

# Specifying Reusable Components

Nadia Polikarpova, Carlo A. Furia, and Bertrand Meyer

Chair of Software Engineering, ETH Zurich, Switzerland

{nadia.polikarpova, carlo.furia, bertrand.meyer}@inf.ethz.ch

**Abstract.** Reusable software components need well-defined interfaces, rigorously and completely documented features, and a design amenable both to reuse and to formal verification; all these requirements call for expressive specifications. This paper outlines a rigorous foundation to *model-based contracts*, a methodology to equip classes with expressive contracts supporting the accurate design, implementation, and formal verification of reusable components. Model-based contracts conservatively extend the classic Design by Contract by means of expressive models based on mathematical notions, which underpin the precise definitions of notions such as abstract equivalence and specification completeness. Preliminary experiments applying model-based contracts to libraries of data structures demonstrate the versatility of the methodology and suggest that it can introduce rigorous notions, but still intuitive and natural to use in practice.

## 1 Introduction

Software specifications have many uses. The most widely recognized one is verification: comparing two independent descriptions of a software module — a specification and an implementation — is likely to reveal errors, because the probability of making the same error in both is considered low. Besides, the level of abstraction of the specification is usually higher than that of the implementation, which makes the latter a more direct representation of the programmer’s informal intent and further decreases the probability of an error.

Another use, especially important for reusable software components, is providing a precise interface for a client. Without specifications client modules can rely on type information and informal documentation (comments), but in essence they have no guarantees about the functional properties of the component.

This is even more critical in presence of inheritance in object-oriented context. If a client uses a concrete class through an interface of its abstract superclass, the substitution principle [21] ensures that the behavior observed by the client is compatible with the specification he relies on. However, if the superclass has no or a very weak specification, a wide range of unpredictable behaviors are allowed; this makes such an abstract class essentially unusable from the client’s perspective and strongly depreciates the advantages of inheritance.

In spite of many benefits of formal specifications, only a little portion of existing software is actually specified. The situation is different in languages with support for Design by Contract (DbC), such as Eiffel [22], JML [19] and Spec# [3]. An extensive study [5] indicates that Eiffel classes in practice contain substantial amount of contracts, however, as shown in [8] these contracts are generally weak compared to the

programmer's informal understanding of the functional properties of a software module. In many cases the restricted specification language used in DbC does not allow a programmer to express (or at least express easily) all the desired properties.

One of the approaches to increasing the expressiveness of the contract language is using *model classes* — immutable classes designed for specification purposes and representing standard mathematical notions, such as sets, relations or sequences. This approach is used in JML [7, 20] and Eiffel [28]. The availability of model classes gives programmers the required expressive power, however it does not by itself guarantee strong specifications.

Current paper presents *model-based contracts* — an interface specification method based on Design by Contract and model classes. The method provides systematic guidelines for writing strong specifications and a precise definition of specification completeness, which is easy to reason about.

We performed experimental evaluation of the approach on two Eiffel data structure libraries: EiffelBase and its successor EiffelBase2. The former is a standard library used for many years in production software, while the latter in a research project intended as a testbed for specification and verification techniques, but still providing the full functionality of a standard data structures library.

Section 2 motivates the need for more expressive contracts and briefly demonstrates the advantages of using them on a few concrete examples. Section 3 provides details of the specification method, discusses the semantics of model-based contracts with respect to proofs and tests and introduces the notion of specification completeness. Section 4 demonstrates the applicability of the approach on two case studies, section 5 presents related work and section 6 concludes.

All the examples in the paper are from the EiffelBase library if not specified otherwise.

## 2 Motivation and overview

Design by Contract (DbC) is a discipline of analysis, design, implementation, and management of software. It relies on the fundamental idea of defining the role of any component in the system in terms of a *contract* that formalizes the obligations and benefits of that component relative to the rest of the system. Concretely, the contract is as a collection of assertions (*preconditions*, *postconditions*, and *invariants*) that constitute the module's *specification*.

### 2.1 Some limitations of Design by Contract

To emphasize the seamless connection that must exist between specification and implementation, and to make writing contracts palatable to the programmer, DbC uses the same notation for expressions in the implementation and in the specification. This choice successfully encourages programmers to write contracts [5]. On the other hand, it also restricts the assertions that can be expressed — or that can be expressed easily. This restriction ultimately impedes the formalization and verification of full functional correctness and even limits the scope of application of DbC for the correct design of an

```

2 class LINKED_LIST [G]
3   count: INTEGER -- Number of elements
4
5   index: INTEGER -- Current cursor position
6
7   put_right (v: G)
8     -- Add 'v' to the right of cursor.
9     require 0 ≤ index ≤ count
10    do ...
11  ensure
12    count = old count + 1
13    index = old index
14  end
16  duplicate (n: INTEGER): LINKED_LIST
17    -- Copy of sublist of length 'n' beginning at current position
18    require n ≥ 0 do ... ensure Result.index = 0 end
19  end
21 class TABLE [G, K]
22  put (v: G ; k: K)
23    -- Associate value 'v' with key 'k'.
24    require valid_key (k)
25    deferred end
26
27 end

```

**Table 1.** Snippets from the EiffelBase classes `LINKED_LIST` (lines 1–17) and `TABLE` (lines 19–25).

implementation. Let us demonstrate this on a couple of examples from the EiffelBase library [12].

Lines 1–14 in Table 1 show a portion of class `LINKED_LIST`, implementing a dynamic list. Features (members) `count` and `index` record respectively the number of elements stored in the list and the current position of the internal cursor. Routine `put_right` inserts an element  $v$  to the right of the current position of the cursor, without moving it. The postcondition of the routine (clause `ensure`) asserts that inserting an element increments `count` by one but does not change `index`. This is correct, but it does not capture the gist of the semantics of insertion: the list after insertion is obtained by all the elements that were in the list up to position `index`, followed by element  $v$  and then by all elements that were to the right of `index`.

Expressing such complex facts is impossible or exceedingly complicated with the standard assertion language; as a result most specifications are *incomplete* in the sense that they fail to capture precisely the functional semantics of routines. Weak specifications hinder formal verification in two ways. First, establishing weak postconditions is simple, but confidence in the full functional correctness of a verified routine will be low: the quality of specifications limits the value of verification. Second, weak contracts affect negatively verification modularity: it is impossible to establish what a routine  $r$  achieves, if  $r$  calls another routine  $s$  whose contract is not strong enough to document its effect within  $r$  precisely.

Weak assertions limit the potential of many other applications of DbC. Specifications, for example, should document the abstract semantics of operations in deferred classes (classes without an implementation). Weak contracts cannot fully do so; as a result, programmers have fewer safeguards to prevent inconsistencies in the design and fewer chances to make deferred classes useful to clients through polymorphism and dynamic dispatching.

