

A Specification-Based Test Case Generation Method for UML/OCL

Achim D. Brucker¹, Matthias P. Krieger^{2,3}, Delphine Longuet^{2,3},
and Burkhard Wolff^{2,3}

¹ SAP Research, Vincenz-Priessnitz-Str. 1, 76131 Karlsruhe, Germany
achim.brucker@sap.com

² Univ. Paris-Sud, LRI UMR8623, Orsay F-91405*

³ CNRS, Orsay F-91405

{krieger,longuet,wolff}@lri.fr

Abstract. Automated test data generation is an important method for the verification and validation of UML/OCL specifications. In this paper, we present an extension of DNF-based test case generation methods to class models and recursive query operations on them. A key feature of our approach is an implicit representation of object graphs avoiding a representation based on object-id's; thus, our approach avoids the generation of isomorphic object graphs by using a concise and still human-readable symbolic representation.

Keywords: OCL, UML, test case generation, specification-based testing.

1 Introduction

Automated test data generation is an important application domain for OCL specifications. Instead of verifying concrete code via a Hoare-Calculus for a specific programming language against OCL method contracts—a technique developed in detail for OCL in [9, 11]—test generation can be a more light-weighted (but logically less safe) formal method to reveal errors both in specification and implementations. A particular advantage of black-box testing is that implementations may consist of arbitrary mixtures of (dirty) programming languages.

In this paper, we will adapt existing specification-based testing techniques to UML/OCL, i. e., an object-oriented specification formalism centered around the concept of an object-graph as state, state-transitions described by class-models and state-charts (which we will ignore here), and a type system based on subtyping and inheritance. The work presented here is based on the previous work on a formal UML/OCL semantics [7, 11] and attempts to develop, in contrast to prior works such as [4], a *comprehensive* test-generation method for the complete language and for realistic test-scenarios. Overall, our contribution consists in:

1. the extension of specification-based test generation methods to the world of object-oriented specifications,

* This work was partially supported by the Digiteo Foundation.

2. a *deductive*, theorem-prover based test data generation from OCL specifications including language features such as recursive query-operations, and
3. a particular representation of object-graph classes by our novel concept of an *alias closure*; rather than representing the explicit object graphs we represent the object identity by an equivalence relation.

This paper is written with hindsight to the HOL-TESTGEN system [10], into which we will implement the technique presented here in a future step. This implies that our technique must fit to the underlying logical framework Isabelle/HOL (an embedding of UML/OCL has been presented in [11]), and that it can be organized into the generation phases of HOL-TESTGEN.

2 A Gentle Introduction to a Formal OCL 2.2 Semantics

In this section, we briefly present a formal semantics for OCL 2.2 [20], see [7] for details. With respect to the syntax, we use the mathematical notation of HOL-OCL [8] which allows for a concise presentation of OCL constraints.

2.1 Higher-Order Logic

Higher-order Logic (HOL) [12] is a classical logic with equality enriched by total parametrically polymorphic higher-order functions. It is more expressive than first-order logic, e. g., induction schemes can be expressed inside the logic. Pragmatically, HOL can be viewed as a typed functional programming language like Haskell extended by logical quantifiers.

HOL is based on the typed λ -calculus, i. e., the *terms* of HOL are λ -expressions. Types of terms may be built from *type variables* (like α, β, \dots , optionally annotated by Haskell-like *type classes* as in $\alpha :: \text{order}$ or $\alpha :: \text{bot}$) or *type constructors* (like `bool` or `nat`). Type constructors may have arguments (as in α list or α set). The type constructor for the function space \Rightarrow is written infix: $\alpha \Rightarrow \beta$; multiple applications like $\alpha_1 \Rightarrow (\dots \Rightarrow (\alpha_n \Rightarrow \alpha_{n+1}) \dots)$ have the alternative syntax $[\alpha_1, \dots, \alpha_n] \Rightarrow \alpha_{n+1}$. HOL is centered around the extensional logical equality $_ = _$ with type $[\alpha, \alpha] \Rightarrow \text{bool}$, where `bool` is the fundamental logical type. We use infix notation: instead of $(_ = _) E_1 E_2$ we write $E_1 = E_2$. The logical connectives $_ \wedge _, _ \vee _, _ \Rightarrow _$ of HOL have type $[\text{bool}, \text{bool}] \Rightarrow \text{bool}$, $\neg _$ has type $\text{bool} \Rightarrow \text{bool}$. The quantifiers $\forall _ . _$ and $\exists _ . _$ have type $[\alpha \Rightarrow \text{bool}] \Rightarrow \text{bool}$. The quantifiers may range over types of higher order, i. e., functions or sets.

