classes and procedures. For class patterns, examples of subpatterns have been given. Subpatterns may also be used for procedure patterns. For attributes, subpatterns may add new attributes and extend definitions of virtual patterns in the superpattern. In addition, a subpattern may specify further imperatives which have to be combined with the imperatives of the superpattern. The combination of the imperatives is handled by the inner construct. Consider the following objects:

```
mutex: @semaphore; sharedVar: @integer
```

The variable sharedVar is shared by a number of concurrent processes. Mutual access to the variable is handled by the semaphore mutex. Update of sharedVar should then be performed as follows:

```
mutex.P; m+sharedVar -> sharedVar; mutex.V
```

3.8.1 Abstract procedure pattern

This pattern of actions must be used whenever sharedVar and other shared objects have to be accessed. Instead of manipulating the semaphore directly it is possible to encapsulate these operations in an abstract procedure pattern. The pattern entry can describe this encapsulation:

```
entry: (# do mutex.P; inner; mutex.V #)
```

Execution of entry locks mutex before the inner and releases it afterwards. inner may then in subpatterns of entry be replaced by arbitrary imperatives. The subpattern updateShared of entry updates sharedVar:

```
updateShared: entry
  (# m: @integer
  enter m
  do sharedVar+m-> sharedVar
  #)
```

Execution of an imperative

```
123 -> updateShared
```

will then result in execution of the actions

```
mutex.P; sharedVar+123->sharedVar; mutex.V
```

We may now define an abstract superpattern corresponding to a monitor:

```
monitor:
  (# mutex: @semaphore;
     entry: (# do mutex.P; inner; mutex.V #);
     init:< (# do mutex.V(*initially open*); inner #)
  #);
```

A (singular) monitor object may now be declared like shared below:

```
shared: @monitor
  (# var: @integer;
     update: entry(# m: @integer enter m do var+m->var #);
     get: entry(# v: @integer do var->v exit v #)
  #)
```

Semaphores are the basic mechanism in BETA for synchronization. They can express most synchronization problems, but may be complicated to use. It is therefore mandatory that high level abstraction mechanisms like monitor can be defined. In section 9 below, further details about concurrency in BETA will be given.

3.9 Control Patterns

Sub (procedure) patterns are used intensively in BETA for defining control patterns (control structures). This includes simple control patterns like cycle, forTo, etc. It also includes so−called iterators on data objects like list, set and register. A pattern describing a register of objects may have the following interface:

3.9.1 scan is a control pattern

```
register:
  (# has: (# E: ^type; B: @boolean enter E[] do  exit B #);
     insert: (# E: ^type enter E[] do  #);
     delete: (# E: ^type enter E[] do  #);
     scan: (# current: ^type do  inner  #);

  #)
```

A number of details have been left out from the example. These include the representation and implementation of the register. A register may include instances of the pattern type, which has not been specified. Type is an example of a virtual class pattern which will be introduced later. For the moment type is assumed to stand for the pattern object which is a superclass of all patterns, i.e. a register may include instances of all patterns. An instance of register may be declared and used as follows:

```
employees: @register;

mary[]−>employees.insert;
(if boss[]−>employees.has then  if)
```

The control pattern scan may be used as follows:

3.9.2 Using a control pattern

```
0−>totalSalary;
employees.scan
  (# do current.computeSalary+totalSalary−>totalSalary #);
totalSalary−>screen.putint
```

This works as follows:

- The imperative employees.scan(# #) is an example of a singular action object as mentioned in section 4.
- The do−part of scan has an inner imperative which is executed for each element in the register. The details of this are not shown, but it may be implemented as a loop that steps through the elements of the register and executes inner for each element.
- The attribute current of scan is used as an index variable that for each iteration refers to the current element of the register. This may be implemented by assigning the reference of the current element to current before inner is executed.
- The effect of executing the above singular action object is that current.computeSalary+totalSalary−>totalSalary is executed for each element in the register.

3.10 Nested Patterns

One of the characteristics of Algol−like languages is block−structure, which allows for arbitrary nesting of procedures. The possibility of nesting has been carried over to BETA where patterns can be arbitrarily nested. Block−structure is a powerful mechanism that extends the modeling capabilities of languages. However, besides Simula and BETA, none of the mainstream object−oriented languages supports block−structure. In most object−oriented languages, an object may be characterized by data attributes (instance variables) and procedure attributes. In BETA, an object may in addition be characterized by class pattern attributes.

3.11 BETA supports general block−structure

In the examples presented so far, there have been two levels of nesting. The outer level corresponds to class patterns, like employee, and the inner level corresponds to procedure patterns, like computeSalary. In procedural languages like Algol and Pascal it is common practice to define procedures with local procedures. This is also possible in BETA.

