

VALloy

Virtual Functions Meet a Relational Language

Darko Marinov and Sarfraz Khurshid

MIT Laboratory for Computer Science
200 Technology Square
Cambridge, MA 02139 USA
{marinov,khurshid}@lcs.mit.edu

Abstract. We propose VALloy, an extension to the first order, relational language Alloy. Alloy is suitable for modeling structural properties of object-oriented software. However, Alloy lacks support for dynamic dispatch, i.e., function invocation based on actual parameter types. VALloy introduces virtual functions in Alloy, which enables intuitive modeling of inheritance. Models in VALloy are automatically translated into Alloy and can be automatically checked using the existing Alloy Analyzer. We illustrate the use of VALloy by modeling object equality, such as in Java. We also give specifications for a part of the Java Collections Framework.

1 Introduction

Object-oriented design and object-oriented programming have become predominant software methodologies. An essential feature of object-oriented languages is *inheritance*. It allows a (sub)class to inherit variables and methods from superclasses. Some languages, such as Java, only support single inheritance for classes.

Subclasses can *override* some methods, changing the behavior inherited from superclasses. We use C++ term *virtual functions* to refer to methods that can be overridden. Virtual functions are *dynamically dispatched*—the actual function to invoke is selected based on the dynamic types of parameters. Java only supports single dynamic dispatch, i.e., the function is selected based only on the type of the *receiver* object.

Alloy [6] is a first order, declarative language based on relations. Alloy is suitable for specifying structural properties of software. Alloy *specifications* can be analyzed automatically using the Alloy Analyzer (AA) [5]. Given a finite *scope* for a specification, AA translates it into a propositional formula and uses SAT solving technology to generate *instances* that satisfy the properties expressed in the specification.

Alloy supports some features of object-oriented design. However, Alloy does not have built in support for dynamic dispatch. Recently, Jackson and Fekete [4] presented an approach for modeling parts of Java in Alloy, pointing out that modeling “the notion of equality is problematic”.

In Java, the `equals` method, which allows comparing object values, as opposed to using the `'=='` operator, which compares object identities, is overridden in majority of classes. Good programming methodology suggests that `equals` be overridden in all immutable classes [19]. This method is pervasively used, for example in the Java Collections Framework [17] for comparing elements of collections. Any `equals` method must satisfy a set of properties, such as implementing an equivalence relation; otherwise, the collections do not behave as expected. However, getting `equals` methods right is surprisingly hard.

We present VALloy, an extension to Alloy that enables intuitive modeling of dynamic dispatch. VALloy introduces in Alloy virtual functions and related inheritance constructs. We give VALloy a formal semantics through a translation to Alloy. The translation is similar to compilation of object-oriented languages, involving creation of virtual function tables. Since VALloy models can be automatically translated to Alloy, they can also be automatically analyzed using the existing AA.

Having an easy way to model dynamic dispatch is important for several reasons. First, it enables automatic analysis of models of overridden methods. Second, it allows modeling comparisons based on object values and developing specifications for collections that use these comparisons, such as Java collections. Third, such specifications can be used to test the actual implementations, for example using the TestEra framework [12].

The rest of this paper is organized as follows. Section 2 gives an example that illustrates the key constructs of VALloy. Section 3 defines a semantics for VALloy through a translation to Alloy. Section 4 presents VALloy specifications that partially model Java-like collections. Section 5 reviews related work. Section 6 discusses future work, and Section 7 presents our conclusions.

2 Example

We illustrate VALloy by modeling and analyzing an (in)correct overriding of the `equals` method in Java. We first develop an Alloy specification that contains only one `equals` method, and then describe challenges that arise in modeling method overriding. Finally, we present how VALloy tackles these challenges.

2.1 Modeling equals in Alloy

Consider the following `equals` method that appears in `java.awt.Dimension` in the standard Java libraries [17]:

```
class Dimension {
    int width;
    int height;
    public boolean equals(Object obj) {
        if (!(obj instanceof Dimension))
            return false;
        Dimension d = (Dimension)obj;
        return (width == d.width) && (height == d.height);
    }
}
```

We develop in Alloy (not yet VALloy) a specification for the above method. An Alloy specification consists of a sequence of paragraphs that either introduce an *uninterpreted type* or express constraints over the types. We start with the following declarations:

```
sig Object {} // java.lang.Object
sig Dimension extends Object { // java.awt.Dimension
  width: Integer,
  height: Integer
}
```

Each *signature*, introduced by the keyword `sig`, denotes a set of atomic individuals. In this specification, atoms in `sig Object` model Java objects. The signature `Dimension` is declared to be a subset of `Object`. Alloy subsets model Java subclassing with typing rules being as follows.