Feature `put` in class `TABLE` (lines 16–19 in Table 1) is an example of such a phenomenon. It is unclear how to express the abstract semantics of `put` with standard contracts. In particular, the absence of a postcondition leaves it undefined what should happen when an element is inserted with a key that is already associated to some other element: should `put` replace the previous element with the new one or cancel the insertion

```

2 note model: sequence, index
3 class LINKED_LIST [G]
4   sequence: MML_SEQUENCE [G]
5   -- Sequence of elements
6   do ... end
7
8   count: INTEGER -- Number of elements
9   ensure Result = sequence.count end
10
11  index: INTEGER -- Current cursor position
12
13  put_right (v: G)
14  -- Add 'v' to the right of cursor.
15  require 0 ≤ index ≤ count
16  do ...
17  ensure
18    sequence = old ( sequence.front (index)
19      .extended (v) + sequence.tail (index + 1) )
20    index = old index
21  end
22 end

```

```

24 note model: map
25 class TABLE [G, K]
26   map: MML_MAP [G, K]
27   -- Map of keys to values
28   deferred end
29
30   put (v: G ; k: K)
31   -- Associate value 'v' with key 'k'.
32   require map.domain [k]
33   deferred
34   ensure
35     map = old map.replaced_at (k, v)
36   end
37 end

```

**Table 2.** Classes `LINKED_LIST` (left) and `TABLE` (right) with model-based contracts.

of the new element? Indeed, some heirs of `TABLE` implement `put` with a replacement semantics (such as class `ARRAY`), while others disallow overriding of preexisting mappings with `put` (such as class `HASH_TABLE`). Some classes (including `HASH_TABLE`) even introduce another feature `force` that implements the replacement semantics. This obscures the behavior of routines to clients and makes it questionable whether `put` has been introduced at the right point in the inheritance hierarchy.

## 2.2 Enhancing Design by Contract with models

This paper presents an extension of DbC that addresses the aforementioned problems. The extension conservatively enhances DbC with *model classes*: immutable classes representing mathematical concepts that provide for more expressive specifications. Wrapping mathematical entities with classes supports richer contracts without need to extend the notation, which remains the one familiar to programmers as in DbC. Contracts using model classes are called *model-based contracts*.

Table 2 shows an extensions of the examples in Table 1 with model-based contracts. `LINKED_LIST` is augmented with a query `sequence` that returns an instance of class `MML_SEQUENCE`, a model class representing a mathematical sequence of elements of homogeneous type; the implementation, omitted for brevity, builds `sequence` according to the actual content of the list. The meta-annotation `note` declares the two features `sequence` and `index` as `model` of the class; every contract will rely on the abstraction they provide. In particular, the postcondition of `put_right` can precisely describe the effect of the routine: the new `sequence` is the concatenation of the `old` `sequence` up to `index`, extended with element `v`, with the tail of the `old` `sequence` starting after `index`. We can assert that the new postcondition — including the clause about `index` — is *complete* with respect to the model of the class, because it completely defines the effect of `put_right` on

the abstract model. This notion of completeness is a powerful guide to writing accurate specification that makes for well-defined interfaces and verifiable classes.

The mathematical notion of a *map* — encapsulated by the model class `MML_MAP` — is the natural model for the class `TABLE`. Feature *map* cannot have an implementation yet, because `TABLE` is deferred and hence it is not committed to any representation of data. Nonetheless, the mere availability of a model class supports complex specifications already at this abstract level. In particular, writing a complete postcondition for routine *put* requires to commit to a specific semantics for insertion. The example in Table 2 chooses the replacement semantics; correspondingly, all heirs of `TABLE` will have to conform to this semantics, guaranteeing a coherent reuse of `TABLE` throughout the class hierarchy.

### 3 Foundations of model-based contracts

#### 3.1 Specifying classes with models

This subsection describes a rigorous approach to equipping classes with expressive contracts.

**Interfaces, references, and objects.** The definitions of abstract objects and models (introduced in the remainder) rely on the following simple assumptions about classes. Each class  $C$  defines a notion of *reference equality*  $\equiv_C$  and of *object equality*  $\overset{\circ}{=}_C$ ; both are equivalence relations. Two objects  $o_1, o_2 \in C$  of class  $C$  can be *reference equal* (written  $o_1 \equiv_C o_2$ ) or *object equal* (written  $o_1 \overset{\circ}{=}_C o_2$ ). Reference equality is meant to capture whether  $o_1$  and  $o_2$  are aliases for the same physical object, whereas object equality is meant to hold for (possibly) physically distinct objects with the same actual content. The following discussion is however independent of the particular choice of reference and object equality.

The principle of information hiding prescribes that each class define an interface: the set of its publicly accessible features [22]. It is good practice to partition features into queries and commands; queries are functions of the object state, whereas commands modify the object state but do not return any value.  $I_C = Q_C \cup M_C$  denotes the interface of a class  $C$  partitioned in queries  $Q_C$  and commands  $M_C$ .<sup>1</sup> It is convenient to partition all queries into *value-bound* queries  $Q_C^o$  and *reference-bound* queries  $Q_C^r$ . Value-bound queries should create fresh objects to return (or more generally objects that were unknown to the client before calling the query), whereas reference-bound queries give the client direct access, through a reference, to parts of the target object or of the query arguments. In other words, clients of a value-bound query are insensitive to whether they received a unique fresh object or they are just sharing a reference to a previously existing one. The chosen partitioning between value-bound and reference bound queries does not affect the following discussion, although it is usually quite natural to adhere to this informal distinction when designing a class.

<sup>1</sup> Constructors need no special treatment and can be modeled as queries returning new objects.

*Example 1.* Query *item* (Table 3) is reference-bound, as the client receives the very same physical object that was earlier inserted in the list. Query *duplicate* (Table 3) is instead value-bound, as it returns a copy of a portion of the list.

The classification in value-bound and reference-bound extends naturally to *arguments* of features: if the feature does not rely on having a direct reference to the actual argument (as opposed to a copy of it), the argument is value-bound; otherwise, it is reference-bound.

**Abstract object space.** The interface  $I_C$  induces an equivalence relation  $\approx_C$  over objects of class  $C$  called *abstract equality* and defined as follows:  $o_1 \approx_C o_2$  holds for two objects  $o_1, o_2 \in C$  iff for any applicable sequence of calls to commands  $m_1, m_2, \dots \in M_C^*$  and a query  $q \in Q_C$ , the qualified calls  $o_1.m_1; o_1.m_2; \dots o_1.q$  and  $o_2.m_1; o_2.m_2; \dots o_2.q$  (with physically identical actual arguments where appropriate) respectively return objects  $t_1$  and  $t_2$  of some class  $T$  such that: if  $q$  is reference-bound then  $t_1 \equiv_T t_2$ , and if  $q$  is value-bound then  $t_1 \overset{\circ}{=}_T t_2$ . Intuitively, two objects are equivalent with respect to  $\approx_C$  if a client cannot distinguish them by any sequence of calls to public features.

Abstract equality defines an *abstract object space*: the quotient set  $A_C = C / \approx_C$  of  $C$  (as a set of objects) by  $\approx_C$ . As a consequence, two objects are equivalent w.r.t.  $\approx_C$  iff they have the same *abstract (object) state*. Any concrete set that is isomorphic to  $A_C$  is called a *model* of  $C$ .

*Example 2.* A *queue* class typically consists of the queries *item*, *count*, and *empty* — returning the next element to be dequeued, the total number of elements in the queue, and a fresh empty queue — and the commands *put* and *remove* — to enqueue an element and dequeue the next element. If *remove* were not part of the interface, any element in the queue but the least recently inserted one would be inaccessible to clients; the model of such a class would then be a pair of type  $\mathbb{N} \times G$  recording the current number of elements and the latest enqueued element of generic type  $G$ . Including *remove* in the interface, as it usually is the case for queues, allows clients to read the whole sequence of enqueued elements. Hence, two queues with full interfaces are indistinguishable iff they have the very same sequence of elements; the model of a queue class with full interface is then an abstract sequence of type  $G^*$ .

