

Abstract Datatypes, Multiple Views and Multiple Inheritance in Scratchpad II

Extended Abstract

Stephen M. Watt 8-862-3405 (SMWATT at YKTVMZ)

Richard D. Jenks 8-862-1233 (JENKS at YKTVMZ)

Computer Algebra Group, Mathematical Sciences Department
IBM Thomas J. Watson Research Center
P.O. Box 218, Yorktown Heights, New York 10598

Scratchpad II is an abstract datatype language developed at Yorktown Heights for the implementation of a new computer algebra system. It provides packages of polymorphic functions and parameterized, abstract datatypes with operator overloading and multiple inheritance. To express the intricate inter-relationships between the datatypes necessary for the description of mathematical objects, a number of techniques based on the notion of *category* have been used. Categories are used to enforce relationships between type parameters and to provide the mechanism for multiple inheritance. They also allow the language to be statically type checked and the generation of efficient code. This paper describes the role of categories in Scratchpad II.

Domains

In an environment where each operation may have many meanings, it is often useful to give a "domain of computation" to specify the operations of interest. A *domain* is a Scratchpad II object which provides a set of operations and attributes. Domains are used to implement abstract datatypes and packages.

In Scratchpad II every object belongs to a unique domain. For example, 2 belongs to the domain Integer. Domains are first class run-time values which belong to the domain Domain. Objects belonging to a domain may be manipulated by operations it exports. For example, two of the operations Integer provides are

```
"+": (Integer, Integer) -> Integer  
"=": (Integer, Integer) -> Boolean
```

The operations exported by a domain need not combine only members of the domain itself. The signatures of the exported operations may involve any types whatsoever. For example, the domain RationalNumber exports the operations

```
"/": (Integer, Integer) -> RationalNumber  
characteristic: () -> NonNegativeInteger
```

Note that RationalNumber does not appear in the signature for characteristic. A domain which does not export any operations for creating or manipulating members is a *package*.

Domains may be created by functions, providing parameterized types and packages of polymorphic functions. As an example of a parameterized type, when Integer is passed to the function Stack a domain which we denote Stack(Integer) is returned.

The need for polymorphic functions stems from the desire to implement a given algorithm only once, and to be able to use the program for any values for which it makes sense. For example, the Euclidean algorithm for computing greatest common divisors can be used for values belonging to any type which is a Euclidean domain. The following package takes a Euclidean domain as a type parameter and exports the operations gcd and lcm on that type.

```

GCDpackage(R: EuclideanDomain): with
  gcd: (R, R) -> R
  lcm: (R, R) -> R
  == add
  gcd(x,y) == -- Euclidean algorithm
    x:= unitNormal.x.coef
    y:= unitNormal.y.coef
    while y != 0 repeat
      (x,y):= (y,x rem y)
      y:= unitNormal.y.coef
    x
  lcm(x, y) ==
    u: Union(R, "failed") := y exquo gcd(x,y)
    x * u::R

```

The exported operations can equally well be used for many types, the integers or polynomials over GF(7) being two examples. Although the same gcd program is used in both cases, the operations it uses (rem, unitNormal, etc.) come from the type parameter R.

Subdomains

An object may belong to any number of subdomains. The number 2, for example, belongs to many subdomains of Integer, including PositiveInteger, EvenInteger and SmallInteger. The domain Integer belongs to many subdomains of Domain, including Monoid, AbelianGroup, Ring and Algebra(Integer).

A *subdomain* consists of

- a domain
- a boolean function that characterizes which members of the domain belong to the subdomain
- additional operations defined on the subdomain.

A variable may be declared to belong to a domain or to a subdomain. When a variable is declared to belong to a subdomain, its domain is that which the subdomain restricts. That is, values lying in a subdomain also lie in the domain and may be used in all the appropriate domain operations.

The subdomain may provide operations which supplement or supersede those of the domain but which are restricted to values lying in the subdomain. This restriction is for two reasons. First, operations which are closed on the domain may not be closed on the subdomain. Second, if a value has been determined to lie in a subdomain or is the result of a closed subdomain operation, then the subdomain may provide a more efficient implementation than the domain operation.

Often the subdomain predicate can be determined at compile-time. In places where this cannot be done, a run-time check may be inserted when a value must belong to a subdomain. Certain subdomains which are known to the compiler can be more highly optimised than others.

Categories

While polymorphic packages allow the implementation of algorithms in a general way, it is necessary to ensure that these algorithms may *only* be used in meaningful contexts. It would *not* be meaningful to try to use GCDpackage above with Stack(Integer) as the parameter. In order to restrict the use to cases where it makes sense Scratchpad II has the notion of categories.

A *category* in Scratchpad II is a restriction on the class of all domains. It specifies what operations a domain must support and certain properties the operations must satisfy. A category may be created using a category constructor such as the one below.

```

OrderedSet(): Category == Set with
  -- operations
  "<": ($,$) -> Boolean
  max: ($,$) -> $
  min: ($,$) -> $
  -- attributes
  irreflexive "<" -- not (x < x)
  transitive "<" -- x < y and y < z implies x < z
  total "<" -- not(x < y) and not(y < x) implies x=y

```

OrderedSet gives a category which extends the category Set by requiring three additional operations and three properties, or *attributes*.