3.12 Nested Class Patterns

The possibility of nesting classes is a powerful feature which is not possible in languages like C++ and Eiffel. The following example shows a class pattern that describes a product of our company:

```
productDescription:
  (# name: @text;
     price: @integer;
     noOfSoldUnits: @integer;
     order:
       (# orderDate: @date;
          c: ^customer;
          print:<
            (#
            do name[] −> puttext;
               'Price: '−>puttext; price −> putint; ' '−>put;
               ' No of units sold: '−>puttext;
               noOfSoldUnits−>putint; ' '−>put;
               orderDate.print;
               C.print;
               inner
            #)
       #)
  #);
```

One of the attributes of a productDescription object is the class pattern order. An instance of order describes an order made on this product by some customer. The attributes of an order object include the date of the order, the number of units ordered, the customer ordering the product, and a print operation. Consider the objects:

```
P1,P2: @product; o1,o2: @P1.order; o3,o4: @P2.order
```

The objects o1 and o2 are instances of P1.order whereas o3 and o4 are instances of P2.order. The block−structure makes it possible to refer to global names in enclosing objects. In the above example, the print operation refers to names in the enclosing order object. This resembles most object−oriented languages where operations inside a procedure refer to names in the enclosing object. The print operation, however, also refers to names in the surrounding productDescription object. Execution of say o1.print will thus print the values of P1.name, P1.price, P1.noOfSoldUnits, o1.orderDate, and o1.c.

4 Virtual Pattern

4.1 Structural equivalence is used in BETA

In the example in section 3 it was mentioned that a redefinition of a virtual procedure pattern is not a redefinition (overriding) as in C++. In fact a virtual pattern in BETA can only be extended and cannot be completely redefined. The rationale behind this is that a subpattern should have the same properties as its superpattern including which imperatives are executed. Ideally a subpattern should be behaviorally equivalent to its superpattern. This will, however, require a correctness proof. The subpattern mechanism of BETA supports a form of structural equivalence between a subpattern and its superpattern.

Consider the following patterns:

4.2 Patterns A and AA

```
A: (# V:< (# x:  do I1; inner; I2 #) #);
AA: A(# V::< (# y:  do I3; inner; I4#) #)
```

The pattern A has a virtual procedure attribute V. V has an attribute x and its do−part contains the execution of I1; inner; I2. The subpattern AA of A extends the definition of V. The extended definition of V in AA corresponds to the following object−descriptor (except for scope rules):

4.3 Combined descriptor

```
(# x: ; y:  do I1; I3; inner; I4; I2 #)
```

As may be seen the V attribute of AA has the attributes x and y and the do−part consists of I1; I3; inner; I4; I2. The definition of V is an extension of the one from A and not a replacement.

The subpattern AB of A describes another extension of V:

4.4 Pattern AB

```
AB: A(# V::< (# z:  do I5; inner; I6 #) #)
```

Here V corresponds to the following object descriptor:

```
V: (# x: ; z:  do I1; I5; inner; I6; I2 #)
```

The definition of V may be further extended in subpatterns of AA also as shown in the definition AAA:

4.5 Pattern AAA

```
AAA: AA(# V::< (# q:  do I7; inner; I8 #) #)
```

The definition of V corresponds to the following object−descriptor:

```
V: (# x: ; y: ; q:  do I1; I3; I7; inner; I8; I4; I2 #)
```

As may be seen, the pattern V is a combination of the definitions of V from A, AA and AAA.

4.6 Final binding

The definition of V may be extended using a final binding (::) in subpatterns of A as shown in the definition AC:

```
AC: A(# V::(# q:  do I2; inner; I4 #) #)
```

The final binding of V means that V cannot be extended in subpatterns of AC. The extended definition of V in AC corresponds to the following object−descriptor (except for scope rules):

```
(# x: ; y:  do I1; I3; inner; I4; I2 #)
```

The virtual mechanism in BETA guarantees that behavior defined in a superpattern cannot be replaced in a subpattern. This form of structural equivalence is useful when defining libraries of patterns that are supposed to execute a certain sequence of actions. In C++, the programmer must explicitly invoke the actions from the superclass by means of superclass::functionname. This is illustrated by the example in the next section.

The inner construct is more general than shown above, since a pattern may have more than one inner and inner may appear inside control structures and nested singular object descriptors.

4.7 Virtual Procedure Pattern

The attribute computeSalary of pattern employee is an example of a virtual procedure pattern. In this example the do−part of the virtual definition in employee is very simple, only consisting of an inner−imperative. The extended definitions of computeSalary in worker and salesman both include the code noOfHours*80 and 0−>totalHours. This code may instead be defined in the definition of computeSalary in employee as shown below:

4.7.1 ComputeSalary is a virtual procedure pattern

```
employee:
  (#
     computeSalary:<
       (# salary: @integer
       do noOfHours*80−>salary; inner; 0−>totalHours
       exit salary
       #)
  #);
  worker: employee
    (#
       computeSalary::< (# do seniority*4+salary−>salary; inner #)
    #);
  salesman: employee
    (#
       computeSalary::<
         (#
         do noOfSoldUnits*6+salary −>salary;
            0 −>noOfSoldUnits;
```

```
            inner
         #)
    #)
```

The extended definitions of computeSalary in worker and salesman have an inner to enable further extensions of computeSalary in subpatterns of worker and salesman.

4.8 Virtual Class Pattern

Virtual patterns may also be used to parameterize general container patterns such as the register pattern described above. For the register pattern we assumed the existence of a type pattern defining the elements of the register, i.e. elements of a register must be instances of the pattern type. The pattern type may be declared as a virtual pattern attribute of register as shown below:

4.8.1 type is a virtual class pattern

```
register:
  (# type:< object;
     insert:< (# e: ^type enter e[] do #)

  #)
```

The declaration type:< object specifies that type is either the pattern object or some subpattern of object. In the definition of register, type may be used as an alias for object, e.g. references qualified by type are known to be at least Objects. Since object is the most general superpattern, type may potentially be any other pattern. The virtual attribute type may be bound to a subpattern of object in subpatterns of register. The following declaration shows a pattern workerRegister which is a register where the type attribute has been bound to worker:

```
workerRegister: register
  (# type::< worker;
     findOldestSeniority:
       (# old: @integer
       do scan
           (# do (if current.seniority > old then
                     current.seniority->old
           if)#)
       exit old
       #)
  #);
```

In the definition of workerRegister, the virtual pattern type may be used as a synonym for the pattern worker. This means that all references qualified by type may be used as if they were qualified by worker. The reference current of the scan operation is used in this way by the operation findOldestSeniority which computes the oldest seniority of the register. The expression current.seniority is legal since current is qualified by type which in workerRegister is at least a worker.