As all the following examples will suggest, the most natural design choice implements object equality to have the same semantics as abstract equality. Notice, however, that complying or not with this rule of thumb does not affect the soundness of the definitions in the present paper, nor does introduce circularities in the definition of abstract equality.

**Model classes.** The model of a class  $C$  is expressed as a collection  $D_C = D_C^1, D_C^2, \dots, D_C^n$  of *model classes*.<sup>2</sup> Model classes are immutable classes designed for specification purposes; essentially, they are wrappers of rigorously defined mathematical entities: elementary sorts such as Booleans, integers, and object references, as well as more

<sup>2</sup> The model may include the same class multiple times

complex structures such as sets, bags, relations, maps, and sequences. The MML library [27] provides a variety of such model classes, equipped with features that correspond to common operations on the mathematical structure they represent, including first-order quantification. For example, class `MML_SET` models sets of elements of homogeneous type; it includes features for operations such as membership and quantification over all elements of the set that satisfy a certain predicate (passed as a function object).

*Example 3.* As we discussed in Example 2, a sequence is a suitable model for a queue; it can be represented by class `MML_SEQUENCE`. To represent the model of a linked list with internal cursor, we can combine a sequence of class `MML_SEQUENCE` with an element of class `INTEGER` to represent the position of the cursor; this assumes that no information about the pointer structure of the list in the heap is accessible through the interface of the class.

**Model queries.** Every class  $C$  provides a collection of public *model queries*  $S_C = s_C^1, s_C^2, \dots, s_C^n$ , one for each component model class in  $D_C$ . Each model query  $s_C^i$  returns an instance of the corresponding model class  $D_C^i$  that represents the current value of the  $i$ -th component of the model. (Informally, the values returned by model queries are analogues to the coefficients expressing the abstract state as a combination of independent basis vectors spanning the whole space). Since the abstract object state should always be defined between operations and should not depend on the state of any other object, model queries are typically argumentless and without precondition. Clauses in the class invariant can constrain the values of the model queries to match precisely the abstract states of the model. For example, model query *index*: `INTEGER` returning the cursor position of the `LINKED_LIST` in Table 1 should be constrained by an invariant clause  $0 \leq \text{index} \leq \text{sequence.count} + 1$ . A meta-annotation `note model`:  $s_C^1, s_C^2, \dots$  lists all model queries of the class (see Table 2 for an example).

Programmers can add model queries incrementally to classes developed with DbC. In fact, it is likely that some model queries are already used in the implementation before models are added explicitly; for example feature *index* of class `LINKED_LIST` (Table 2). Additional model queries return the remaining components of the model for specification purposes, such as *sequence* in `LINKED_LIST`.

Our approach prefers to implement new model queries as functions rather than attributes. This choice facilitates a purely descriptive usage of references to model queries in specifications. In other words, instead of augmenting routine bodies with bookkeeping instructions that update model attributes, routine postconditions are extended with clauses that describe the new value returned by model queries in terms of the old one. This has the advantage of enforcing a cleaner division between implementation and specification, while better modularizing the latter at routine level (properties of model attributes are typically gathered in the class invariant). A meta-annotation of the form `note specification` tags model queries that are not meant for use in implementation; runtime checking of annotations calling these model queries can be disabled if performance is a concern.

**Model-based contracts.** Let  $C$  be a class equipped with model queries whose interface  $I_C$  is partitioned into queries  $Q_C$  and commands  $M_C$ .  $Q_C$  now includes the model

```

35 note model: sequence, index
36 class LINKED_LIST [G]
37 ...
38 has (v: G): BOOLEAN
39   -- Does list include 'v'? (Reference equality)
40   do ...
41   ensure Result iff sequence.has (v) end
42
43 item: G
44   -- Value at cursor position
45   require
46     sequence.domain [index]
47   ensure
48     Result = sequence [index]
49   end

51 duplicate (n: INTEGER): LINKED_LIST [G]
52   -- A copy of at most 'n' elements
53   -- starting at cursor position
54   require n ≥ 0
55   do ...
56   ensure
57     Result.sequence = sequence.interval (index, index + n - 1)
58     Result.index = 0
59   end
60
61 make_empty
62   -- Create an empty list
63   ensure sequence.is_empty and index = 0
64   end
65 ...
66 end

```

**Table 3.** Snippets of class `LINKED_LIST` with model-based contracts (continued from Table 2).

queries  $S_C \subseteq Q_C$  together with other queries  $R_C = Q_C \setminus S_C$  (note that this does not change the abstract space according to the definitions given at the beginning of the section). Queries in  $R_C$  are called *standard queries*. The rest of the section contains guidelines to writing model-based contracts for commands in  $M_C$  and queries in  $R_C$ .

- The *precondition* of a feature is a constraint on the abstract values of its value-bound arguments and, possibly, on the actual references to its reference-bound arguments. The target object, in particular, can be considered an implicit value-bound argument. For example, the precondition  $map.domain [k]$  of feature *put* in class `TABLE` (Table 2), refers to the abstract state of the target object, given by the model query  $map$ , and to its actual reference-bound argument  $k$ .
- *Postconditions* should refer to abstract states only through model queries. This emphasizes the components of the abstract state that are set by a command or a query, which in turn facilitates understanding and reasoning on the semantics of a feature.
- The *postcondition of a command* defines a relation between the prestate and the poststate of its arguments and the target object; prestate and poststate refer respectively to the state before and after executing the command. More precisely, the postcondition mentions only abstract values of its value-bound arguments and possibly the actual references to its reference-bound arguments; the target object is considered value-bound both in the prestate and in the poststate.

It is common that a command only affects a few components of the abstract state and leaves all the others unchanged. Accordingly, the *closed world assumption* is convenient: the value of any model query  $s \in S_C$  that is not mentioned in the postcondition is assumed not to be modified by the command, as if  $s = \mathbf{old} \ s$  were a clause of the postcondition. When the closed world assumption is wrong, explicit clauses in the postcondition should establish the correct semantics. If a command may modify the value of a model query  $s$  but the actual new value is not known precisely and  $s$  is not mentioned in other clauses of the postcondition, add a clause *relevant* ( $s$ ) to the postcondition of the command (in terms of implementation, *relevant* is just a constant function that returns true). If a command does not affect



<pre> 2 <b>note model:</b> bag 3 <b>class</b> COLLECTION [G] 4   bag: MML_BAG [G] 5 6   <b>is_empty:</b> BOOLEAN 7     <b>ensure</b> Result = bag.is_empty <b>end</b> 8 9   wipe_out 10  <b>ensure</b> bag.is_empty <b>end</b> 11 12  put (v: G) 13  <b>ensure</b> bag = old bag.extended (v) <b>end</b> 14 <b>end</b> </pre>	<pre> 16 <b>note model:</b> sequence 17 <b>class</b> DISPENSER [G] 18 <b>inherit</b> COLLECTION [G] 19 20  sequence: MML_SEQUENCE [G] 21 22  <b>invariant</b> 23    bag.domain = sequence.range 24    bag.domain.for_all ( agent (x: G): BOOLEAN 25                        bag [x] = sequence.occurrences (x) ) 26 <b>end</b> </pre>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**Table 4.** Snippets of classes `COLLECTION` (left) and `DISPENSER` (right) with model-based contracts.

the value a model query  $s$  but the postcondition of the command mentions  $s$ , add a clause  $s = \mathbf{old} \ s$  to the postcondition of the command.