Signatures declared without `extends` are *basic* signatures. Basic signatures are disjoint from one another and represent Alloy types. Subsets do not introduce new Alloy types. The type of an atom is the basic signature it belongs to; all atoms in the above specification, including those in `Dimension`, have Alloy type `Object`. We can reconstruct Java type, i.e., class, of modeled Java objects based on their (sub)set membership.

Fields `width` and `height` introduce relations between `Dimension` atoms and `Integer` atoms, where `Integer` is predefined in Alloy. More precisely, each field introduces a function that maps `Dimension` atoms to `Integer` atoms.

We next add to the specification a model of the above `equals` method:

```
fun Dimension::equals(obj: Object) {
  obj in Dimension // instanceof
  this.width = obj.width && this.height = obj.height
}
```

The Alloy *function* `equals` records constraints that can be invoked elsewhere in the specification. This function has two arguments: `obj` and the implicit `this` argument, introduced with `'::'`. The function body constrains `obj` to be an atom of `Dimension`, effectively modeling Java's `instanceof`. This constraint is conjoined with the other that requires the fields to be the same. However, the above declaration does not constrain `this` to be an atom of `Dimension`; the declaration is equivalent to `fun Object::equals(obj: Object)`.

We next use the Alloy Analyzer (AA) to automatically check properties of the above specification. Each `equals` method should satisfy a set of properties: implement an equivalence relation and be consistent with `hashCode` [17]. The following Alloy assertion requires the function `equals`, which models the method `equals`, to be an equivalence relation:

```
assert equalsIsEquivalence {
  all o: Object | // reflexivity
    o..equals(o)
  all o1, o2: Object | // symmetry
    o1..equals(o2) => o2..equals(o1)
  all o1, o2, o3: Object | // transitivity
    o1..equals(o2) && o2..equals(o3) => o1..equals(o3)
}
```

The operator `'..'` invokes Alloy functions (using static resolution). AA checks the above assertion and reports that there are no counterexamples.

2.2 Overriding

Consider `Dimension3D`, a subclass of `java.awt.Dimension` that adds a field `depth` and overrides `equals`:

```
class Dimension3D extends java.awt.Dimension {
    int depth;
    boolean equals(Object obj) {
        if (!(obj instanceof Dimension3D))
            return false;
        Dimension3D d = (Dimension3D)obj;
        return super.equals(obj) && depth = d.depth;
    }
}
```

In order to check the `equals` method in `Dimension3D`, we would like to add the following to the Alloy specification presented so far:

```
sig Dimension3D extends Dimension {
    depth: Integer
}
// duplicate function names are NOT allowed in Alloy
fun Dimension3D::equals(obj: Object) {
    obj in Dimension3D
    // super.equals needs to be inlined because
    // there is no built in support for super
    this.width = obj.width && this.height = obj.height
    this.depth = obj.depth
}
```

However, this does not produce the intended model of overriding. In fact, this is not even a legal Alloy specification—each Alloy specification must have unique function names.¹ We could try renaming one of the `equals` functions, but it does not directly solve the problem of modeling overriding. Namely, the invocations `o.equals(o')` should choose the function based on the Java type/class of `o`. Since Alloy has no built in support for dynamic dispatch, we would need to model it manually for each function. Instead, we propose that it be done automatically.

2.3 Modeling equals in VALloy

VALloy introduces a natural way to model dynamic dispatch in Alloy. The following VALloy specification models the above Java classes:

```
class Object {}
virtual fun Object::equals(obj: Object) { this = obj }

class Dimension {
    width: Integer,
    height: Integer
}
virtual fun Dimension::equals(obj: Object) {
    obj in Dimension
    this.width = obj.width && this.height = obj.height
}

class Dimension3D extends Dimension {
    depth: Integer
}
virtual fun Dimension3D::equals(obj: Object) {
    obj in Dimension3D
    super..equals(obj) && this.depth = obj.depth
}
```

¹ That is why we do not initially add `equals` function for `Object`.