The type discipline rules out paradoxes such as Russel’s paradox in untyped set theory. Sets of type α set can be defined isomorphic to functions of type $\alpha \Rightarrow \text{bool}$; the definition of the elementhood $_ \in _$, the set comprehension $\{ _ . _ \}$, $_ \cup _$ and $_ \cap _$ is then standard.

2.2 Valid Transitions and Evaluations

We recall that OCL expressions form a typed assertion language whose syntactic elements are composed of (a) operators on built-in data structures such

as Boolean or collection types like **Set** or **Bag**, (b) operators of the user-defined data-model such as attribute accessors, type-casts and tests, and (c) user-defined, potentially recursive, side-effect-free method calls.

The topmost goal of the formal semantics for OCL expressions is to define the notion of a *valid transition* over states of a system; even concepts like *object invariants* can be derived from this notion. Let σ be a pre-state and σ' a post-state and let ϕ be a Boolean OCL expression, then we write

$$(\sigma, \sigma') \models \phi$$

for “the transition from σ to σ' is valid in ϕ .” A formula ϕ is valid if and only if its evaluation in the transition $\tau = (\sigma, \sigma')$ yields true. As all types in HOL-OCL are extended by the special element \perp denoting undefinedness, we define formally:

$$\tau \models \phi \equiv (I[\llbracket \phi \rrbracket] \tau = \perp_{\text{true}}).$$

In OCL, the evaluation of all expressions can result in an undefinedness element called **invalid** which we will write \perp for short. The test for definedness (**not** $_.$ **oclIsInvalid()**) will be written $\partial _$ and is defined by $\partial X \equiv \text{not } (X \triangleq \perp)$. Here, $_ \triangleq _$ denotes the strong equality, which is a reflexive, symmetric and transitive congruence relation; therefore, the strong equality allows for substituting equals with equals in any OCL expression, even if the expressions are undefined. In contrast, the standard equality in OCL, i. e. $_ \doteq _$, is *strict*, which means $x \doteq \perp$ is strongly equal to $\perp \doteq x$ which is strongly equal to \perp .

Since all operators of the assertion language depend on the context τ and results can be \perp , all expressions can be viewed as *evaluations* from τ to a type α_{\perp} . All types of expressions are of a form captured by the type abbreviation:

$$V(\alpha) = \sigma \times \sigma \Rightarrow \alpha_{\perp},$$

where $\sigma \times \sigma$ stands for the type of a pair of system states (i. e., the type of τ).

2.3 Semantics of Object Invariants and Operation Contracts

The OCL semantics [20, Annex A] uses different interpretation functions for invariants and pre-conditions; instead, we achieve their semantic effect by a syntactic transformation $_.$ **pre** which replaces all accessor functions $_.$ **i** by their counterparts $_.$ **i@pre**. For example, $(\text{self}.\text{i} > 5)_{\text{pre}}$ is just $(\text{self}.\text{i}@pre > 5)$. The operation $_.$ **allInstances()** is also substituted by its **@pre** counterpart. Thus, we can re-formulate the semantics of the two OCL top-level constructs, invariant specification and method specification, as follows:

$$\begin{aligned} I[\llbracket \text{context } c : C \text{ inv } n : \phi(c) \rrbracket] \tau &\equiv \\ \tau \models (C.\text{allInstances()} \rightarrow \text{forall}(x | \phi(x))) \wedge & \quad (1) \\ \tau \models (C.\text{allInstances()} \rightarrow \text{forall}(x | \phi(x)))_{\text{pre}} & \end{aligned}$$

The standard forbids expressions containing **@pre** constructs in invariants or preconditions syntactically; thus, mixed forms cannot arise. Since operations

have strict semantics in OCL, we have to distinguish for a specification of an *op* with the arguments a_1, \dots, a_n the two cases where all arguments are defined (and *self* is non-null), or not. In the former case, a method call can be replaced by a *result* that satisfies the contract, in the latter case the argument is \perp :