- The *postcondition of a query* defines the result as a function of the target object and its arguments (with the usual discipline of mentioning only abstract values of value-bound arguments and target object and possibly actual references to reference-bound arguments). Value-bound queries define the abstract state of the result, whereas reference-bound queries describe an actual reference to it. For example, compare the postcondition of the reference-bound query *item* from class `LINKED_LIST` (Table 3), which precisely defines a reference to the returned list element, with the postcondition of the value-bound query *duplicate* in the same class, which specifies the abstract state of the returned list.
- A clear-cut separation between queries and commands assumes *abstract purity* for all queries: executing a query leaves the abstract state of all its arguments and of the targeted object unchanged.

**Inheritance and model-based contracts.** A class  $C'$  that inherits from a parent class  $C$  may or may not re-use  $C$ 's model queries to represent its own abstract state. For every model query  $s_C \in S_C$  of the parent class that is not among the heir's model queries  $S_{C'}$ ,  $C'$  should provide a *linking invariant* to guarantee consistency in the inheritance hierarchy. The linking invariant is a formula that defines the value returned by  $s_C$  in terms of the values returned by the model queries  $S_{C'}$  of the inheriting class. This guarantees that the new model is indeed a specialization of the previous model, in accordance with the notion of subtyping inheritance.

A properly defined linking invariant ensures that every inherited feature has a definite semantics in terms of the new model. However, the new semantics may be weaker in that a command whose contract in the parent class characterized it as a function, becomes characterized as a relation in the child class; that is, incompleteness is introduced (see Section 3.2).

*Example 4.* Consider class `COLLECTION` in Table 4, a generic container of elements whose model is a bag. Class `DISPENSER` inherits from `COLLECTION` and specializes it

by introducing a notion of insertion order; correspondingly, its model is a sequence. The linking invariant of **DISPENSER** defines the value of the inherited feature *bag* in terms of the new feature *sequence*: the domain of *bag* coincides with the range of *sequence*, and the number of occurrences of any element  $x$  in *bag* correspond to the number of occurrences of the same element in *sequence*.

The linking invariant ensures that the semantics of features *is\_empty* and *wipe\_out* is unambiguously defined also in **DISPENSER**. On the other hand, the model-based contract of command *put* in **COLLECTION** and the linking invariant are insufficient to characterize the effects of *put* in **DISPENSER**, as the position within the sequence where the new element is inserted is irrelevant for the bag.

### 3.2 Completeness of contracts

The notion of *completeness* for the specification of a class gives an indication of how accurate are the contracts of that class with respect to the model. An incomplete contract does not fully capture the effects of a feature, suggesting that the contract may be more detailed or, less commonly, that the model of the class — and hence its interface — is not abstract enough. Unlike the notion of *sufficient completeness* for algebraic specifications [14] — that serves a similar purpose —, the present definition of completeness is structurally similar to the concept of completeness for a set of axioms, and a dual notion of soundness complements it. For simplicity, the following definitions do not mention feature arguments; introducing them is, however, routine.

**Soundness and completeness of a model-based contract.** Let  $f$  be a feature of class  $C$ . The specification of  $f$  denotes two predicates  $\mathbf{pre}_f$  and  $\mathbf{post}_f$ .  $\mathbf{pre}_f$  represents the set of objects of class  $C$  that satisfy the precondition. If  $f$  is a query returning object of class  $T$ ,  $\mathbf{post}_f$  is of type  $C \times T$  and denotes the pairs of target and returned objects. If  $f$  is a command,  $\mathbf{post}_f$  is of type  $C \times C$  and denotes the pairs of target objects before and after executing the command.<sup>3</sup>

- The *precondition* of a feature  $f$  (query or command) is *sound* iff: for every  $o_1, o_2 \in C$  such that  $o_1 \succ_C o_2$  it is  $\mathbf{pre}_f(o_1) \Leftrightarrow \mathbf{pre}_f(o_2)$ .<sup>4</sup>
- The *postcondition* of a command  $m$  is *sound* iff: for every  $o, o'_1, o'_2 \in C$  such that  $o'_1 \succ_C o'_2$  it is  $\mathbf{post}_m(o, o'_1) \Leftrightarrow \mathbf{post}_m(o, o'_2)$ .  
The *postcondition* of a command  $m$  is *complete* iff: for every  $o, o'_1, o'_2 \in C$  such that  $\mathbf{post}_m(o, o'_1)$  and  $\mathbf{post}_m(o, o'_2)$  it is  $o'_1 \succ_C o'_2$ .
- The *postcondition* of a value-bound query  $q$  is *sound* iff: for every  $o \in C$  and  $t_1, t_2 \in T$  such that  $t_1 \succ_T t_2$  it is  $\mathbf{post}_q(o, t_1) \Leftrightarrow \mathbf{post}_q(o, t_2)$ .  
The *postcondition* of a value-bound query  $q$  is *complete* iff: for every  $o \in C$  and  $t_1, t_2 \in T$  such that  $\mathbf{post}_q(o, t_1)$  and  $\mathbf{post}_q(o, t_2)$  it is  $t_1 \succ_T t_2$ .
- The *postcondition* of a reference-bound query  $q$  is *sound* iff: for every  $o \in C$  and  $t_1, t_2 \in T$  such that  $t_1 \equiv_T t_2$  it is  $\mathbf{post}_q(o, t_1) \Leftrightarrow \mathbf{post}_q(o, t_2)$ .  
The *postcondition* of a reference-bound query  $q$  is *complete* iff: for every  $o \in C$  and  $t_1, t_2 \in T$  such that  $\mathbf{post}_q(o, t_1)$  and  $\mathbf{post}_q(o, t_2)$  it is  $t_1 \equiv_T t_2$ .

<sup>3</sup> These definitions imply the absence of side-effects in evaluating assertions.

<sup>4</sup> Completeness of preconditions is not an interesting notion and hence it is not defined.

Informally, a sound assertion is one that is consistent with the notion of equivalence that is appropriate: sound postconditions of commands and value-bound queries do not distinguish between objects with the same abstract state; sound postconditions of reference-bound queries do not distinguish between aliases.<sup>5</sup>

A postcondition is complete if all the pairs of objects that satisfy it are equivalent (according to the right model of equivalence). This means that the complete postcondition of a command defines the effects of the command as a mathematical *function* (as apposed to a relation) from the prestate to the abstract poststate. Similarly, the complete postcondition of a query defines the value of the result as a *function* of the abstract state of argument-bound arguments and of actual references to reference-bound arguments.

*Example 5.* The contracts of features *is\_empty*, *wipe\_out*, and *put* in class `COLLECTION` (Table 4) are sound and complete; the postcondition of *put*, in particular, is complete as it defines the new value of *bag* uniquely. In the heir class `DISPENSER`, however, the inherited postcondition of *put* becomes incomplete: the linking invariant does not uniquely define *sequence* from *bag*, hence inequivalent sequences (for example, one with *v* inserted at the beginning and another one with *v* at the end) satisfy the postcondition.