In subpatterns of workerRegister it is possible to make further bindings of type thereby restricting the possible members of the register. Suppose that manager is a subpattern of worker. A manager register may then be defined as a subpattern of workerRegister:

```
managerRegister: workerRegister(# type::< manager #)
```

In the definition of managerRegister, type may be used as a synonym for manager, i.e. all references qualified by type are also qualified by manager.

4.8.2 General parameterized patterns

Virtual patterns make it possible to define general parameterized patterns like register and to restrict the member type of the elements. In this way virtual class patterns provide an alternative to templates as found in C++.

5 Coroutines and Concurrency

A BETA object may be the basis for an execution thread. Such a thread will consist of a stack of objects currently being executed. An object which can be used as the basis for an execution thread has to be declared as an object of kind component as shown in the following declaration:

5.1 Components with execution threads

```
A: @|activity
```

The symbol '|' describes that the object A is a component. A component (thread) may be executed as a coroutine or it may be forked as a concurrent process. Consider the following description of activity:

```
activity:
  (#
  do cycle
      (#
      do getOrder; suspend;
         processOrder; suspend;
         deliverOrder; suspend
  #)#)
```

The component object may be invoked by an imperative

```
A
```

which implies that the do−part is executed. The execution of A is temporarily suspended when A executes a suspend−imperative. In the above example this happens after the execution of getOrder. A subsequent invocation of A will resume execution after the suspend−imperative. In the above example this means that processOrder will be executed. If B is also an instance of activity, then the calling object may alternate between executing A and B:

```
cycle(# do A;  B;  #)
```

The above example shows how to use components as deterministic coroutines in the sense that the calling object controls the scheduling of the coroutines. In section 9.1 below another example of using coroutines will be given.

It is also possible to execute component objects concurrently. By executing

```
A[]->fork; B[]->fork
```

the component objects A and B will be executed concurrently. As for the deterministic coroutine situation, A and B will temporarily suspend execution when they execute a suspend−imperative. Further examples of concurrent objects will be given below in section 9.2.

5.2 Coroutines

Deterministic coroutines have demonstrated their usefulness through many years of usage. Below we give a typical example of using coroutines.

Suppose we have a register for the permanent workers and another one for the hourly paid workers. Suppose also that it is possible to sort these registers according to a given criterion like the total hours worked by the employee. Suppose that we want to produce a list of names of all employees sorted according to the total hours worked. This may be done by merging the two registers. A register object has a scan operation that makes it possible to go through all elements of the register. Instead we define an operation of register in the form of a coroutine getNext, which delivers the next element of the register when called:

```
register:
  (#
     getNext: | @
       (# elm: ^employee
       do scan(# do current[]->elm[]; suspend #);
          none->elm[]
       exit elm[]
       #);
  #);
  pReg: @permanentRegister; hReg: @hourlyPaidRegister;

  pReg.getNext->e1[]; hReg.getnext->e2[];
  L: cycle
    (#
    do (if e1[] = none then (*empty hReg*); leave L if);
       (if e2[] = none then (*empty pReg*); leave L if);
       (if e1.totalHours < e2.totalHours then
           e1.print; pReg.getNext->e1[]
        else
           e2.print; hReg.getNext->e2[]
       if)
    #)
```

5.2.1 Suspending and resuming

The attributes getNext of the objects pReg and hReg have their own thread of execution. When called in an imperative like pReg.getNext−>e1[], the thread is executed until it either executes a suspend or terminates. If it executes a suspend, it may be called again in which case it will resume execution at the point of suspend. The first time getNext is called, it will start executing scan. For each element in the register, it will suspend execution and exit the current element via the exit variable elm[]. When the register is empty, NONE is returned.

5.3 Concurrency

As previously mentioned, it is possible to perform concurrent execution of components by means of the fork operation as sketched in the following example:

```
(# S1: @| (#  do  #);
   S2: @| (#  do  #);
   S3: @| (#  do  #)
  do S1[] -> fork; S2[] -> fork; S3[] -> fork;
#)
```

The execution of S1, S2 and S3 will take place concurrently with each other and with the object executing the fork operations. Concurrent objects may access the same shared objects without synchronization, but may synchronize access to shared objects by means of semaphores. In section 5 above the pattern semaphore has been described. It is well known that a semaphore is a low level synchronization mechanism which may be difficult to use in other than simple situations. For this reason the Mjølner library has a number of patterns defining higher level synchronization mechanisms. This library includes a monitor pattern as described in section 5 above. The library also includes patterns defining synchronization in the form of rendezvous as in Ada.