The `class` declaration in VALloy corresponds to the Alloy declaration `disj sig`, where `disj` indicates that the declared subset is disjoint from other `disj` subsets of its parent set. As in Java, VALloy classes by default extend `Object`.

The `virtual` function modifier² is the main VALloy extension to Alloy. This modifier declares a function that is dynamically dispatched at invocation, based on the VALloy class of the receiver. VALloy allows virtual functions to have the same name. The above example also shows the keyword `super` that VALloy provides for modeling `super` as found in Java.

2.4 Checking VALloy specifications

Every VALloy specification can be automatically translated into an Alloy specification. Section 3 presents the translation and the resulting Alloy specification for our running example.³

We use AA to automatically check the above assertion `equalsIsEquivalence`. Note that the invocations in the assertion do not need to change; the translation properly models dynamic dispatch. AA generates a counterexample:

```
Object_2: Dimension3D {
  width = 0,
  height = 1,
  depth = 3
}
Object_1: Dimension {
  width = 0,
  height = 1
}
```

These two objects violate the symmetry property: `Object_1.equals(Object_2)`, but *not* `Object_2.equals(Object_1)`. This is because `equals` of `Dimension` is oblivious of the field `depth` declared in `Dimension3D`. This counterexample shows that it is hard to extend the `java.awt.Dimension` class and preserve the properties of `equals`.

A way to provide an overridable implementation of `equals` in Java is to use the `getClass` method instead of the `instanceof` primitive [14]. In the running example, it requires changing `equals` of `java.awt.Dimension` to use the expression `obj.getClass() == this.getClass()` instead of `obj instanceof Dimension`. A similar change should be made in `Dimension3D`, unless it is declared `final`, and therefore cannot be extended.

Modeling this change in VALloy is straightforward: change `obj` in `Dimension` with `obj.getClass() = this.getClass()` in the function `Dimension::equals`. VALloy provides the function `getClass` that models the `final` method `getClass` from the class `java.lang.Object`. We translate the changed VALloy specification into Alloy and again use AA to check the equivalence assertion. This time AA reports that there are no counterexamples.

² VALloy borrows the modifier name from C++.

³ We have not yet implemented the translation; we perform it manually.

3 VALloy

This section presents VALloy as an extension to Alloy. We define a formal semantics for VALloy by giving a translation of VALloy specifications to Alloy specifications. Details of Alloy semantics can be found in [6].

VALloy adds the following to Alloy:

- `virtual` function modifier that declares a function whose invocation depends on the class of the receiver;
- `class` declaration that introduces VALloy classes;
- `super` keyword that directly correspond to Java;
- `getClass` function that corresponds to the `getClass` method of the class `java.lang.Object`.

These constructs are syntactically added to Alloy in the obvious way.

3.1 Translation Example

We give a semantics to the new constructs through a translation into Alloy. The translation algorithm operates in six steps, which we first describe through examples. Figures 1 and 2 show the Java code and VALloy specification from Section 2. For this example, the translation proceeds as follows.

Step 1. Compute the hierarchy of `class` declarations:

```
Object
+-- Dimension
    +-- Dimension3D
```

Step 2. Construct `sig Class` and `sig Object` based on the above hierarchy:

```
sig Class { ext: option Class }
static part sig Object_Class, Dimension_Class,
    Dimension3D_Class extends Class {}
fact Hierarchy {
  no Object_Class.ext
  Dimension_Class.ext = Object_Class
  Dimension3D_Class.ext = Dimension_Class
}
sig Object { class: Class }
fact ObjectClasses {
  (Object - Dimension).class = Object_Class
  (Dimension - Dimension3D).class = Dimension_Class
  Dimension3D.class = Dimension3D_Class
}
fun Object::getClass(): Class { result = this.class }
```

Atoms in `Class` and the *fact* `Hierarchy` represent the VALloy `class` declarations. (A *fact* in Alloy expresses constraints that must hold for all instances of the specification.) For each atom `c` in `Class`, `c.ext` gives the `Class` atom that corresponds to the superclass of `c`.⁴ The keyword `static` constrains each of the

⁴ For simplicity, we only present single inheritance, where the hierarchy can only be a tree. In multiple inheritance, each class can have a set of superclasses.

```

class Object {
    boolean equals(obj: Object) {
        return this == obj;
    }
}
class Dimension {
    int width;
    int height;
    boolean equals(obj: Object) {
        if (obj.getClass() != this.getClass())
            return false;
        Dimension d = (Dimension)obj;
        return width == d.width &&
            height == d.height;
    }
}
class Dimension3D extends Dimension {
    int depth;
    boolean equals(obj: Object) {
        if (obj.getClass() != this.getClass())
            return false;
        Dimension3D d = (Dimension3D)obj;
        return super.equals(obj) &&
            depth == d.depth;
    }
}