**Soundness and completeness in practice.** As the previous example suggests, reasoning informally — but precisely — about soundness and completeness of model-based contracts is often straightforward and intuitive, especially if the guidelines of Section 3.1 have been followed. Completeness captures the uniqueness of the (abstract) state described by a postcondition, hence query postconditions in the form **Result** = *exp* (*s*, *a*) or **Result**.*s* = *exp* (*s*, *a*) and command postconditions in the form *s* = *exp* (**old** *s*, *a*) — where *exp* is a side-effect free expression, *s* denotes the value returned by the model query of some argument, and *a* is a reference-bound argument — are painless to check for completeness.

*Example 6.* Consider the following example, from class `ARRAY` whose model is a map.

```

2  fill (v: G ; l, u: INTEGER) -- Put 'v' at all positions in [l, u].
3  require map.domain [l] and map.domain [u]
4  ensure map.domain = old map.domain
5      ( map | {MML_INT_SET} [[l, u]] ).is_constant (v)
6      ( map | (map.domain - {MML_INT_SET} [[l, u]]) ) =
7          old ( map | (map.domain - {MML_INT_SET} [[l, u]]) )
8  end

```

Pre and postconditions are sound because they both refer only to model queries, or functions thereof. The following reasoning shows that the postcondition is also complete: a map is uniquely defined by its domain and by a value for every key in the domain. The first clause of the postcondition defined the domain completely. Then, let *k* be any key in the domain. If *k* ∈ [*l*, *u*] then the second clause defines *map* (*k*) = *v*; otherwise *k* ∉ [*l*, *u*], and the third clause postulates *map*(*k*) unchanged.

<sup>5</sup> Postconditions of argumentless reference-bound queries are trivially sound for sensible definitions of reference equality.

Soundness is an irrenounceable requirement for pre and postconditions in the presence of model-based contracts, as it boils down to writing contracts that are consistent with the chosen level of information hiding.

On the other hand, how useful is completeness in practice? As a norm, completeness is a valuable yardstick to evaluate whether the contracts are sufficiently detailed. This is not enough to guarantee that the contracts are correct — and meet the original requirements — but the yardstick is serviceable methodologically to focus on what a routine really achieves and how that is related to the abstract model. As a result, inconsistencies in specifications are less likely to occur, and the impossibility of systematically writing complete contracts is a strong indication that the model is incorrect, or the implementation is faulty. Either way, a warning is available before attempting a correctness proof.

While complete postconditions should be the norm, there are recurring cases where incomplete postconditions are unavoidable or even preferable. Three major sources of benign incompleteness are the following.

- Inherently *nondeterministic or stochastic* specifications. For example, a class for random number generation can use a sequence as model, but its specification should not define the precise content of the sequence unambiguously.
- Usage of *inheritance* to factor out common parts of (complete) specifications. For example, class `DISPENSER` in Table 4 is a common ancestor of `STACK` and `QUEUE`. If its interface includes features *item*, *put* and *remove*, its model must be isomorphic to a sequence. Then, it becomes impossible to write a complete postcondition for *put* in `DISPENSER`: the specification of *put* cannot define precisely where an element is added to the sequence; a choice compatible with the semantics of `STACK` will be incompatible with `QUEUE` and vice versa.
- Imperfections in *information hiding*. For example, class `ARRAYED_LIST` is an array-based implementation of lists which exports a query *capacity* returning the size of the underlying array; this piece of information is then part of the model of the class. Default constructors set *capacity* to an initial fixed value. Their postconditions, however, do not mention this default value, hence they are incomplete. The rationale behind not revealing this information is that clients should not rely on the exact size of the array when they invoke the constructor.

In all these cases, reasoning about completeness is still likely to improve the understanding of the classes and to question constructively the choices made for interfaces and inheritance hierarchies.

### 3.3 Verification: proofs and runtime checking

This subsection outlines the main ideas behind using model-based contracts for verification with formal correctness proofs and with runtime checking for automated testing. Its goal is not to detail any particular proof or testing technique, but rather to sketch how to express the semantics of model-based contracts within standard verification frameworks.

<pre> 2 note mapped_to: "Sequence G" 3 class MML_SEQUENCE [G] 4 ... 5 extended (x: G): MML_SEQUENCE[G] 6   -- Current sequence extended with 'x' at the end 7   note mapped_to: "Sequence.extended(Current, x)" 8   do ... end 9 end </pre>	<pre> 9 type Sequence T = [int] T ; 10 function Sequence.extended (T) (Sequence T, T) 11   returns (Sequence T); 12 axiom (∀ (T) s: Sequence T, x:T • {Sequence.extended(s, x)} 13   Sequence.extended(s, x) =s[Sequence.count(s)+1 := x]); 14 axiom (∀ (T) s: Sequence T, x: T • 15   {Sequence.count(Sequence.extended(s, x))} 16   Sequence.count(Sequence.extended(s, x)) = 17     Sequence.count(s)+1; 18 ... </pre>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

**Table 5.** Snippets from class `MML_SEQUENCE` (left) and the corresponding Boogie theory (right).

**Proofs.** The *axiomatic* treatment of model classes [6, 27, 9] is quite natural: the semantics of a model class is defined directly in terms of a theory expressed in the underlying proof language, rather than with “special” contracts. The mapping is often straightforward, and has the advantage of reusing theories that are optimized for effective usage with the proof engine of choice. In addition, the immutability (and value semantics) of model classes makes them very similar to mathematical structures and facilitates a straightforward translation into mathematical theories.

In this respect, we are currently developing an accurate mapping of model classes and model-based contracts into Boogie [2]. First, the mapping introduces axiomatic definitions of MML model classes as Boogie theories; annotations in the form `note mapped_to` connect MML classes to the corresponding Boogie types. For example, Table 5 shows how a portion of the `MML_SEQUENCE` model class translates into a Boogie theory: a mapping type `[int] T` represents sequences of elements of generic type `T`, and a few axioms constrain a function `Sequence.extended` to return values in accordance with the MML semantic of feature `extended`.

Then, each model query in a class with model-based contracts maps to a Boogie function that references a representation of the heap; some axioms connect the value returned by the function to other features in the translated class. For example, the model query `sequence` in `LINKED_LIST` becomes `function LinkedList.sequence(HeapType, ref) returns (Sequence ref)`.

Finally, model-based contracts are translated into Boogie formulas according to the `mapped_to` annotations in model classes. For example, the postcondition clause: `sequence = old (sequence.front (index).extended (v)+ sequence.tail (index + 1))` of `put_right` in `LINKED_LIST` (Table 2) maps to the Boogie formula:

$$\text{LinkedList.sequence(Heap, Current)} = \text{Sequence.concat} ( \text{Sequence.extended} ( \text{Sequence.front} ( \text{LinkedList.sequence}(\text{old}(\text{Heap}), \text{Current}), \text{LinkedList.index}(\text{old}(\text{Heap}), \text{Current}), v ), \text{Sequence.tail} ( \text{LinkedList.sequence}(\text{old}(\text{Heap}), \text{Current}), \text{LinkedList.index}(\text{old}(\text{Heap}), \text{Current}) + 1 ) );$$

**Runtime checking and testing.** Most model classes represent *finite* mathematical objects, such as sets of finite cardinality, sequences of finite length, and so on. All these

classes can have an implementation of their operations which is executable in finite time; this supports the runtime checking of assertions that reference these model classes.