5.4 Monitor Example

The following example describes a company with a number of salesmen, workers and carriers. The salesmen obtain orders from customers and store them in an order pool. The workers obtain orders from the order pool, process them and deliver the resulting item in an item pool. The carriers pick up the items from the item pool and bring them to the customer. Salesmen, workers and carriers are described as active objects whereas the order− and item pools are represented as monitor objects.

```
(# salesman: employee
     (# getOrder: (#  exit anOrder[] #)
     do cycle (# do getOrder −> jobPool.put #)
     #);
   S1,S2, : @|salesman;
   jobPool: @monitor
     (# jobs: @register(# type::< order #);
        put: entry
          (# ord: ^order enter ord[] do ord[] −>jobs.insert #);
        get: entry
          (# ord: ^order do jobs.remove −> ord[] exit ord[] #)
     #);
   worker: employee
     (# processJob: (#  enter anOrder[] do  exit anItem[] #)
     do cycle(# do jobPool.get −> processJob −> itemPool.put #)
     #);
   W1,W2,: @| worker;
   itemPool: @monitor(#  #);
   carrier: employee
     (# deliverItem: (# enter anItem[] do  #)
     do cycle(# do itemPool.get −>DeliverItem #)
     #);
   C1,C2, : @| carrier;
do jobPool.init; itemPool.init;
   conc(# do S1[]−>start;  W1[]−>start;  C1[]−>start;  #)
#)
```

5.4.1 Procedure pattern conc

The procedure pattern conc is another example of a high−level concurrency pattern from the Mjølner library. It does not terminate execution until components being started (by S1[]−>start, etc.) have terminated their execution.

5.4.2 Rendezvous Example

Next we show an example of using the library patterns for describing synchronized rendezvous. The example shows a drink machine that provides coffee and soup. A customer operates the machine by pushing either makeCoffee or makeSoup. If makeCoffee has been pushed, then the customer may obtain the coffee by means of getCoffee. Similarly if makeSoup has been pushed then the soup may be obtained by means of getSoup.

The system pattern has a port attribute which may be used to define synchronization ports. The drink machine described below has three such ports, activate, coffeeReady, and soupReady. A port object has a pattern attribute entry which may be used to define procedure patterns associated with port. For the port activate, two procedure patterns makeCoffee and makeSoup are defined. For coffeeReady and soupReady, the procedure patterns getCoffee and getSoup are defined.

An execution of a port−entry operation like aDrinkMachine.makeCoffee will only be executed if the drinkMachine has executed a corresponding accept by means of activate.accept.

- Initially a drinkMachine is ready to accept either makeCoffee or makeSoup.
- If e.g. makeCoffee is executed, then when 'the coffee has been made', the drinkMachine is willing to accept the operation getCoffee. This is signaled by executing an accept on the port coffeeReady. Technically this is implemented by assigning a reference to coffeeReady to the port reference drinkReady. The do−part of drinkMachine then makes an accept on drinkReady.
- When the operation getCoffee, has been executed, the drinkMachine is again ready to accept a new operation associated with the activate port.

```
drinkMachine: system
  (# activate: @port;
     makeCoffee: activate.entry
       (# do  coffeeReady[]->drinkReady[] #);
    makeSoup: activate.entry(# do  soupReady[]->drinkReady[] #);
    coffeeReady, soupReady: @port;
    getCoffee: coffeeReady.entry(# do  exit someCoffee [] #);
    getSoup: soupReady.entry(# do  exit someSoup [] #);
    drinkReady: ^port
  do cycle(# do activate.accept; drinkReady.accept #)
  #)
```

The drinkMachine may be used in the following way:

```
aDrinkMachine: @| drinkMachine

  aDrinkMachine.makeCoffee;  aDrinkMachine.getCoffee;
  aDrinkMachine.makeSoup;  aDrinkMachine.getSoup;
```

As may be seen the use of the patterns system, port and entry makes it possible to describe a concurrent program in the style of Ada tasks that synchronize their execution by means of rendezvous. A port object defines two semaphores for controlling the execution of the associated entry patterns. The actual details will not be given in this language introduction.

It is possible to specialize the drinkMachine into a machine that accepts further operations:

```
extendedDrinkMachine: drinkMachine
  (# makeTea: activate.entry(# do  teaReady[]->drinkReady[] #);
     teaReady: @port;
     getTea: teaReady.entry(#  exit someTea[] #)
  #)
```

The extendedDrinkMachine inherits the operations and protocol from drinkMachine and adds new operations to the protocol.

The basic mechanisms in BETA for providing concurrency are component−objects (providing threads), the fork−imperative (for initiating concurrent execution) and the semaphore (for providing synchronization). As has been mentioned already, these mechanisms are inadequate for many situations. The abstraction mechanisms of BETA make it possible to define higher−level abstractions for concurrency and synchronization.

5.4.3 More information

Please see the manual [MIA 90−8] for details about the concurrency library.

6 Inheritance

The subpattern mechanism combined with the possibility of redefining/extending virtual procedures is widely recognized as a major benefit of object−oriented languages. This mechanism is often called inheritance since a subpattern is said to inherit properties (code) from its superpattern. Inheritance makes it easy to define new patterns from other patterns. In practice this has implied that subpatterns are often used for sheer inheritance of code without any concern for the relation between a pattern and its subpatterns in terms of generalization/specialization. The use of multiple inheritance is in most cases justified in inheritance of code and may lead to complicated inheritance structures.