```

Fig. 1. Java code

```

class Object {}
virtual fun Object::equals(obj: Object) {
    this = obj
}

class Dimension {
    width: Integer,
    height: Integer
}
virtual fun Dimension::equals(obj: Object) {
    obj.getClass() = this.getClass()
    this.width = obj.width
    this.height = obj.height
}

class Dimension3D extends Dimension {
    depth: Integer
}
virtual fun Dimension3D::equals(obj: Object) {
    obj.getClass() = this.getClass()
    super.equals(obj)
    this.depth = obj.depth
}

assert equalsIsEquivalence {
    all o: Object | // reflexivity
        o.equals(o)
    all o1, o2: Object | // symmetry
        o1.equals(o2) => o2.equals(o1)
    all o1, o2, o3: Object | // transitivity
        o1.equals(o2) && o2.equals(o3) =>
            o1.equals(o3)
}

```

Fig. 2. Valloy specification

```

sig Class { ext: option Class }
static part sig
    Object_Class, Dimension_Class,
    Dimension3D_Class extends Class {}
fact Hierarchy {
    no Object_Class.ext
    Dimension_Class.ext = Object_Class
    Dimension3D_Class.ext = Dimension_Class
}
sig Object { class: Class }
fact ObjectClasses {
    (Object - Dimension).class = Object_Class
    (Dimension - Dimension3D).class =
        Dimension_Class
    Dimension3D.class = Dimension3D_Class
}

fun Object::getClass(): Class {
    result = this.class
}

fun Object::equals(obj: Object) {
    this.class = Object_Class =>
        this..Object_equals(obj)
    this.class = Dimension_Class =>
        this..Dimension_equals(obj)
    this.class = Dimension3D_Class =>
        this..Dimension3D_equals(obj)
}

```

```

fun Object::Object_equals(obj: Object) {
    this = obj
}

disj sig Dimension extends Object {
    width: Integer,
    height: Integer
}
fun Object::Dimension_equals(obj: Object) {
    obj.getClass() = this.getClass()
    this.width = obj.width
    this.height = obj.height
}

disj sig Dimension3D extends Dimension {
    depth: Integer
}
fun Object::Dimension3D_equals(obj: Object) {
    obj.getClass() = this.getClass()
    this..Dimension_equals(obj)
    this.depth = obj.depth
}

assert equalsIsEquivalence {
    all o: Object | // reflexivity
        o..equals(o)
    all o1, o2: Object | // symmetry
        o1..equals(o2) => o2..equals(o1)
    all o1, o2, o3: Object | // transitivity
        o1..equals(o2) && o2..equals(o3) =>
            o1..equals(o3)
}

```

Fig. 3. Translated Alloy specification

declared subsets to contain exactly one atom, and the keyword `part` declares a partition—the subsets are disjoint and their union is the whole set.

For each atom `o` in `Object`, `o.class` gives the corresponding `Class` atom. This correspondence is set with `fact ObjectClasses` based on the VALloy `class` hierarchy. (The ‘-’ operator denotes set difference in Alloy.) This translation step also introduces the function `getClass`.

Step 3. Change `class` declarations into `disj sig` declarations, adding `extends Object` where required:

```
disj sig Dimension extends Object { ... }
disj sig Dimension3D extends Dimension { ... }
```

This step does not change field declarations.⁵

Step 4. Rename each virtual function so that all functions in the specification have unique names:

```
fun Object::Object_equals(obj: Object) { this = obj }
fun Object::Dimension_equals(obj: Object) { ... }
fun Object::Dimension3D_equals(obj: Object) { ... }
```

This step also removes the modifier `virtual`, translating dynamically dispatched VALloy functions into statically dispatched Alloy functions.

Step 5. Add, for each overridden function name, a *dispatching* function, i.e., a new Alloy function that models dynamic dispatch:

```
fun Object::equals(obj: Object) {
  this.class = Object_Class =>
  this..Object_equals(obj)
  this.class = Dimension_Class =>
  this..Dimension_equals(obj)
  this.class = Dimension3D_Class =>
  this..Dimension3D_equals(obj)
}
```

This step is the crux of the translation. It allows function invocations in VALloy to be written in the usual Alloy notation, but it models dynamic dispatch semantics—the actual function is selected based on the class of the receiver.