Testing techniques can leverage runtime checkable contracts to fully automate the testing process: generate objects by randomly calling constructors and commands; check the precondition of a routine on the generated objects to filter out valid inputs for the routine; execute the routine body on a valid input and check the validity of the postcondition on the result; any postcondition violation on a valid input is a fault in the routine.

This approach to contract-based testing has proved extremely effective at uncovering plenty of bugs in production code [23], hence it is an excellent “lightweight” precursor to correctness proofs. Contract-based testing, however, is only as good as the contracts are; the weak postconditions of traditional DbC, in particular, leave many real faults undetected. Runtime checkable model-based contracts can help in this respect and boost the effectiveness of contract-based testing by providing more expressive, and complete, specifications. Section 4 describes some testing experiments that support this claim.

**Consistency of tests and proofs.** Using contract-based testing as a precursor to correctness proofs poses the problem of consistency between two semantics given to model classes: the runtime semantics given by an executable implementation and the proof semantics given by a mapping to a logical theory. Under reasonable assumptions about the execution environment, consistency must ensure that a component is proven correct against its model-based specification if and only if testing the component never detects a violation of its model-based contracts. Establishing this consistency amounts to proving that: (1) the implementation of each model class is consistent with the mapping of the class to a logical theory; and (2) the implementation of each model query satisfies its specification. Future work will detail and address these problems.

## 4 Model-based contracts at work

This section describes experiments in developing model-based contracts for real object-oriented software written in Eiffel. The experiments target two non-trivial case studies based on data-structure libraries (described in Section 4.1) with the goal of demonstrating that deploying model-based contracts is feasible, practical, and useful. Section 4.2 discusses the successes and limitations highlighted by the experiments.

### 4.1 Case studies

The first case study targeted EiffelBase [12], a library of general-purpose data structures widely used in Eiffel programs; EiffelBase is representative of mature Eiffel code exploiting extensively traditional DbC. We selected 7 classes from EiffelBase, for a total of 304 features (254 of them are public) over more than 5700 lines of code. The 7 classes include 3 widely used container data structures ([ARRAY](#), [ARRAYED\\_LIST](#), and [LINKED\\_LIST](#)) and 4 auxiliary classes used by the containers ([INTEGER\\_INTERVAL](#), [LINKABLE](#), [ARRAYED\\_LIST\\_CURSOR](#), and [LINKED\\_LIST\\_CURSOR](#)). Our experiments

```

2 note model: set, relation
3 class SET [G]
4 ...
5 has (v: G): BOOLEAN
6   -- Does this set contain 'v'?
7   ensure
8     Result = not (set * relation.image_of(v)).is_empty
9   end
10
11 set: MML_SET [G] -- The set of elements
12 relation: MML_RELATION [G, G]
13   -- Equivalence relation on elements
14 end

16 note model: map
17 class BINARY_TREE [G]
18 ...
19 add_root (v: G)
20   -- Add a root with value 'v' to an empty tree
21   require map.is_empty
22   ensure map.count = 1 and map [Empty] = v
23   end
24
25 map: MML_MAP [MML_SEQUENCE[BOOLEAN], G]
26   -- Map of paths to elements
27 end

```

**Table 6.** Examples of nonobvious models: classes `SET` and `BINARY_TREE` from EiffelBase2.

systematically introduced models and conservatively augmented the contracts of all public features in these 7 classes with model-based specifications.

The second case study developed EiffelBase2, a new general-purpose data structure library. The design of EiffelBase2 is similar to that of its precursor EiffelBase; EiffelBase2, however, has been developed from the start with expressive model-based specifications and with the ultimate goal of proving its full functional correctness — backward compatibility is not one of its primary aims. This implies that EiffelBase2 rediscusses and solves any deficiency and inconsistency in the design of EiffelBase that impedes achieving full functional correctness or hinders the full-fledged application of formal techniques. EiffelBase2 provides containers such as arrays, lists, sets, tables, stacks, queues, and binary trees; iterators to traverse these containers; and comparator objects to parameterize containers with respect to arbitrary equivalence and order relations on their elements. The current version of EiffelBase2 includes 46 classes with 460 features (403 of them are public) totalling about 5800 lines of code; these figures make EiffelBase2 a library of substantial size with realistic functionalities. The latest version of EiffelBase2 is available at <http://eiffelbase2.origo.ethz.ch>.

## 4.2 Results and discussion

This section addresses the following questions based on the experience with the two case studies of EiffelBase and EiffelBase2.

- How many different model classes are needed to write model-based contracts?
- How many contracts can be complete?
- Do executable accurate model-based contracts boost contract-based testing?

**How many model classes?** Model-based contracts for EiffelBase used model classes for Booleans, integers, references, (finite) sets, relations, and sequences. EiffelBase2 additionally required (finite) maps, bags, and infinite maps and relations for special purposes (such as modeling comparator objects). These figures suggest that a moderate

number of well-understood mathematical models suffices to specify a general-purpose library of data structures.

Determining to what extent this is generalizable to software other than libraries of general-purpose data structures is an open question which belongs to future work. Domain-specific software may indeed require complex domain-specific model classes (e.g., real-valued functions, stochastic variables, finite-state machines), and application software that interacts with a complex environment may be less prone to accurate documentation with models. However, even if writing model-based contracts for such systems proved exceedingly complex, some formal model is required if the goal is formal verification. In this sense, focusing model-based contracts on library software is likely to have a great payoff through extensive reuse: the many clients of the reusable components can rely on expressive contracts not only as detailed documentation but also to express their own contracts and interfaces by combining a limited set of well-understood, highly dependable components.

Another interesting remark is that the correspondence between the limited number of model classes needed in our experiments and the classes using these model classes is far from trivial: data structures are often more complex than the mathematical structures they implement. Consider, for example, class `SET` in Table 6: EiffelBase2 sets are parameterized with respect to an equivalence relation, hence the model of `SET` is a pair of a mathematical set and a relation. Another significant example is `BINARY_TREE` (also in Table 6): instead of introducing a new model class for trees or graphs, `BINARY_TREE` concisely represents a tree as a map of paths to values; the model of a path is in turn a sequence of Booleans.

**How many complete contracts?** Reasoning informally, but rigorously, about the completeness of postconditions — along the lines of Section 3.2 — proved to be straightforward in our experiments. Only 18 (7%) out of 254 public features in EiffelBase with model-based contracts and 17 (4%) out of 403 public features in EiffelBase2 have incomplete postconditions. All of them are examples of “intrinsic” incompleteness mentioned at the end of Section 3.2; EiffelBase2, in particular, was designed trying to minimize the number of features with intrinsically incomplete postconditions.

These results indicate that model-based contracts make it feasible to write systematically complete contracts; in most cases this was even relatively straightforward to achieve. Unsurprisingly, using model-based contracts dramatically increases the completeness of contracts in comparison with standard DbC. For example, 42 (66%) out of 64 public features of class `LIST` in the original version of EiffelBase (without model-based contracts) have incomplete postconditions, including 20 features (31%) without any postcondition.