6.1 Classification and inheritance

In BETA subpatterns are intended for representing classification and inheritance of code is a (useful) side effect. In BETA it is not possible to define a pattern with multiple superpatterns corresponding to multiple inheritance. There are indeed cases where it is useful to represent classification hierarchies that are not tree structured. However, a technical solution that justifies the extra complexity has not yet been found.

6.2 Inheritance from part−objects

BETA does support multiple inheritance, but in the form of inheritance from part−objects. A compound object inherits from its parts as well as its superpattern. The reason that this has not been more widely explored/accepted is that in most languages inheritance from part−objects lacks the possibility of redefining/extending virtual procedures in the same way as for inheritance from superpatterns. Block−structure and singular objects make this possible in BETA.

Assume that we have a set of patterns for handling addresses. An address has properties such as street name, street number, city, etc., and a virtual procedure for printing the address. In addition we have a pattern defining an address register.

```
address:
  (# streetName: @text; streetNo: @integer; city: @text;
     print:<
       (#
       do inner;
          streetName−>puttext;
          streetNo−>putint; (*etc.*)
       #);
  #);
addressRegister: register(# element::< address #)
```

We may use the address pattern for defining part−objects of employee/company objects:

```
employee:
  (# name: @text; {the name of the employee*)
     adr: @address(# print:: (# do name−>puttext #)#)
  #);
company:
  (# name: @text; (*the name of the company*)
     adr: @address(# print:: (# do name−>puttext #) #)
  #);
```

The object adr of employee is defined as a singular address object where the virtual print pattern is defined to print the name of the employee. As can be seen it is possible to define a part−object and define its virtual procedures to have an effect on the whole object. The company pattern is defined in a similar way.

It is possible to handle the address aspect of employees and companies. An example is an address register:

```
AReg: @addressRegister;

employee1.adr[]->AReg.insert; employee2.adr[]->AReg.insert;
company1.adr[]->AReg.insert; company2.adr[]->AReg.insert;
AReg.scan(# do current.print #)
```

The AReg register will contain address objects which are part of either employee objects or company objects. For the purpose of the register this does not matter. When the print procedure of one of these address objects is invoked it will call the print procedure associated with either employee or company. The scanning of the AReg register is an example of invoking the print pattern.

The example shows that in BETA inheritance from part−objects may be used as an alternative to inheritance from superpatterns. The choice in a given situation depends of course on the actual concepts and phenomena to be modeled. In the above example it seems reasonable to model the address as a part instead of defining employee and company as specializations of address.

In general it is possible to specify multiple inheritance from part−objects since it is possible to have several part−objects like the address object above. This form of multiple inheritance provides most of the functionality of multiple inheritance from C++ and Eiffel. It is simpler since the programmer must be explicit about the combination of virtual operations. It does, however, not handle so−called overlapping superclasses. The programmer must also explicitly redefine the attributes of the component classes. This may be tedious if there is a large number of attributes. However, a renaming mechanism for making this easier has been proposed for BETA, but it is not yet implemented in the Mjølner System. Multiple inheritance from part−objects should be used when there is a part−of relationship between the components and the compound. This also covers situations where implementations are inherited. It should not be used as a replacement for multiple specialization hierarchies.

A common example of using multiple inheritance is modeling windows with titles and borders. This may be modeled using block−structure. Since a window may have a title, a border or both, the following class hierarchy using multiple inheritance is often used:

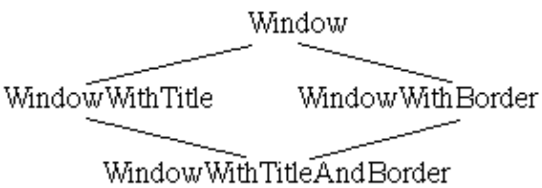

In BETA this can be described using nested patterns:

```
window:
  (# title: (#  #);
     border: (#  #);

  #);
aWindow: @window(# T: @title; B: @border #)
```

The descriptions for title and border are made using nested patterns. For a given window, like aWindow, a title object and a border object may be instantiated. If e.g. two titles are needed, two instances of title are made. This example illustrates another situation where multiple inheritance may be avoided.

7 BETA Terminology

The following is a short description of important concepts used in the BETA language. Please note, that these descriptions are deliberately informal. The precise meanings of these terms must be found in [MMN 93].

Contents

7.1 Modelling

Object−oriented programming

A *program execution* is viewed as a *physical model* or *representation* of part of the world. *Objects* on the computer model phenomena in the world; *attributes* of objects model properties of phenomena.

Computer

Real world

Object

Phenomenon

Attribute

Property

Pattern

Concept

BETA program execution

A collection of *objects*. Some represent phenomena while others are simply part of the implementation.

Object

Computer representation of a real world phenomenon. Its structure consists of attributes and actions.

Pattern

Computer representation of a real world concept. Objects defined according to the pattern are called *instances* or *pattern defined objects*: A pattern is to its instances as a concept is to its phenomena.

Singular object

An object representing a singular "one−of−a−kind" phenomenon − the object is not defined as an instance of some pattern.

State of an object

The combined values of its *measurable properties* at some point in time.

Measurable property

A property which has a measurable value. The value may vary over time.

Part object

An object which is part of another object. Part objects are used to model part or aggregation hierarchies.

Separate object

An autonomous self−contained object which is not a part object.

Reference to separate object

An attribute which is a reference to a separate object.

