

Value Types and Object Classes

Ole-Johan Dahl
Dept. of Informatics, University of Oslo

March 22, 1999

Abstract

Elements of the specification and programming language ABEL are explained by examples. It is argued that a programming language should allow a maximum of programming detail to be expressed applicatively, however, imperative mechanisms are useful for controlling efficiency in terms of computation time and memory space, and expressing interactions with the computer environment as well as concurrent computer processes. The applicative sublanguage should be well separated from the imperative language mechanisms. An important side issue is the use of data pointers in ways that are not detrimental to the program logic. Concurrency is not considered.

1 Introduction

Computer programs control enormous sequences of state transitions in vast computer stores. We develop such programs by introducing program components which somehow implement problem oriented concepts, thereby aggregating, or even hiding, transition sequences and data elements. In our reasoning about a program we may thus replace complicated implementations by simpler external “abstract” specifications, possibly based on complete reinterpretations of data elements. The actual program development can be in terms of the latter. Usually several layers of abstraction occur.

An applicative program is an expression consisting of function applications, whose free variables represent the program input and whose value is the program output. Each function is implemented in such a way that the only effect of a function application, observable within the programming language, is to provide its value, and the state transitions which occur during its evaluation are therefore irrelevant (given that sufficient data storage and processing speed are available). Consequently any applicative expression may be seen as a representation of its value, which implies that the substitution property obtains. Thus, any applicative program may be analyzed by standard axiomatic techniques (assuming evaluation by a sufficiently powerful computer).

Imperative programming mechanisms are, however, needed in general, for several pragmatic reasons:

- for the explicit control of efficiency with respect to computation time and memory space,
- for directly expressing sequences of interactions with the program environment, as well as with concurrent processes and processors, and
- for reasons of style. Some algorithms are best expressed in terms of sequences of actions, and, furthermore, some applications are explicitly concerned with changes of state, such as data base systems.

For ease of reasoning about imperative programs it is important that statements are free of “side effects”. This means that expressions should remain purely applicative,

and there should be no data aliases giving rise to side effects of assignment operations, even if data pointers are used for efficiency reasons. Modularization is important.

In the following sections elements of a specification and programming language ABEL (Abstraction Building Experimental Language) are explained through simple examples. ABEL has been developed at Oslo University, mainly by O. Owe and the author, as part of a research effort in program specification and verification (see e.g. [1]) In addition to the design objectives mentioned above, consistency in specification and programming is supported through strong typing, made more flexible through subtypes, function polymorphism, and overloading. There are three language levels, confined to within different kinds of modules:

1. **property** modules containing non-constructive specifications by axioms in first-order predicate style (not further discussed here),
2. **type** modules containing constructive specifications of one or more value types and associated functions, in a style of applicative programming, and
3. **class** modules defining object classes and associated procedures.

Any module may be parameterized by types and may import other modules. The style of constructive specifications forms an applicative programming language suited for mechanized reasoning aids through the use of “generators” for the definition of data types and “generator induction” in function definitions.

Class modules introduce an imperative and object oriented programming style, including implicit object pointers. The logical complications caused by data aliasing through pointers may be prevented using a restricted programming style which is enforceable by syntactic checks. Then the pointers are entirely transparent, and for the purpose of reasoning there is a simple correspondence between classes and certain types. For a given class C the corresponding type, say T_C , may be obtained by a syntactic transformation replacing updating procedure attributes by their “effect functions” expressing new object values. Thus, an update of a C -object (pointed to by the variable X , say **call** $X.p(\dots)$), may be thought of as a value assignment, $X := f_p(X, \dots)$, where X is seen as a variable of the type T_C , and f_p is the effect function of the procedure p .

It is sometimes necessary, for efficiency reasons, to generate object list structures with converging pointers. If the ensuing alias problems can be confined to within some object class, then the logic of the using program will remain simple; the internal logic of the class, however, becomes more complicated and so does its relationship to a corresponding type.

2 Data types by generator functions

The idea is to define the value set associated with a type implicitly, by listing a set of functions, the *generators*, in terms of which all such values can be expressed.