$$\begin{aligned}
& I[\text{context } C :: \text{op}(a_1, \dots, a_n) : T \\
& \quad \text{pre } \phi(\text{self}, a_1, \dots, a_n) \\
& \quad \text{post } \psi(\text{self}, a_1, \dots, a_n, \text{result})] \tau \equiv \forall s, x_1, \dots, x_n. \\
& \quad \Delta(s, x_1, \dots, x_n) \wedge \tau \models \phi(s, x_1, \dots, x_n)_{\text{pre}} \\
& \quad \longrightarrow \tau \models \psi(s, x_1, \dots, x_n, s.\text{op}(x_1, \dots, x_n)) \\
& \quad \wedge \neg \Delta(s, x_1, \dots, x_n) \longrightarrow \tau \models s.\text{op}(x_1, \dots, x_n) \triangleq \perp
\end{aligned} \tag{2}$$

where $\Delta(s, x_1, \dots, x_n)$ is an abbreviation for $\tau \models s \not\equiv \text{null} \wedge \tau \models \partial s \wedge \tau \models \partial x_1 \wedge \dots \wedge \tau \models \partial x_n$. This definition captures the two cases: if the arguments of an operation are defined and, moreover, *self* is not `null`, the result of a method call must satisfy the specification; otherwise the operation will be strict and return invalid \perp . By these definitions an OCL specification, i. e., a sequence of invariant declarations and operation contracts, can be transformed into a set of (logically conjoined) statements which is called the *context* Γ_τ . The *theory* of an OCL specification is the set of all valid transitions $\tau \models \phi$ that can be derived from Γ_τ . For the logical connectives of OCL, a conventional Gentzen-style calculus for pairs of the form $\Gamma_\tau \vdash \phi$ can be developed that allows for inferring valid transitions from Γ_τ by deduction (cf. [11]). Due to the inclusion of arithmetic, any calculus for OCL is necessarily incomplete. It is straight-forward to extend our notion of context to multi-transition contexts such as:

$$\Gamma \equiv \{(\sigma, \sigma') \models \phi, (\sigma', \sigma'') \models \psi\}$$

such that we can reason over systems executing several transitions.

2.4 Strict Operations and Their Role in Reasoning

The OCL standard [20] defines most operations as strict, not just the special case of the strict equality $x \doteq y$ mentioned earlier. Overall, we have the rule

$$f(x_1, \dots, \perp, \dots, x_n) \triangleq \perp. \tag{3}$$

A notable exception from this rule are the logical connectives, which are a three-valued strong Kleene-logic; e. g., \perp **and** **false** \triangleq **false** and analogously \perp **or** **true** \triangleq **true**. Overall, using a three-valued logic is a burden if a simple compilation of OCL to standard automated theorem provers is envisaged. Looking at the wealth of tools (that are specialized for two-valued logics) like Kodkod [22] or Z3 [17], this is perceived as a major drawback of OCL by many.

The methodology of OCL (in particular the strictness of most operations and the fact that most OCL expressions, e. g., invariants, are, by definition, defined)

enforces that a reduction to a two-valued representation is always possible; it suffices to apply the case-distinction:

$$\tau \models \phi(\perp) \vee (\tau \models \partial E \wedge \tau \models \phi(E)) \quad (4)$$

exhaustively to all sub-expressions E (the $\tau \models \phi(\perp)$ -parts will either reduce quickly due to Fact 3 to $\tau \models \perp$ which is just false or again be subject to Fact 4). The result are formulae of the form: $\tau \models \partial E_1 \wedge \dots \wedge \tau \models \partial E_n \wedge \tau \models \phi$ or just $\Delta_\phi \wedge \tau \models \phi$ for short. In this form—called Δ -long-form—all implicit definednesses in a valid OCL-formula are made explicit. We call the process of constructing a Δ -long-form Δ -saturation. We do not distinguish between $\{\tau \models E_1 \wedge \tau \models E_2\} \cup \Gamma$ and $\{\tau \models E_1, \tau \models E_2\} \cup \Gamma$.

This process can be optimized: if we have, for example, as consequence of an invariant $\tau \models f(a) \doteq b$ in our context Γ (meaning that it holds and, thus, evaluates to true), we can infer that $\tau \models \partial f(a)$ and $\tau \models \partial b$. From there, we can further infer $\tau \models \partial a$ (if f is strict). The same holds for the common connective $\tau \models X$ **and** Y (but not for **_ or _**). Once that the implicit knowledge on definedness is established, rules of the following form can be applied:

$$\begin{array}{ll} \tau \models \mathbf{not} X = \neg \tau \models X & \text{if } \tau \models \partial X \\ \tau \models X \mathbf{ or } Y = \tau \models X \vee \tau \models Y & \text{if } \tau \models \partial X \text{ and } \tau \models \partial Y \\ \tau \models X \mathbf{ and } Y = \tau \models X \wedge \tau \models Y & \text{if } \tau \models \partial X \text{ and } \tau \models \partial Y \\ \tau \models X \doteq Y = \tau \models X \triangleq Y & \text{if } \tau \models \partial X \text{ and } \tau \models \partial Y \end{array}$$

By applying this form of equations to Γ and ϕ , we transform them into sets of judgments of the form $\tau \models \phi$, i. e., perfect two-valued statements that can be treated by conventional SMT solvers like Z3 (provided that we add an appropriate background theory that axiomatizes the basic operations of the OCL language).