Kinds of actions

The actions of a phenomenon in a real world system often take place *concurrently* (i.e. in parallel) with those of other phenomena in the system. A single phenomenon normally *alternates* among its own actions.

7.2 Declarations and Object Descriptors

Declaration or attribute declaration

An association or *binding* of a name to some entity. The syntactic construct used is the colon ":" as in, <name>: <entity**>**. For attributes of an object descriptor, these are sometimes referred to as the *attribute name* and *attribute description*, respectively.

Pattern declaration

A declaration binding a *pattern name* to an object descriptor, describing the structure of *instances of the pattern*. Pattern declarations serve as templates for generating objects having a given structure.

Syntax is:

```
<name>: <object−descriptor>
```

Singular object declaration

Declaration of a singular object binding the object name to the singular object description (an object descriptor).

Syntax is:

```
<name>: @<object−descriptor>
```

Attribute reference

An occurrence of an attribute's name in an object descriptor.

Local attribute reference of a pattern

A reference from within a pattern's object descriptor to an attribute declared inside the same object descriptor.

Global attribute reference

Any attribute reference which is not local.

Object−descriptor

Used to describe the structure of objects and consists of a *prefix part* and a *main part*.

Prefix part

Part of object descriptor used to specify the superpattern of the descriptor. The prefix part is specified by a pattern name (or is empty).

Main part

Used to describe the additional structure of objects. Has the syntactic form(# E #) and consists of an *attribute part and an action part.*

Attribute part

Part of object descriptor used to describe the object's attributes. Consists of a list of *attribute declarations*.

Action part

Part of object descriptor used to describe the actions to be performed when the object is executed. Consists of three parts: *enter−part*, *do−part*, *exit−part*.

Enter part

Part of *action part* describing the *enter parameters*.

Do part

Part of *action part* consisting of a list of *imperatives*.

Exit part

Part of *action part* describing the exit *parameters*.

Program

An object descriptor that can be compiled and executed.

7.3 Reference Attributes

Reference attribute

An attribute that denotes an object. Reference attributes can be either *static references* or *dynamic references*.

Static reference

A reference attribute that constantly denotes the same object. Such objects are often referred to as *static objects*. In cases where these objects are used to model part (or aggregation) hierarchies, they are referred to as *part objects*, that is, they are part of an *enclosing object*.

Static reference declaration

Used to define static reference attributes.

Syntax is:

```
 <name>: @<ptn.name or
obj.descriptor>
```

Dynamic reference

A reference attribute that denotes a object. The reference is variable in that it may denote different objects over time. Initially it denotes *NONE* which represents "no object."

Dynamic reference declaration

Used to define dynamic reference attributes.

Syntax is:

```
 <name>: ^<pattern name>
```

Indexed collection of static / dynamic references

A repetition (or *array*) of object references referred to by a single name plus an index. The size of a repetition A is denoted by A.range. A[1] refers to the first element in the repetition, A[A.range] to the last.

Syntax is:

```
Name: [eval] @<ptn.name or obj.descriptor>
Name: [eval] ^<ptn.name>
```

The size of the repetition can be dynamically extended by:

```
<number> -> A.extend
```

Qualification or qualifying pattern

The pattern name appearing in a reference attribute declaration. It restricts the set of objects that can be denoted by the reference.

Remote access

Used to denote attributes within an enclosing object.

Syntax is:

```
reference.attribute
```

Computed Remote access

Used to denote attributes within objects that are returned as the result of evaluations.

Syntax is:

```
(evaluation).attribute
```

7.4 Pattern Attributes

Pattern reference

A reference attribute that denotes a pattern. The structure of the pattern is represented locally using a *structure object*. Such objects include a reference back to the object of which the pattern is an attribute. This reference is called the *origin* of the pattern.

Pattern reference declaration

Used to define a pattern.

Syntax is:

```
<name>: <object descriptor>
```

Pattern variable declaration

Used to define pattern variable attributes. A pattern variable may denote different patterns during the execution. The qualification restricts the set of patterns which may be denoted by the pattern variable.

Syntax is:

```
<name>: ##<pattern name>
```

Class pattern

Generally, a pattern used to model physical objects.

Procedure pattern

Generally, a pattern used to model action sequences.

Function pattern

A procedure pattern which computes and returns a value. Such patterns always have an exit part.

Basic pattern

A pattern that is predefined within the BETA language. Examples are integer, real, boolean, and char. Relevant operations include: +, −, *, div, mod, and, or, not, true, false, =, <, >, <>, <=, >=.

7.5 Imperatives

Imperative

Describes an action; *executing* the imperative causes the action. Imperatives appear in the do−part of an object. Kinds of imperatives include *evaluations*, *reference assignments*, *dynamic object creation*, and *control structures*.

Evaluation imperative

An imperative that can cause state changes and may produce a value when executed.

Value assignment

An evaluation imperative that sets (changes) the value of an attribute.

Syntax is:

```
3 -> I
```

Reference assignment

An *imperative* used to change the value of a dynamic reference.

Syntax is:

```
objRef[] -> dynObjRef[]
```

objRef may be any object reference but dynObjRef *must be a dynamic object reference.*

Pattern assignment

An *imperative* used to change the pattern denoted by a pattern variable.

Syntax is:

```
ref## -> dynPatRef##
```

Ref may be the name of a pattern variable, the name of an object, or the name of a pattern but dynPatRef *must be a dynamic pattern reference.*

Multiple assignment

An evaluation imperative that causes several assignments.