This basic idea occurs in some form in several languages for abstract specification as well as programming, notably the LARCH Shared Language (J. Guttag et al, [4]), the IOTA system (R. Nakajima et al, [6]) OBJ versions (J. Goguen et al, [3]) and dialects of ML (R. Milner [7], G.Huet and others [5]).

For typing purposes functions are characterized syntactically by *profiles* specifying their domain and codomain:

$$\mathbf{func} \ f : T_1 \times T_2 \times \dots \times T_n \longrightarrow T$$

The *generator basis* for a type T is a list of functions with the codomain T , the chosen T -generators.

$$\mathbf{genbas} \ g_1, g_2, \dots, g_m$$

If types other than T occur in the generator domains (possibly identified through a type parameter) they should be previously defined. They are called *underlying types*.

Neg1 representing the set of strictly negative ones. Let *Nat* correspond to the non-negative integers, $Nat = Zero \cup Nat1$. Then a one-to-one basis may be defined by restricting the domain of the successor function S^\wedge to *Nat*, and using a negation function N^\wedge with the domain *Nat1*. The codomain of each generator then will be one of the three “basic” subtypes.

```

type Int by Zero, Neg1, Neg1 with  $Nat = Zero \cup Nat1$  :
    func  $0 : \longrightarrow Zero$ ,  $S^\wedge : Nat \longrightarrow Nat1$ ,  $N^\wedge : Nat1 \longrightarrow Neg1$ 
    1-1 genbas  $0, S^\wedge, N^\wedge$ 
giving  $GU_{Int} = \{0; S0, SS0, SSS0, \dots; NS0, NSS0, NSSS0, \dots\}$ 

```

The generator universe consists of those basic expressions which are consistent with strong typing. They are clearly in a one-to-one correspondence with the integers.

In addition to *Nat*, another two “intermediate” syntactic subtypes of *Int* may be defined: $Neg = Zero \cup Neg1$ and $Nzo = Neg1 \cup Nat1$. They are useful for characterizing functions partial on *Int*, but total on a restricted domain. For instance, the division operator \wedge/\wedge is total on the domain $Int \times Nzo$.

Whereas syntactic subtypes, together with their main type, form a hierarchy of types defined simultaneously by a single generator basis, a *semantic subtype* is formed by explicitly restricting a given type. For instance, the set of integers implemented on a given computer might be the semantic subtype

```

type ImpInt ==  $x : Int$  where  $-32768 \leq x \leq 32767$ 

```

It is possible to simulate an one-to-one generator basis for $Set\{T\}$ in a semantic subtype consisting of unique “canonic forms”, e.g. basic expressions of the form

$$add(\dots add(add(\emptyset, t_1), t_2), \dots, t_n), \text{ where } n \geq 0 \text{ and } t_1 < t_2 < \dots < t_n,$$

assuming that $\wedge < \wedge$ is a total order on the otherwise unspecified type *T*. (Such an assumption may be expressed referring to a suitable **property** module parameterized by *T*.) In the subtype the generator *add* will be replaced by a defined function generating canonic forms. Since all set values are then canonic, equality on sets is reduced to syntactic equality (up to *T*-equality) on the representations.

4 Function definition by generator induction

The subterm relation provides a well-founded partial order on any generator universe: *e* is “less” than $g(\dots, e, \dots)$, for basic *T*-expressions *e* and $g(\dots, e, \dots)$. This order gives rise to an induction principle, called *generator induction*, used to cover the entire generator universe. This principle is useful for the purpose of defining functions, as well as for proving theorems.

In a function definition

```

def  $f(x_1, x_2, \dots, x_n) == \text{expression in } x_1, x_2, \dots, x_n$ 

```

generator induction with respect to an argument of type *T* may be expressed by a “**case-construct**” branching on the value of that argument, with one branch for each *T*-generator. If all the generators are constants, the construct is analogous to that of the Pascal language; for instance, branching on the value of an argument of type *Bool* gives one branch for the value **f** and one for **t**. A value of type *Nat*, defined as above, must either be 0 or it must be of the form **Se** where *e* is a *Nat* value. A corresponding **case-construct** must be a pattern matching mechanism. The “discriminator” heading the second branch will introduce a new variable for the argument of S^\wedge .