**Contract-based testing with model-based contracts.** The standard EiffelBase library has been in use for many years and has been extensively tested, both manually and automatically. Are the expressive contracts based on models enough to boost automated testing finding new, subtle bugs? While preliminary, our experiments seem to answer in the affirmative. Applying the AutoTest testing framework [23] on EiffelBase with model-based contracts for 30 minutes discovered 3 faults; none of them would have



```

2 merge_right (other: LINKED_LIST [G])
3   -- Merge 'other' into current list after cursor position. Do not move cursor. Empty 'other'.
4   do
5     ...
6     other_first_element := other.first_element ; other_count := other.count ; other.wipe_out
7     if before then first_element := other.first_element ; active := first_element
8     else ... end
9     count := count + other_count
10  ensure
11    -- Original contract
12    count = old count + old other.count ; index = old index ; other.is_empty
13    -- Model based contract
14    sequence = old (sequence.front (index) + other.sequence + sequence.tail (index + 1))
15  end

```

**Table 7.** Faulty routine *merge\_right* from class `LINKED_LIST`.

been detectable with standard contracts. Running these tests did not require any modification to AutoTest or model classes, because the latter include an executable implementation.

The 3 faults reveal subtle mistakes that have gone undetected so far. For example, consider the implementation of routine *merge\_right* in Table 7; the routine merges a linked list *other* into the current linked list at the cursor position by modifying references in the chain of elements. The **then** branch of the **if** statement (line 6) deals with the special case where the cursor in the current list is *before* the first element; in this case the first element of the current list (*first\_element*) will point directly to the first element of the other list (*other\_first\_element*). This is not sufficient, as the routine should also link the end of the other list to the front of the current one, otherwise all elements in the current list become inaccessible. The original contract does not detect this fault; the clause  $count = \text{old } count + \text{old } other.count$  is in particular satisfied as *count* is anyway updated (line 8), but its value does not reflect the actual content of the new list. On the contrary, the complete model-based contract (line 13) specifies the desired configuration of the list after executing the command, which leads to easily detecting the error.

## 5 Related work

Every fully formal specification ultimately boils down to a mathematical model, and the research on formal modeling and analysis is so extensive and diverse that it cannot be summarized concisely. This section focuses on a few major approaches to the formal specification of object-oriented abstract data types that adopt a stance similar to that of the present paper: using highly expressive mathematical models geared towards the full functional correctness specification (and verification) of complex data structures.

Hoare pioneered the usage of mathematical models to define and prove correctness of data type implementations [16]. This idea spawned much related work, which can be roughly partitioned in three major lines: algebraic notations, descriptive notations, and design-by-contract approaches. The following subsections shortly summarize the main features of each of these techniques; then, Section 5.4 describes the approaches based on mathematical models that are closest to the present paper.

## 5.1 Algebraic notations

Algebraic notations formalize classes in terms of (uninterpreted) functions and axioms that describe the mutual relationship among the functions. For example, the axiom  $s.\text{insert}(x).\text{member\_of}(x) = \text{True}$  defines the mutual semantics of the operations `insert` and `member_of` of a set data type. The most influential work in algebraic specifications is arguably Guttag and Horning's [14] and Gougen et al.'s [13], which gave a foundation to much derivative work. The former was also made practical in the Larch project [15], and introduced a notion of *completeness* that differs from the one of the present paper (see Section 3.2), and applies to whole types, not single features.

Algebraic notations emphasize the calculational aspect of a specification. This makes them very effective notations to formalize and verify data types at a high level of abstraction. In particular, the close connection between rewriting systems [10] and algebraic definitions enables, in many practical cases, the automated or semi-automated verification of consistency and completeness [14] requirements of abstract specifications. The algebraic approach, on the other hand, does not integrate as well with real programming languages to document implementations in the form of pre and postconditions of single operations.

## 5.2 Descriptive notations

Descriptive notations formalize classes in terms of simpler types — ultimately grounded in simple mathematical models such as sets and relations — and operations defined as input/output relations (that is, pre and postconditions) constrained by logic or arithmetic formulas. For example, the `insert` operation of a set data structure could be defined by the formula  $\forall s, x \bullet \llbracket s.\text{insert}(x) \rrbracket = \llbracket s \rrbracket \cup \{x\}$ , in terms of the union operation applied to a model set  $\llbracket s \rrbracket$ .

Descriptive notations can be used in isolation to build language-independent models, or to give a formal semantics to concrete implementations. Languages and methods such as Z [29], B [1], and VDM [17] pursue the former approach, usually within a top-down development framework. Other specification languages and tools such as RESOLVE [24], AAL [18], and Jahob [30] are examples of the latter approach for the programming languages C<sup>++</sup> and Java.

Descriptive notations are apt to develop correct-by-construction designs and to accurately document implementations, often with the goal of verifying functional correctness. Using them in contracts, however, introduces a new notation on top of the programming language, which requires additional effort and expertise from the programmer and makes it more difficult to maintain the specification synchronized with the actual implementation. This weakness is shared by algebraic notations alike.

## 5.3 Design-by-contract approaches

Design by contract [22] introduces formal specifications in programs using the same notation for implementation and annotations, in an attempt to make writing the contracts as congenial as possible to programmers. For example, the `insert` operation of a multiset class could have a postcondition clause such as `count > 0` that defines an effect of

the insert operation in terms of the value returned by another function count of the same class. The Eiffel programming language [11] epitomizes the design by contract methodology, together with similar solutions for other languages such as APP [26] for C, Spec# [2] for C#, and many others.

As we discussed also in the rest of the paper, using a subset of the programming language in annotations helps programmers writing them [5], but it often does not provide enough expressive power to formalize “complete” functional correctness, or requires cumbersome workarounds to capture the semantics of mathematical concepts in terms of programming language constructs. Going back to the example of the set class, it is impossible to express directly (with quantification over the domain) the fact that insert does not remove any element that was in the set before insertion. The semantics of quantification could still be expressed as iteration over the data structure. This, however, is unintuitive and programmers tend not to write such assertions [25]; furthermore, it does not quite solve the problem but only reduces it to the — arguably simpler but still error-prone — problem of ensuring that the iteration over the data structure realizes the intended quantification semantics without incurring in misleading side effects.

#### 5.4 Model-based annotation languages

The Java Modeling Language (JML) [20, 19] is likely the approach that shares the most similarities with ours: JML annotations are based on a subset of the Java programming language and the JML framework provides a library of model classes mapping mathematical concepts. While sharing a common outlook, the approaches in JML and in the present paper differ in several details pertaining scope and technical aspects.

At the technical level, JML prefers model variables [7] while our approach leverages model queries that return the value of immutable model classes; each approach has its merits, but model queries have the advantage of supporting an axiomatic definition that is easily grounded in an underlying mathematical theory, and facilitate a seamless integration with traditional contracts — also typically based on queries. Section 3.1 discusses other advantages of model queries. A notational difference is that JML extends Java’s expressions with notations for logic operators and quantifiers, while our method does not extend Eiffel’s syntax and reuses notation such as agents to express quantifications and other aspects that belong to expressive specifications.

In terms of scope, our approach strives to be more methodological and systematic, with the primary target of fully contracting a complete library of data structures. Our method tries to keep the additional effort required to the programmer to a minimum; this is the case, for example, with frame conditions that are extracted automatically from postconditions in many cases (although our solution is still partial and certainly requires further investigation). Finally, let us remark that our usage scenarios are multi-faceted, ranging from specification and design (also supporting notions such as completeness), to verification, runtime checking, and automated testing.