Syntax is:

```
3 -> I -> J
```

Dynamic object creation / generation

Imperatives used to create new *dynamic objects*.

Syntax is:

```
&Pat   or   &Pat[]
```

Value equality

True when two references denote objects that have the same state.

Syntax is:

```
A = B
```

Reference equality

True when two references denote the same object.

Syntax is:

```
A[] = B[]
```

Pattern equality

True when two pattern references denote the same pattern.

Syntax is:

```
A## = B##
```

Note that < and <= are also defined for pattern comparisons based on the inheritance hierarchy.

Procedure call

An evaluation imperative that causes invocation of a procedure pattern.

Syntax is:

```
&ProcPat
```

or

```
 (arg1,arg2) -> &ProcPat
```

Function call

An evaluation imperative that causes invocation of a function pattern.

Syntax is:

```
 (arg1,arg2) -> &FuncPat ->
result
```

Control structure

An imperative that controls the flow of executions.

For imperative

A control structure used to support *iteration*. A list of imperatives are executed repeatedly while an index steps from 1 up to the number of iterations.

Syntax is:

```
(for Index: Range repeat
      Imperative-list
```

for)

General−if imperative

A control structure used to support *selection*. Based on evaluating a condition evaluation and comparing it to the values of a number of selection evaluations, one of a set of imperative−lists is executed.

Syntax is:

```
(if E0
  // E1 then I1
  // E2 then I2
  E
  // En then In
  else I
  if)
```

Simple−if imperative

A control structure used to support boolean *selection*. Based on evaluating a condition evaluation and testing if it is true or false, one of two imperative−lists is executed.

Syntax is:

```
(if E then
     I1
```

```
  else I2
 if)
```

Labelled imperative

A means of naming an imperative. References to the label (via *jump imperatives*) can be made from within the imperative.

Syntax is:

```
 L: Imperative
```

or

```
 L: (# ... do ... #)
```

Jump imperative

Causes flow of control to "jump" to another location. A jump imperative is one of a *Leave imperative* or a *Restart imperative*.

Leave imperative

Causes termination of the execution of a labelled imperative; execution resumes after the labelled imperative. This imperative can only appear within the labelled imperative.

Syntax is:

```
leave L
```

Restart imperative

Causes restarting of the execution of a labelled imperative, that is, jump is to the start of the imperative. Can only appear within the labelled imperative.

Syntax is:

```
restart L
```

7.6 Block Structure and Scoping

Block structure

The nesting of one structure in another in the text of a program. In BETA, object descriptors and imperatives can be nested inside of other object descriptors and imperatives. It is the job of the programmer to use *indentation* to make such nesting visible to readers. In the following example, Deposit's object descriptor is nested inside of Account's.

```
 Account:
   (# Deposit:
        (# E
        do E
        #);
   #);
```

Declaration of a name

An association of a name with some defining expression.

Syntax is:

```
<name>: E
```

Recall that colon ":" always signals a declaration of some kind.

Application of a name

Any occurrence of a name in a program which is not a declaration. Note that this does not include keywords of the BETA syntax (e.g. if, for, repeat, do), but does include predefined pattern and attribute names (e.g. char, putInt, stream).

Scope of a declaration

The part of the program text "covered" by the declaration, that is, where applications of the declared name refer to the given declaration. In BETA, the scope of a declaration is the object descriptor it appears in. The exception to this is that the declaration may be "hidden" by declarations of the same name in nested object descriptors or labelled imperatives. Note that the declared name can also be applied outside its object descriptor using remote access. We say that a name is *local* to the object descriptor in which it is declared and *global* to any nested object descriptors (for which it is not hidden).

7.7 Inserted Objects

Inserted item

A means of generating (and executing) a procedure object allocated as part of the enclosing object.

Syntax is:

```
A -> P -> B
```

or

```
A -> P(# E #) -> B
```

This differs from dynamic generation, &P, in that the instance of P is generated only once rather than each time the imperative is executed. Note that inserted items should not be used to define recursive procedures. That is, an inserted instance of P may be specified in the action part of P.

7.8 Inheritance

Direct subpattern

A pattern P is a *direct subpattern* of Q if P extends (specialises) the definition of Q. Q is called the *direct superpattern* of P and instances of P are also instances of Q.

Syntax is:

```
P: Q(# E #)
```

Q is called the *prefix pattern* (or simply *prefix*), while the contents of (# E #) is called the *main−part* of P. The prefix Q means that P's object descriptor *inherits* all of Q's declarations in addition to any new ones defined in P's main−part.

Subpattern

A pattern P is a *subpattern* of Q if it is either a direct subpattern of Q or a subpattern of a direct subpattern of Q. Likewise, Q is a *superpattern* of P if it is either a direct superpattern of P or a superpattern of the direct superpattern of P. A pattern can have at most one direct superpattern.

Abstract superpattern

A pattern used only as a superpattern for other patterns, that is, it is not intended to be used to generate objects. If P is declared without the use of a superpattern, P: (# E #), then P is assumed to be a subpattern of the most general abstract superpattern, Object. Note that the basic patterns, Integer, Real, Boolean, Char and Real are not subpatterns of Object.

Superpattern as qualification

If R is a dynamic reference qualified by the pattern Q (i.e. R: ^Q) and Q is a superpattern of P, then instances of both P and Q can be assigned to R. However, only attributes of Q (and of superpatterns of Q) can be accessed using remote access through R. That is, if attribute A is declared in the main part of P, then the remote access R.A is illegal.