An alternative to a function definition with a **case-construct** in the right hand side, branching on a left hand argument, is to give one defining equation for each branch

with the discriminator thrown back into the left hand side. Such equations are referred to as **case-free** (inductive) axioms.

case-constructs may be nested to any depth, where inner ones may branch on left hand variables or those introduced in outer discriminators. (Branching on expressions other than variables corresponds to equational axioms guarded by explicit conditions.)

Examples

```

func ^ ^ ^ : Bool × Bool → Bool
def  x ^ y == case y of f → f | t → x fo
or:   x ^ f == f,   x ^ t == x

func ^ + ^ : Nat × Nat → Nat
def  x + y == case y of 0 → x | Sy' → S(x+y') fo
or:   A1: x+0 == x,  A2: x+Sy == S(x+y)

func ^ < ^ : Nat × Nat → Bool
def  x < y == case y of 0 → f | Sy' →
                                     case x of 0 → t | Sx' → x' < y' fo fo
or:   L1: x < 0 == f,  L2: 0 < Sy == t,  L3: Sx < Sy == x < y

```

Notice that, using generator induction and recursion when necessary, we are able to define these functions on the basis of generators alone. It is easy to see that the recursion in the axioms A2 and L3 must terminate, for in either case there is an argument which is textually simpler in the recursive application than in the left hand side. Such recursion is said to be *guarded by induction*. (Partial functions may be defined without losing the termination property by introducing a symbol, \perp , standing for an ill-defined expression of arbitrary type.)

5 Term Rewriting

The **case-free** axioms of a function definition may be taken as rules for term rewriting, read from left to right. They may be applied to arbitrary well-defined expressions without changing their meaning. Provided that any recursion is guarded by induction, these rules form a so called *convergent* system, which means that the rewriting process must necessarily terminate with a unique final result independent of the rewriting order.

If the given initial expression is variable-free, the final result is a basic expression, i.e. *the value* of the given one:

$$\begin{aligned} \text{SS0} + \text{SS0} &\xrightarrow{\text{A2}} \text{S}(\text{SS0} + \text{S0}) \xrightarrow{\text{A2}} \text{SS}(\text{SS0} + 0) \xrightarrow{\text{A1}} \text{SSSS0} \\ \text{SS0} < \text{SSS0} &\xrightarrow{\text{L3}} \text{S0} < \text{SS0} \xrightarrow{\text{L3}} 0 < \text{S0} \xrightarrow{\text{L2}} \text{t} \end{aligned}$$

If the given expression contains variables the rewriting process will often amount to a *simplification*:

$$\text{Sx} < \text{y} + \text{S0} \xrightarrow{\text{A2}} \text{Sx} < \text{S}(y+0) \xrightarrow{\text{A1}} \text{Sx} < \text{Sy} \xrightarrow{\text{L3}} x < y$$

If a Boolean expression can be rewritten to **t**, it is thereby proved as a theorem. It is sometimes possible to strengthen a convergent rewriting system, by including well chosen lemmas as additional rewrite rules, without losing convergence. In general, however, term rewriting is not sufficient for proof purposes; one needs the additional power provided by induction principles.

Proof by generator induction combines very nicely with generator inductive definition techniques. The structure of a generator inductive proof is determined by the relevant generator basis: there must be one separate proof for each generator, in which the theorem is modified by replacing the induction variable by the generator applied to

fresh variables. Induction hypotheses are allowed for those generator arguments which are of the same type as the induction variable (or of a subtype). In typical cases most of the work can be done by term rewriting based on (inductive) function definitions.

Example

The expression $x < \mathbf{S}x$ is irreducible by the inductive axioms above. In order to prove that it is true for all *Nat*-values x we may use generator induction on x . Given a generator basis for *Nat* consisting of 0 and \mathbf{S}^{\wedge} , the structure of the proof is similar to that of ordinary mathematical induction:

$$\begin{array}{l} \underline{0} : 0 < \mathbf{S}0 \xrightarrow{L2} \mathbf{t} \\ \underline{\mathbf{S}y} : \text{Assuming IH: } y < \mathbf{S}y, \text{ prove } \mathbf{S}y < \mathbf{S}\mathbf{S}y. \\ \mathbf{S}y < \mathbf{S}\mathbf{S}y \xrightarrow{L3} y < \mathbf{S}y \xrightarrow{\text{IH}} \mathbf{t} \end{array}$$

In summary, generator inductive definition techniques provide a very general framework for concept modelling. Term rewriting based on such definitions, combined with generator inductive proof techniques, provide powerful and efficient means of semi-mechanical reasoning about properties of functions and of programs using them.