Step 6. Replace each invocation on `super` with an invocation to the corresponding, previously renamed, static function:

```
fun Object::Dimension3D_equals(obj: Object) {
  obj..getClass() = this..getClass()
  this..Dimension_equals(obj) && this.depth = obj.depth
}
```

This completes the translation. Figure 3 shows the full resulting Alloy specification. Note that the translation does not change the assertion; the invocations `o..equals(o’)` remain written in the most intuitive manner, but they have dynamic dispatch semantics.

⁵ For simplicity, we do not present modeling `null`, which would require slightly changing field declarations.

3.2 General class hierarchy

To illustrate the general translation of `class` hierarchy, consider the following excerpt from a VALloy specification:

```
class O {}          virtual fun O::hC() { /*0*/ }
class C extends O {} virtual fun C::hC() { /*C*/ }
class C1 extends C {} // C1 does not define fun hC
class C2 extends C {} virtual fun C2::hC() { /*C2*/ }
class D extends O {} // D does not define fun hC
class D1 extends D {} virtual fun D1::hC() { /*D1*/ }
```

For this hierarchy, the translation generates the following `sig Class` and `sig Object`:

```
sig Class { ext: option Class }
static part sig O_Class, C_Class, C1_Class, C2_Class,
              D_Class, D1_Class extends Class {}

fact Hierarchy {
  no O_Class.ext
  C_Class.ext = O_Class
  C1_Class.ext = C_Class
  C2_Class.ext = C_Class
  D_Class.ext = O_Class
  D1_Class.ext = D_Class
}

sig Object { class: Class }
fact ObjectClasses {
  (O - C - D).class = O_Class
  (C - C1 - C2).class = C_Class
  C1.class = C1_Class
  C2.class = C2_Class
  (D - D1).class = D_Class
  D1.class = D1_Class
}
```

For the function `hC`, the translation generates the following Alloy functions:

```
fun O::O_hC() { /*0*/ }
fun O::C_hC() { /*C*/ }
// there is no O::C1_hC()
fun O::C2_hC() { /*C2*/ }
// there is no O::D_hC()
fun O::D1_hC() { /*D1*/ }
fun O::hC() {
  this.class = O_Class => this..O_hC()
  this.class = C_Class => this..C_hC()
  this.class = C1_Class => this..C_hC() /* not C1 */
  this.class = C2_Class => this..C2_hC()
  this.class = D_Class => this..O_hC() /* not D */
  this.class = D1_Class => this..D1_hC()
}
```

3.3 Summary

To summarize, the translation from VALloy to Alloy proceeds in the following six steps:

1. Compute the hierarchy of `class` declarations.
2. Construct `sig Class` and `sig Object`.
3. Change `class` into `disj sig` declarations.
4. Rename uniquely each virtual function.
5. Add dispatching functions.
6. Replace `super` with an appropriate static invocation.

4 Collections

This section presents VALloy models for some collection classes. Our main focus is comparison based on object values. We ignore the orthogonal issue of modeling state, i.e., sharing and object interactions. An approach for modeling state in Alloy is discussed in [4], and we can apply the same approach to VALloy.

4.1 Sets

We develop a VALloy specification for sets whose membership is based on object values, not object identities. As in Java, elements of the sets are objects of classes that (in)directly extend `Object` and potentially override `equals`.

We first declare a VALloy class for sets:

```
class Set { s: set Object }
```

For each atom `a` in `Set`, `a.s` is the (Alloy) set of objects in the (modeled) set `a`. To constrain set membership to be based on object values, we introduce the following fact:

```
fact SetIsBasedOnEquals {
  all a: Set | all disj e1, e2: a.s | !e1..equals(e2)
}
```

This fact requires distinct elements in each set to be not equal with respect to `equals`. For example, this rules out the set `a` such that `a.s={d1,d2}`, where `d1` and `d2` are distinct atoms (i.e., `d1!=d2`) of `Dimension`, but `d1.width=3`, `d1.height=8` and also `d2.width=3`, `d2.height=8`, which makes `d1..equals(d2)`. Note that `a.s` is a valid Alloy set.