The present paper extends in scope the previous work of ours on model-based classes [28, 27], and systematically applies the results to the re-design and re-implementation of a rich library of data structures. The experience gained in this practical application also prompted us to refine and rediscuss aspects of the previous approach, as we discussed at length in the rest of the paper.

## 6 Conclusions and future work

Present work makes the following contributions:

- A method for writing strong interface specifications for reusable object-oriented components; we give a systematic formal description of the method and define a notion of specification completeness, which is easy to reason about.
- A library of reusable components supplied with strong specification, which demonstrates applicability of the proposed specification method and its benefits for automated contract-based testing.

There are many directions of future work that we wish to pursue. First, the proposed specification method has to be tried out on more libraries and application from diverse problem domains. Second, a user study is needed to justify our intuition that model-based contracts are easy to write, understand and reason about in practice.

As the ultimate goal we see model-based contracts as a part of an integrated verification environment: a software development environment that employs a wide range of tools and techniques to assist a programmer in constructing correct software. Achieving this goal requires a lot of work in the direction of both proofs and testing.

On the side of proofs our method has to be extended to non-interface specifications: abstraction functions, representation invariants, loop invariants are often more complex than public contracts. We also have to refine the model-based approach to specifying frame properties.

To enable automated proofs we have to implement the translation of model-based contracts from Eiffel to Boogie. We also have to provide full Boogie theories for all classes in the MML library. With these tools at hand we plan to prove correctness of the EiffelBase2 library.

Except for correctness proofs, other interesting types of proofs can be done: for example, formal proofs of specification completeness as defined in this paper as well as proofs of consistency between the implementations of model classes and their mappings into different theories.

In the direction of testing the main issue is the runtime efficiency of model-based contracts. We plan to experiment with different implementation of model classes and try to minimize the runtime penalty caused by their immutability.

## References

1. J.-R. Abrial. *The B-book: assigning programs to meanings*. Cambridge University Press, New York, NY, USA, 1996.
2. M. Barnett, R. DeLine, M. Fähndrich, B. Jacobs, K. R. M. Leino, W. Schulte, and H. Venter. The Spec# programming system: Challenges and directions. In *Verified Software: Theories, Tools, Experiments, First IFIP TC 2/WG 2.3 Conference (VSTTE 2005)*, volume 4171 of *Lecture Notes in Computer Science*, pages 144–152. Springer, 2008.
3. M. Barnett, K. R. M. Leino, and W. Schulte. The Spec# programming system: An overview. In *CASSIS 2004*, volume 3362 of *LNCS*. Springer, 2004.

4. M. Barnett, B. yuh Evan Chang, R. Deline, B. Jacobs, and K. R. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *Formal Methods for Components and Objects: 4th International Symposium, FMCO 2005, volume 4111 of Lecture Notes in Computer Science*, pages 364–387. Springer, 2006.
5. P. Chalin. Are practitioners writing contracts? In *Rigorous Development of Complex Fault-Tolerant Systems (RODIN Book)*, volume 4157 of *Lecture Notes in Computer Science*, pages 100–113. Springer, 2006.
6. J. Charles. Adding native specifications to jml. In *In Workshop on Formal Techniques for Java-like Programs (FTfJP)*, 2006.
7. Y. Cheon, G. Leavens, M. Sitaraman, and S. Edwards. Model variables: cleanly supporting abstraction in design by contract. *Softw. Pract. Exper.*, 35(6):583–599, 2005.
8. I. Ciupa, B. Meyer, M. Oriol, and A. Pretschner. Finding faults: Manual testing vs. random+ testing vs. user reports. In *Proceedings of ISSRE (International Symposium on Software Reliability) 2008*, 2008.
9. A. Darvas and P. Müller. Faithful mapping of model classes to mathematical structures. In *SAVCBS '07: Proceedings of the 2007 conference on Specification and verification of component-based systems*, pages 31–38, New York, NY, USA, 2007. ACM.
10. N. Dershowitz and J.-P. Jouannaud. Rewrite systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, pages 243–320. Elsevier and MIT Press, 1990.
11. ECMA International. *Standard ECMA-367. Eiffel: Analysis, Design and Programming Language*. 2nd edition, June 2006.
12. <http://freeelks.svn.sourceforge.net>.
13. J. A. Gougen, J. W. Thatcher, and E. G. Wagner. An initial algebra approach to the specification, correctness, and implementation of abstract data types. In R. Yeh, editor, *Current Trends in Programming Methodology*, volume IV, pages 80–149. Prentice Hall, 1978.
14. J. V. Guttag and J. J. Horning. The algebraic specification of abstract data types. *Acta Inf.*, 10:27–52, 1978.
15. J. V. Guttag, J. J. Horning, S. J. Garl, K. D. Jones, A. Modet, and J. M. Wing, editors. *Larch: Languages and Tools for Formal Specification*. Springer-Verlag, 1993.
16. C. A. R. Hoare. Proof of correctness of data representations. *Acta Inf.*, 1:271–281, 1972.
17. C. B. Jones. *Systematic software development using VDM*. Prentice-Hall, 2nd edition, 1990.
18. S. Khurshid, D. Marinov, and D. Jackson. An analyzable annotation language. In *OOPSLA*, pages 231–245, 2002.
19. G. T. Leavens, A. L. Baker, and C. Ruby. Preliminary design of JML: a behavioral interface specification language for java. *SIGSOFT Softw. Eng. Notes*, 31(3):1–38, 2006.
20. G. T. Leavens, Y. Cheon, C. Clifton, C. Ruby, and D. R. Cok. How the design of JML accommodates both runtime assertion checking and formal verification. *Sci. Comput. Program.*, 55(1-3):185–208, 2005.
21. B. H. Liskov and J. M. Wing. A behavioral notion of subtyping. *ACM Trans. Program. Lang. Syst.*, 16(6):1811–1841, 1994.
22. B. Meyer. *Object-oriented software construction*. Prentice Hall, 2nd edition, 1997.
23. B. Meyer, A. Fiva, I. Ciupa, A. Leitner, Y. Wei, and E. Stapf. Programs that test themselves. *Computer*, 42(9):46–55, 2009.
24. W. F. Ogden, M. Sitaraman, B. W. Weide, and S. H. Zweben. The RESOLVE framework and discipline. *ACM SIGSOFT Software Engineering Notes*, 19(4):23–28, 1994.
25. N. Polikarpova, I. Ciupa, and B. Meyer. A comparative study of programmer-written and automatically inferred contracts. In *ISSTA '09: Proceedings of the eighteenth international symposium on Software testing and analysis*, pages 93–104, New York, NY, USA, 2009. ACM.
26. D. S. Rosenblum. Towards a method of programming with assertions. In *ICSE*, pages 92–104, 1992.

27. B. Schoeller. *Making classes provable through contracts, models and frames*. PhD thesis, ETH Zurich, 2007.
28. B. Schoeller, T. Widmer, and B. Meyer. Making specifications complete through models. In *Architecting Systems with Trustworthy Components*, pages 48–70, 2004.
29. J. Woodcock and J. Davies. *Using Z: specification, refinement, and proof*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1996.
30. K. Zee, V. Kuncak, and M. C. Rinard. Full functional verification of linked data structures. In *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation (PLDI'08)*, pages 349–361. ACM, 2008.