6 Type modules

The ideas of the last sections may be seen as techniques of abstract specification, but they can also be taken as an *applicative programming language*. Once a representation of expressions has been chosen, convenient for internal machine manipulation, say list structures, a language interpreter could essentially take the form of a term rewriting algorithm. Another possibility is to view function definitions as algorithms working on values which are basic expressions.

Type modules in ABEL are in some ways comparable to Simula classes. Values of the type defined correspond to class objects, and the functions introduced in the module (generators as well as defined functions) correspond to procedure attributes. They might even be made to look “local to” the individual values if some special “mixfix” syntax, say dot notation, is used to identify a distinguished argument of the type in question (if there is such an argument).

Embedding such user defined types in an imperative program, they may be used for the typing of updatable program variables. The assignment operations should have the usual meaning of value copying.

Example

```

type Stack{T} by Estack, Stack1 ==
module
func nil :  $\rightarrow$  Estack, push : Stack  $\times$  T  $\rightarrow$  Stack1
1-1 genbas nil, push
func top : Stack1  $\rightarrow$  T      def top(push(S, x)) == x
func pop : Stack1  $\rightarrow$  Stack def pop(push(S, x)) == S
endmodule

var S : Stack{Int} = push(push(nil, 3), 5);
var T := S; T := push(T, 17); T := pop(T); T := pop(T);

```

Figure 1 shows the data configuration after variable initialization and after the second update (unbroken lines), respectively after the first and third updates (dashed lines). Notice that the updating of T does not change the value of S , although these variables share data structure. Thus, implementing value assignments and parameter

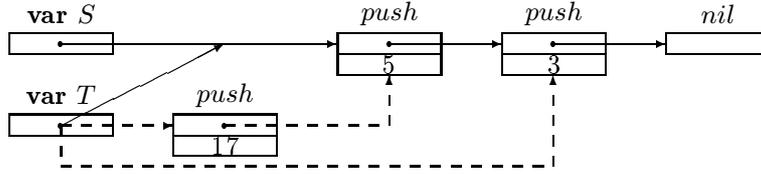


Figure 1: Pushing and popping stacks.

transmission by pointer copying does not harm the program logic. This is true as long as values are represented by *immutable* data structures. The list structures directly represent basic $Stack\{Int\}$ expressions (with an ad-hoc representation of integers, rather than a list structure representing $\mathbf{SSS0}$ for 3, and so forth).

It should be noticed that the functions *pop* and *top* are total functions on the domain of non-empty stacks, but are undefined for empty stacks. This is the reason why the syntactic subtypes *Estack* and *Stack1* have been introduced. (*Stack1* has only one generator, therefore one **case-free** inductive axiom is sufficient for either function.) Whereas values of a subtype are also values of its supertype, the reverse is not generally true. However, it may be practical to include a mechanism for “coercion” to subtypes. For instance, behind the scenes application of the function

```

func StackToStack1 : Stack  $\rightarrow$  Stack1
def StackToStack1(S) == case S of nil  $\rightarrow \perp$  | others  $\rightarrow S$  fo

```

would make it legal to apply *pop* and *top* to values of type *Stack*, i.e. to stacks not syntactically identifiable as being non-empty.

Unfortunately the requirement that values be immutable in general leads to gross inefficiency when manipulating high volume data, having to recreate large objects in order to make small local changes. In fact, the efficiency of the operations of pushing and popping stacks is atypical. A more normal situation occurs in the following example of adding new elements at the far end of a stack.

```

func enq : Stack  $\times$  T  $\rightarrow$  Stack1
def enq(S, x) == case S of nil  $\rightarrow$  push(S, x)
                | push(S', y)  $\rightarrow$  push(enq(S', x), y) fo

```

```

var S : Stack{Int} = push(push(nil, 3), 5);
var T := S; T := enq(T, 17);

```

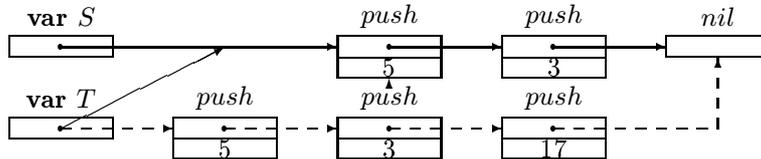


Figure 2: Enqueueing.