A declaration is necessary in order for a domain to belong to a given category: having the necessary operations and attributes does not suffice. This is because the attributes are not intended to give a complete set of axioms, but merely to make explicit certain facts that may be queried later on. It is usually the case that belonging to a category implies that a domain must satisfy conditions that are not mentioned as attributes. For example, in OrderedSet there is no attribute relating min and "<", although such relations are implied.

A type parameter to a domain or package constructor is usually required to belong to an appropriate category. For example, in the previous section the parameter R to GCDpackage was declared to belong to the category EuclideanDomain.

Separation of Contract and Implementation

The use of categories in restricting the type parameters to a domain or package constructor allows algorithms to be specified in very general contexts. For example, since all table types belong to a common category, algorithms can be written that do not need to know the actual implementation (for example, whether it is a hash table in memory or a file on disk). As an example of algebraic generality, consider the domain of linear ordinary differential operators, which is declared as follows

```
LinearOrdinaryDifferentialOperator(A, M): T == B where
  A: DifferentialRing
  M: Module(A) with deriv: $ -> $

  T == GeneralPolynomialWithoutCommutativity(A, NonNegativeInteger) with
    D:      () -> $
    ".":    ($, M) -> M
  ...
  ...
```

This domain defines a ring of differential operators which act upon an A-module, where A is a differential ring. The type of the coefficients, A, is declared to belong to the category DifferentialRing and type of the operands, M, is declared to belong to the category Module(A) with a derivative operation. The constructed domain of operators is declared to belong to a category of general polynomials with coefficients A and two additional operations. The operation D creates a differential operator and "." provides the method of applying operators to elements of M.

Multiple Views and Multiple Inheritance

It is often necessary to view a given domain as belonging to different categories at different times. Sometimes we want to think of the domain Integer as belonging to Ring, sometimes as belonging to OrderedSet, and at other times as belonging to other categories. For a domain to have multiple views, it should be declared to belong to the Join of the appropriate categories. For example, the following keyed access file datatype may be viewed either as a table or as a file:

```
KeyedAccessFile(Entry: Set): T == B where
  FileRec ==> Record(key: String, entry: Entry)
  ErrorMsg ==> String

  T == Join(FileCategory(LibraryName, FileRec),
            TableCategory(String, Entry, ErrorMsg)) _
    with
      pack: $ -> $

  B == add ...
```

An important use of categories is to supply default implementations of operations. So long as certain primitive operations are provided by a domain, others can be implemented categorically. For example, supplying only "<" allows definitions of ">", "<=" and ">=". Thus a domain may inherit operations from a category. The use of Join provides multiple inheritance.

Compile-Time Binding

A domain object contains a number of vectors of function/environment pairs. When operation bindings are not known at compile-time (as for the operations exported by a type parameter), the operations are performed by calling the function in the appropriate slot of a vector. When operation bindings are known at compile time, a code fragment for the operation may be placed in line.

The scope rules in Scratchpad II and the operations applicable to domain values have been designed so that when a domain is known to belong to a particular category, it is also known exactly what operations it exports. From this, the precise location of each function is determined. Thus when a function is called using the general mechanism, a hard-coded offset is used.

Scratchpad II is implemented on top of Lisp/VM. Since the exact number and type of arguments are known at compile time, much of the usual function can be omitted. This, in conjunction with compile-time knowledge of function offsets, makes function calling in Scratchpad II faster than that of the underlying Lisp system.

Bibliography

- [1] Jenks, R. D. and Trager, B. M., "A Language for Computational Algebra," *Proceedings of SYMSAC '81, 1981 Symposium on Symbolic and Algebraic Manipulation*, Snowbird, Utah, August, 1981. Also *SIGPLAN Notices*, New York: Association for Computing Machinery, November 1981, and *IBM Research Report RC 8930* (Yorktown Heights, New York).
- [2] Jenks, R. D., "A Primer: 11 Keys to New Scratchpad," *Proceedings of EUROSAM '84, 1984 International Symposium on Symbolic and Algebraic Computation*, Cambridge, England, July 1984
- [3] Computer Algebra Group, *Basic Algebraic Facilities of the Scratchpad II Computer Algebra System*, Yorktown Heights, New York: IBM Corporation, March 1986.
- [4] Computer Algebra Group, *An Overview of the Scratchpad II Language and System*, Yorktown Heights, New York: IBM Corporation, April 1986.
- [5] Computer Algebra Group, *Scratchpad II Examples from INPUT Files*, Yorktown Heights, New York: IBM Corporation, August 1986.
- [6] Sutor, R. S., ed. *The Scratchpad II Newsletter*, Vol. 1, No. 1, Yorktown Heights, New York: IBM Corporation, September 1, 1985.
- [7] Sutor, R. S., ed. *The Scratchpad II Newsletter*, Vol. 1, No. 2, Yorktown Heights, New York: IBM Corporation, January 15, 1986.
- [8] Sutor, R. S., ed. *The Scratchpad II Newsletter*, Vol. 1, No. 3, Yorktown Heights, New York: IBM Corporation, May 15, 1986.