Action specialisation

The use of a subpattern to extend the action part of a pattern. Action specialisation can involve any or all of the enter−part, exit−part and do−part. The enter and exit parts of instances of P (again, a subpattern of Q) consist of Q's enter and exit parameters together with those defined by P. Extending the do−part of Q requires the use of the **inner** imperative in Q's action part. Executing the do−part of an instance of P proceeds by executing Q's do−part and executing P's do−part each time inner is encountered.

Syntax is:

```
Q: (# E do E inner E #);
P: Q(# E do E #);
```

7.9 Virtual Patterns

Virtual pattern

A pattern attribute V of a pattern Q is *virtual* if it is only partially defined in Q. That is, the definition of V can be extended in subpatterns of Q.

Syntax is:

```
Q: (# V:< S #)
Q: (# V:< S0(# E #) #)
Q: (# V:< (# E #) #)
```

In the first of the three forms, we say that the virtual V is *qualified* by the pattern S, in the second and third forms, we say that V is *directly qualified*.

Further binding of a virtual pattern

The means by which a virtual attribute V of a pattern Q is extended in a subpattern P of Q.

Syntax is:

```
P: Q(# V::< S1 #)
P: Q(# V::< S1(# E #) #)
P: Q(# V::< (# E #) #)
```

S1, S1(# E #), or (# E #) is called the *extended descriptor* of V. If we're using either the first or second form, and if V is qualified by S in the pattern Q, then S1 must be a subpattern of S. In the case of the third form there are no constraints on Q's declaration of V. If X is an instance of P, then X.V specialises (that is, adds properties to) the definition of V in Q. Note that V is now a virtual pattern in P (as well as Q) and can continue to be further bound in subpatterns of P.

Final binding of a virtual pattern

The means by which a virtual attribute V of a pattern Q is extended in a subpattern P of Q, and at the same time made non−virtual.

Syntax is:

```
R: P(# V:: S2 #)
R: P(# V:: S2(# E #) #)
R: P(# V:: (# E #) #)
```

Final binding is identical to further binding, except that with final binding, V is no longer virtual.

8 BETA Quick Reference Card

A summary of all special characters in BETA, and a short list of the syntax of the language is given below along with a short description of their semantics:

Special characters	**Semantics**
:	Declaration
: @	Static object reference declaration
: ^	Dynamic object reference declaration
: ##	Pattern reference declaration
: @|	Static component declaration
: ^|	Dynamic component declaration
: [range]	Declaration of repetition. range must be an integer evaluation
:<	Virtual declaration
::<	Extended binding of virtual declaration
::	Final binding of virtual declaration
&	Dynamic creation of item; new
&|	Dynamic creation of component
−>	Assignment
[]	Reference
##	Pattern reference
(#	Object descriptor begin
#)	Object descriptor end
//	Selection in if−imperative

Keywords
do else enter exit inner leave none repeat restart suspend then this (if if) (for for) Additional keywords (for their usage, see below)

Short syntax	Semantics
P: (# E do E #)	Definition of a pattern
PP: P(# E do E #)	Definition of a subpattern
enter E	Specification of enter−parameters
exit E	Specification of exit−parameters
inner P	Execute the actions in the subpattern. P is an optional name of an enclosing pattern.
this(P)	Denotation of this object
this(P)[]	Reference to this object
E.P	Remote name
(E).P	Computed remote name
L: Imp	In action part: labelled imperative
L: (# E do E #)	In action part: labelled imperative (descriptor)
leave L	Terminate labelled imperative or object instance L
restart L	Goto beginning of labelled imperative or object instance L
suspend	Component suspension

E1 −> E2	Assignment imperative
(if E // E1 then Imp1 // En then Impn else Imp if)	General selection imperative: Sequential evaluation of E, E1, E En First Impi is executed where Ei=E If no Ei=E, then Imp is executed 'else Imp' is optional
(if E then **Imp**1 else **Imp**2 if)	Simple if imperative: Evaluation of E (must exit a single boolean value); Execute Imp1 if E is true; Otherwise Imp2 is executed 'else Imp2' is optional
(for I: range repeat Imp for)	Repetition imperative: I is a locally scoped integer variable within Imp. Execute Imp with I assigned each value in [1..range]
NONE	The nil reference value
R[i:j]	Repetition slice
R[i]	Indexed repetition element
(e1, e2, E, en)	Evaluation list

Please note, that the above description is by no means complete, and in some cases ambiguous. The ultimate reference is naturally the BETA grammar as defined in the BETA book [MMN 93].

Index

The entries in the alphabetic index consists of selected words and symbols from the body files of this manual − these are in **bold** font − as well as the identifiers defined in the public interfaces of the libraries − set in regular font.
In the manual, the entries, which can be found in the index are typeset like this. This can help localizing the identifier, when the link from the index if followed − especially in the case where the browser does not scroll the line to the top, e.g. because there is less than a page of text left.
In the small table of letters and symbols below, each entry links directly to the section of the index containing entries starting with the corresponding letter or symbol.

& > (− / : @ [^ A B C D E F G H I J K L M N O P Q R S T U V W X Y Z |

#

&

(

−

/

:

>

@

[

^

A

B

C

D

E

F

G

I

J

K

L

M

N

O

P

Q

R

S

T

V

|