Notice that each *push* operation, invoked recursively in *enq*, must generate a new record in the list structure. As a result the whole stack is copied.

To summarize, values are timeless concepts, like the number 3, and must be represented in a computer by immutable data structures. Values of high volume, such as large arrays or voluminous list structures, therefore must be updated by (more or less) complete recreation, which may entail much inefficiency. On the other hand, whereas

assignment and parameter transmission of values conceptually are by copying, they may be realized efficiently through pointer copying at the discretion of a compiler without harming the program logic.

7 Object classes

Efficiency considerations show that there is a need for locally updatable data structures accessed through pointers. By introducing a local update mechanism for structures defined inductively as above, we obtain a kind of simplified and generalized Simula objects. We choose to let the pointers remain implicit in order to emphasize the similarity between values and objects. Thus, we may define a class of objects in the same way as a value type, simply by replacing the keyword **type** by the keyword **class**.

A local update of an object is achieved by updating a component of the object. Thus, for an object variable X the component names introduced in **case** discriminators for X should count as updatable program variables:

```
case X of ... | g(..., y, ...) → y := ... | ... fo
```

In order to protect the integrity of class objects it is reasonable to allow local updates only textually inside the class in question (and in subclasses).

class modules may have the same kind of contents as **type** modules, but they may in addition contain *procedure declarations* with imperative bodies, parameterized by values (**val**) and/or variables (**var**). Thus, in order to define a class of “stack” objects with a more efficient mechanism for entering elements at the far end, we may copy the module part of figure 1 and add a procedure for the “enqueueing” of elements. However, rather than copying all that program text we may introduce a Simula inspired prefixing mechanism and define a “subclass” of the stack type of figure 1.

```
class STACK{T} by ESTACK, STACK1 ==
Stack{T} module
proc Enq(var S : STACK, val x : T) ==
case S of nil → S := push(S, x) | push(S', y) → call Enq(S', x) fo
endmodule
```

```
var S : STACK{Int} = push(push(nil, 3), 5);
var T := S; call Enq(T, 17);
```

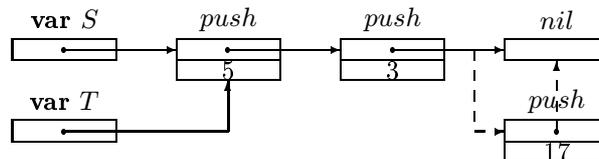


Figure 3: More efficient enqueueing.

It is still necessary to search for the far end of the stack, but at least we do not have to copy any part of the old stack in order to insert the new element. However, as a result of updating *inside* the stack object the program logic suffers: changing T has a side effect on S . There is nothing at all in the statement **call Enq(T, 17)** to indicate that the object called S should be modified, but that is what happens! The reason is that the preceding statement has made S and T become *aliases*.

In general, simultaneously accessible converging pointers in data structures represent aliases and provide a danger of side effects in component assignments. In the example the alias problem was easy to identify, but in general it may be *very* difficult, especially when reasoning about incorrect programs, as we sometimes have to do.

Fortunately, there exist styles of programming which will guarantee the absence of aliasing, and which may be enforced by syntactic checks on the program text. For instance, the following requirement on object assignments is sufficient:

- The right hand side must contain no object variable outside value subterms, except at most a single occurrence of the left hand variable or a component of that object. (In this context variable subscripts must be ignored, so that for instance $A[i]$ and $A[j]$, where A is an object array, are regarded as the same variable, which they are if $i=j$.)

This restriction is too hard to live with, but it may be relaxed in several ways. One important idea is to introduce multiple assignments in the language. Each object variable of the left hand list (or a component of that object) is then allowed to occur once in the right hand list. So, for instance the statement $X, Y := Y, X$ would be allowed. Obviously, it cannot cause aliasing.