It is now easy to specify some set functions from the `java.util.Set` interface:

```
virtual fun Set::contains(o: Object) {
  some e: this.s | o..equals(e)
}
virtual fun Set::add(o: Object): Set {
  this.s..contains(o) =>
  result.s = this.s,
  result.s = this.s + o
}
virtual fun Set::remove(o: Object): Set {
  result.s = this.s - { e: this.s | e..equals(o) }
}
virtual fun Set::isEmpty() { no this.s }
virtual fun Set::clear(): Set { no result.s }
virtual fun Set::size(): Integer { result = #this.s }
virtual fun Set::subset(a: Set) {
  all e: this.s | a..contains(e)
}
virtual fun Set::equals(o: Object) {
  o in Set
  o..size() = this..size()
  o..subset(this)
}
```

The most interesting function is `equals`, which compares two sets for equality. It checks that both sets have the same number of elements and that `o` is a subset (based on `equals`) of `this`. The function `remove` uses set comprehension to specify an object's removal from a set.

The above Valloy specification closely models `java.util.Set`. The main difference is that this specification is written in a functional style and does not model state modifications. As mentioned, state can be modeled using the approach from [4], which also presents a way to handle iterators. Therefore, we do not model “bulk operations” on sets, such as `addAll`, based on iterators. Instead, we present an analogous function for set union:

```
virtual fun Set::union(a: Set): Set {
    this..subset(result)
    a..subset(result)
    all e: result.s | this..contains(e) || a..contains(e)
}
```

Note that the use of `contains` (and `subset` based on `contains`), which is based on `equals`, enables specifying `union` in a direct way.

4.2 Maps

We next develop a partial Valloy specification for maps, such as `java.util.Map`, that compare keys based on `equals`. In this specification, we reuse the `class Set` defined above to automatically constrain the set of keys:

```
class Map {
    keys: Set
    map: keys.s ->! Object
}
```

We model the mapping from keys to values using an Alloy relation `map`; the *multiplicity marking* ‘!’ indicates that for each key, there is exactly one `Object`. For each atom `a` in `class Map`, `a.map` is the actual mapping. For a key `k`, `a.map[k]` gives the value that `k` is mapped to in map `a`.

We next model the essential map functions:

```
virtual fun Map::get(key: Object): Object {
    this.keys..contains(key) => result = this.map[key]
}
virtual fun Map::put(key: Object, value: Object): Map {
    result.keys = (this.key - { e: this.keys | e..equals(key) }) + key,
    result.map = (this.map - { e: this.keys | e..equals(key) }->Object) + key->value
}
```

The function `get` returns the value that `key` is mapped to, if such a key exists in the map; otherwise, the behavior of `get` is unspecified. (Since Alloy is a relational language, non-determinism comes for free.) We can constrain `get` to be deterministic, e.g., to return an explicit `Null` object, if the key is not in the map.

4.3 Trees

We next use a tree-based implementation of sets to illustrate how properties of abstract data types can be expressed in Valloy. Consider the following declaration for binary trees:

```
class Tree { root: Node }
class Node {
    left: Node,
    right: Node,
    data: Object
}
```

Suppose that these VALloy trees model a Java implementation of sets based on `equals`. We can state the *abstraction function* [19] for these trees in VALloy:

```
fun Tree::abstractionFunction(): Set {
  result.s = this.root.*(left+right).data
}
```

The ‘`*`’ operator is reflexive transitive closure, and `root.*(left+right)` denotes an (Alloy) set of all `Nodes` reachable from the `root`. The set of `Objects` from those nodes is obtained accessing `data`, and the abstraction function constrains this Alloy set to be a `Set`.

We also use VALloy to state *representation invariant* [19] for these trees. Assume that they have the following structural constraints: root nodes are sentinels (and thus never `null`) and leaf nodes point to themselves. The following `repOk` predicate characterizes the representation invariants for a tree:

```
fun Tree::repOk() {
  // no node points to root
  no this.root.^(left+right)
  // acyclic (with self loops for leaves)
  all n: this.root.*(left+right) {
    n.left = n || n !in n.left.*(left+right)
    n.right = n || n !in n.right.*(left+right)
  }
  // no duplicates w.r.t equals()
  some a: Set | a = this..abstractionFunction()
}
```

(The ‘`~`’ operator denotes transpose of a binary Alloy relation.) Beside the structural invariants, a valid tree is required to be a concrete representation of some `Set`. Note how the `equals` constraints from the abstract representation, `Set`, propagate to the concrete representation, `Tree`.