Alias-free programming with object classes is a way to combine (up to a point) different important considerations which are otherwise in conflict:

- simplicity of the program logic, in particular with respect to assignment operations, and consequent ease of reasoning about programs,
- the flexibility and computational efficiency that may be obtained using updatable data structures accessed and related through assignable pointers.

8 Simula-like classes

The class/object mechanisms as exemplified so far may look more like a list processing facility than object orientation in the traditional sense. We now show how a little special purpose syntax may close that gap. We first introduce the concept of “labelled Cartesian products”. A type definition of the form

type $T == \{a_1:T_1 \times a_2:T_2 \times \dots \times a_n:T_n\}$

is shorthand for the following type module (stretching the ABEL syntax a little bit):

```
type  $T ==$ 
module
func  $(\hat{\ }, \hat{\ }, \dots, \hat{\ }) : T_1 \times T_2 \times \dots \times T_n \longrightarrow T$            ( $n$ -tuple)
1-1 genbas  $(\hat{\ }, \hat{\ }, \dots, \hat{\ })$ 
func  $\hat{\ }.a_1 : T \longrightarrow T_1, \dots, \hat{\ }.a_n : T \longrightarrow T_n$        (component selectors)
def  $(x_1, x_2, \dots, x_n).a_1 == x_1 \dots \dots \dots$  def  $(x_1, x_2, \dots, x_n).a_n == x_n$ 
endmodule
```

So, if t is a value of type $\{a_1:T_1 \times a_2:T_2 \times \dots \times a_n:T_n\}$, the expression $t.a_k$, where k is an integer in the range $1..n$, evaluates to the k 'th component of t .

Now (more) Simula-like classes, say a class C , may be defined as a subclass of a labelled Cartesian product, where the labels are made to play the roles of updatable variables within the module body, as follows. Any function or procedure whose name is of the form $.name$ has an implied class C parameter in front of the dot in the form of the **case** discriminator (a_1, a_2, \dots, a_n) (unless an explicit actual parameter is provided). For a procedure the implied formal parameter of the left hand side is considered a variable, which implies that each a_k is updatable in the procedure body.

Let X be an object of class C . It follows from the above trickery that any invocation $X.p(\dots)$ of a procedure or function $.p$, whose body contains (free) occurrences of a_1, a_2, \dots, a_n , will access the components of X . Notice that $X.a_k$ is an application of a selector function, not an updatable variable. It follows that any invariant established initially and maintained by the C -procedures is guaranteed to hold for all C -objects. The invariant may occur as a restricting predicate, defining a semantic subtype of the labelled Cartesian product which is the prefix of the C -module.

No predefined “reference-to-nothing” is needed in the language, provided that explicit initialization of (object) variable declarations is possible, and preferably mandatory. On the other hand, any generator basis must contain at least one relative constant, i.e. a function with no occurrence in its domain of the type(s) being defined, in order to terminate data recursion. For instance, the *Stack* generator *nil* plays the role of a recursion terminator. For a non-recursive Cartesian product the tuple generator is already a relative constant.

9 Aliasing for efficiency

It happens that data structures with converging pointers are necessary in order to obtain good efficiency. For instance, in order to program the enqueueing and popping operations of a FIFO list in constant time, pointers to both ends of the list are required. Consequently there will be pointer convergence at the “far” end of the list and a corresponding alias problem. We show an efficient version of a FIFO queue based on the *STACK* class of section 7, as a Simula-like class.

```

class FIFOQ{T} == {front:STACK{T} × rear:STACK{T}}
                    where last(front) = rear

module
func empq : → FIFOQ      def empq == (nil, nil)
proc .Enq(val x) == var new:STACK = push(nil, x);
                    case rear of nil → front, rear := new, new
                    | push(S, y) → S, rear := new, new fo
proc .Pop == case front of nil → abort
                | push(S, x) → front := S; if front = nil th rear := nil fi fo
endmodule

var Q : FIFOQ{Int} = empq; call Q.Enq(5); call Q.Enq(3):

```

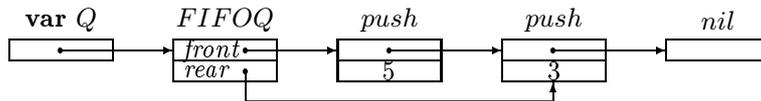


Figure 4: Efficient FIFO queue.

It is assumed that the *STACK* module has been extended by the function *last* in order that its occurrence in the object invariant be meaningful:

```

func last : STACK → STACK
def last(S) == case S of nil → S | push(S', x) →
                case S' of nil → S | others → last(S') fo

```

Reasoning about assignments in the context of pointer aliases is not at all trivial. For instance, *STACK* pointers must be made explicit, which implies that *STACK* generators now have observable side effects. Furthermore, it must be proved that the occurring pointer assignments do not create pointer circles, in which case the *last* function might be ill-defined due to non-terminating recursion, which would make the object invariant ill-defined as well. (Actually, the invariant has to be strengthened in order to be provable.) The “correctness” of *FIFOQ* may nonetheless be defined and proved by comparison with a corresponding type (an extended *Stack* as in section 6 would do), expressing the purpose of *FIFOQ* in a way suitable for easy reasoning while disregarding computational efficiency.

Fortunately, having verified the correctness of *FIFOQ* the users will not be troubled by the internal aliasing; in fact, he may reason as if the *FIFOQ* operations were replaced by corresponding *Stack* operations. Thus, a statement `call Q.Enq(x)`, where `var Q: FIFOQ`, would correspond to $Q := \text{enq}(Q, x)$, interpreting Q as being of type *Stack*.

10 Conclusions

Elements of a language for specification and programming have been explained and exemplified. The language mechanisms are drawn from a language ABEL developed at Oslo University. The techniques of generator inductive type and function definition provides an easy to use style of applicative programming, and at the same time lends itself to efficient reasoning aids based on term rewriting and generator inductive proofs. **type** modules enable natural grouping of functions, and permit associated rules of function overloading (not discussed here), as well as subtyping with inheritance.

It is shown how efficiency in time and space may be controlled using suitable imperative language facilities. A **class** construct serves to introduce locally updatable data objects manipulated through (implicit) pointers, and to keep imperative mechanisms well separated from applicative ones. The logical difficulties caused by pointer aliasing may be avoided entirely using a restricted form of programming (allowing efficiency up to a point), or may be confined to within given class bodies.

A general advice to programmers, borne out by experience, is to put as much of the programming details as feasible into the applicative part, i.e. into types and functions. It is important to keep the applicative part “clean”, avoiding functions with side effects, in order to enhance perspicuity and easy reasoning. Properly restricted object classes introduced for efficiency reasons may be viewed as implementations of corresponding types, so that the using program is logically equivalent to one which is more applicative.

11 Acknowledgements

Bjørn Kristoffersen and Olaf Owe have contributed to the readability of this paper by providing non-trivial feedback.

References

- [1] O. Owe, O.-J. Dahl: Formal Development with ABEL. VDM'91, *Springer Lecture Notes 552*, pp. 320-362. Revised version: Research Report 159, Dept. of Informatics, University of Oslo.
- [2] O.-J. Dahl: *Verifiable Programming*. Prentice Hall, 1992.
- [3] K. Futasugi, J.A. Goguen, J.-P. Jouannaud, J. Meseguer: Principles of OBJ2. Proceedings, 1985 ACM Symposium on Principles of Programming Languages and Programming, W. Brauer, Ed., *Springer Lecture Notes in Computer Science, Volume 194*, pp. 52-66.
- [4] J.V. Guttag, J.J. Horning, J.M. Wing: *Larch in Five Easy Pieces*. Digital Systems Research Center, Palo Alto, California, July 1985.
- [5] G. Huet, ed.: *Logical Foundations of Functional Programming*, Addison Wesley 1990.
- [6] R. Nakajima, M. Honda, H. Nakahara: Describing and Verifying Programs with Abstract Data Types. In Neuhold ed: *Formal Description of Programming Concepts*, North Holland 1978.
- [7] A.J.R.G. Milner: A Proposal for Standard ML. *Proceedings of the ACM Symposium on LISP and Functional Programming*, Austin, 1984.