5 Related work

Recently, Jackson and Fekete [4] proposed an approach for modeling in Alloy object interactions, like those in Java. Their approach models heap using explicit references and captures properties of object sharing and aliasing. However, the approach does not handle inheritance in the presence of method overriding and dynamic dispatch. Their approach is orthogonal to our handling of virtual functions; we are planning to combine these two approaches.

Alloy has been used to check properties of programs that manipulate dynamic data structures. Jackson and Vaziri [7] developed a technique for analyzing bounded segments of procedures that manipulate linked lists. Their technique automatically builds an Alloy model of computation and checks it against a specification. They consider a small subset of Java, without dynamic dispatch.

We developed TestEra [12], a framework for automated testing of Java programs. In TestEra, specifications are written in Alloy and the Alloy Analyzer is used to provide automatic test case generation and correctness evaluation of programs. Writing specifications for Java collections, which use comparisons based on object values, requires modeling the `equals` method in Alloy. This led us

to tackle modeling general Java-like inheritance in Alloy. VALloy presents some ideas toward that goal.

The Java Modeling Language (JML) [10] is a popular specification language for Java. JML assertions use Java syntax and semantics, with some additional constructs, most notably for quantification. Leveraging on Java, JML specifications can obviously express dynamic dispatch. However, JML lacks static tools for automatic verification of such specifications.

The LOOP project [18] models inheritance in higher order logic to reason about Java classes. Java classes and their JML specifications are compiled into logical theories in higher order logic. A theorem prover is used to verify the desired properties. This framework has been used to verify that the methods of `java.util.Vector` maintain the safety property that the actual size of a vector is less than or equal to its capacity [3].

Object-oriented paradigm has been integrated into many existing languages, typically to make reuse easier. For example, Object-Z [15] extends the Z specification language [16], which enables building specifications in an object-oriented style. Object-Z retains the syntax and semantics of Z, adding new constructs. The major new construct is the class schema that captures the object-oriented notion of a class. Object-Z allows inheritance to be modeled, but it lacks tool support for automatically analyzing specifications.

Objects and inheritance have also been added to declarative languages. For example, Prolog++ [13] extends Prolog. OOLP+ [1] aims to integrate object-oriented paradigm with logic programming by translating OOLP+ code into Prolog without meta-interpretation.

Keidar et al. [8] add inheritance to the IOA language [11] for modeling state machines, which enables reusing simulation proofs between state machines. This approach allows only a limited form of inheritance, subclassing for extension: subclasses can add new methods and *specialize* inherited methods, but they cannot override those inherited methods, changing their behavior arbitrarily. VALloy allows subclasses to arbitrarily change the behavior of inherited methods.

6 Future work

VALloy presents our first step towards modeling in Alloy advanced constructs from object-oriented languages. The main focus has been on method overriding in Java. We have therefore designed VALloy to support subclasses that can arbitrarily change behavior of inherited methods.

Our approach can easily be extended to support intuitive modeling of multiple inheritance, such as in C++, and multi-method dispatch, such as in Cecil. Support for method overloading can clearly be added through a simple syntactic manipulation. We omitted support for Java’s interfaces, keeping in line with Alloy’s “micromodularity” philosophy of being a lightweight language amenable to fully automatic analysis. Similarly, we do not consider encapsulation.

We would like to explore modeling in VALloy other constructs of Java. Initially, we plan to add state to VALloy, adapting the approach of [4], or maybe

TAlloy [9]. We also want to add support for exceptions to enable VALloy specifications to directly model exceptional behavior of methods, as well as normal behavior. Further, having exceptions would allow modeling arrays with bound checking. We are also evaluating whether VALloy specifications should include computation, such as control flow, recursion, and multi-threading.

To explore practical value of VALloy, we intend to use it in connection with some existing frameworks. Daikon [2] is a tool for dynamically detecting likely program invariants; we are considering to use it to detect (partial) VALloy specifications of Java classes. TestEra [12] is a framework for automated test generation and correctness evaluation of Java classes; we are considering to use VALloy specifications for TestEra.

7 Conclusions

We described VALloy, an extension to the first order, relational language Alloy. All function invocations in Alloy are static; Alloy has no direct support for dynamic dispatch. VALloy introduces virtual functions in Alloy, which enables intuitive modeling of inheritance, such as that of Java. We illustrated the use of VALloy by modeling a small part of the Java Collections Framework.

We defined a formal semantics for VALloy through a translation to Alloy. VALloy models can be automatically translated into Alloy. The translation is similar to building virtual function tables for object-oriented languages and can benefit from optimizations based on class hierarchy. The translated specifications can be automatically checked using the existing Alloy Analyzer. We believe that VALloy can be effectively used for specification and checking of Java classes.

Acknowledgements

We would like to thank Manu Sridharan and Ang-Chih Kao for comments on an earlier draft of this paper. This work was funded in part by ITR grant #0086154 from the National Science Foundation.

References

1. Mukesh Dalal and Dipayan Gangopahyay. OOLP: A translation approach to object-oriented logic programming. In *Proc. First International Conference on Deductive and Object-Oriented Databases (DOOD-89)*, pages 555–568, Kyoto, Japan, December 1989.
2. Michael D. Ernst. *Dynamically Discovering Likely Program Invariants*. PhD thesis, University of Washington Department of Computer Science and Engineering, Seattle, Washington, August 2000.
3. Marieke Huisman, Bart Jacobs, and Joachim van den Berg. A case study in class library verification: Java’s Vector class. *Software Tools for Technology Transfer*, 2001. (to appear).

4. Daniel Jackson and Alan Fekete. Lightweight analysis of object interactions. In *Proc. Fourth International Symposium on Theoretical Aspects of Computer Software*, Sendai, Japan, October 2001.
5. Daniel Jackson, Ian Schechter, and Ilya Shlyakhter. ALCOA: The Alloy constraint analyzer. In *Proc. 22nd International Conference on Software Engineering (ICSE)*, Limerick, Ireland, June 2000.
6. Daniel Jackson, Ilya Shlyakhter, and Manu Sridharan. A micromodularity mechanism. In *Proc. 9th ACM SIGSOFT Symposium on the Foundations of Software Engineering (FSE)*, Vienna, Austria, September 2001.
7. Daniel Jackson and Mandana Vaziri. Finding bugs with a constraint solver. In *Proc. International Symposium on Software Testing and Analysis (ISSTA)*, Portland, OR, August 2000.
8. Idit Keidar, Roger Khazan, Nancy Lynch, and Alex Shvartsman. An inheritance-based technique for building simulation proofs incrementally. In *Proc. 22nd International Conference on Software Engineering (ICSE)*, pages 478–487, Limerick, Ireland, June 2000.
9. Sarfraz Khurshid and Darko Marinov. TAlloy: Multiple instances in a single Alloy model. (unpublished extended abstract), May 2001.
10. Gary T. Leavens, Albert L. Baker, and Clyde Ruby. Preliminary design of JML: A behavioral interface specification language for Java. Technical Report TR 98-06i, Department of Computer Science, Iowa State University, June 1998. (last revision: Aug 2001).
11. Nancy Lynch. *Distributed Algorithms*. Morgan Kaufmann Publishers, 1996.
12. Darko Marinov and Sarfraz Khurshid. TestEra: A novel framework for automated testing of Java programs. In *Proc. 16th IEEE International Conference on Automated Software Engineering (ASE)*, San Diego, CA, November 2001.
13. Chris Moss. *Prolog++ The Power of Object-Oriented and Logic Programming*. Addison-Wesley, 1994.
14. Mark Roulo. How to avoid traps and correctly override methods from `java.lang.Object`. <http://www.javaworld.com/javaworld/jw-01-1999/jw-01-object.html>.
15. G. Smith. *The Object-Z Specification Language*. Kluwer Academic Publishers, 2000.
16. J. M. Spivey. *The Z Notation: A Reference Manual*. Prentice Hall, second edition, 1992.
17. Sun Microsystems. *Java 2 Platform, Standard Edition, v1.3.1 API Specification*. <http://java.sun.com/j2se/1.3/docs/api/>.
18. Joachim van den Berg and Bart Jacobs. The LOOP compiler for Java and JML. In *Proc. Tools and Algorithms for the Construction and Analysis of Software (TACAS)*, (*Springer LNCS 2031, 2001*), pages 299–312, Genoa, Italy, April 2001.
19. Barbara Liskov with John Guttag. *Program Development in Java: Abstraction, Specification, and Object-Oriented Design*. Addison-Wesley, 2000.