3 Running Example: Linked Lists

In this section, we present a small UML/OCL specification that will serve as a running example for our test case generation technique. We will also discuss the translation of OCL into HOL and discuss the implicit invariants of this example.

3.1 Singly-Linked Lists

Fig. 1 illustrates our running example of a singly-linked list: the list stores integers as data and links between nodes are modeled by an association. As a node does not necessarily need to have a successor, the association end **next** has multiplicity 0..1. An invariant of the class states that the integers are stored in a descending order in the list. We specify an operation **insert** that adds an integer to the list. The postcondition of the **insert** operation states that the set of integers stored in the list in the post-state is the set of stored integers in the pre-state extended by the argument. For defining the set of integers stored in the list, we separately specify the recursive query operation **contents()**.




Fig. 1. A Singly-linked list specified in OCL (excerpt)

In the following, we will describe how to build Γ_τ from this OCL specification via the semantic definitions. We will add to Γ_τ semantic presentations of the specification which are already in a “massaged format” suitable for test case generation later. Since the transition is not changing in the rest of this paper, we will assume one global transition τ (understood to be relative to the specification of this example); we will drop the index and abbreviate $\tau \models \phi$ to just $\models \phi$. In our test case generation approach, we assume that all diagrammatic constraints over the class model are represented as OCL expressions (for details, see [15]). For example, associations are usually represented by collection-valued class attributes together with OCL constraints expressing the multiplicity.

3.2 Translating Invariants into Recursive HOL-Predicates

The example in Fig. 1 only includes one explicit invariant. The multiplicity constraints in the class model constitute invariants semantically. For our example, the multiplicity constraints could be expressed as follows in OCL:

```
inv: (next = null or next <> null) and i <> null
```

In the following, we will assume that attributes and arguments that have a basic datatype (e. g., `Integer`) have a multiplicity of `1..1`, i. e., they cannot be `null`. Thus we can simplify the invariant representing the multiplicity constraints to:

```
inv: (next = null or next <> null)
```

This simplification improves the readability of the formulae in this paper and is not a fundamental restriction of our approach.

For our purposes it will be convenient to convert invariants to recursive predicates and add them to Γ , paving the way for the exploration of input parameters by simply unfolding them rather than making them, based on Fact 1, lengthy arguments over `.allInstances()`. Of course, not any recursive predicate is consistent; however *these* recursive predicates can be derived from the invariants by using a greatest fixed-point construction and proving that the body of the invariant is monotone—the reader interested in the details is referred to HOL-OCL [9] where this is done automatically (albeit for OCL 2.0, i. e., without `null`):

```

 $\forall \mathbf{self}. \models \partial \mathbf{self} \wedge \models \mathbf{self} \neq \mathbf{null} \longrightarrow \models \text{inv}_{\text{Node}}(\mathbf{self})$ 
 $\iff \models \mathbf{self.next} \doteq \mathbf{null} \vee (\models \mathbf{self.next} \neq \mathbf{null}$ 
 $\wedge \models \mathbf{self.i} > \mathbf{self.next.i} \wedge \models \text{inv}_{\text{Node}}(\mathbf{self.next}))$ 

```

Additionally to this recursive predicate, we add to Γ the fact that any defined non-null object will satisfy this invariant:

```

 $\forall \mathbf{self}. \models \partial \mathbf{self} \wedge \models \mathbf{self} \neq \mathbf{null} \longrightarrow \models \text{inv}_{\text{Node}}(\mathbf{self})$ 

```

Our recursive definitions are a conjunction of the explicit invariant and the multiplicity constraints of our example; we used Δ -short form in order not to clutter up our presentation too much, i. e., facts like $\models \partial \mathbf{self.i}$ were omitted. The invariant $\text{inv}_{\text{Node}}@pre$ expresses well-formedness in a pre-state:

```

 $\forall \mathbf{self}. \models \partial \mathbf{self} \wedge \models \mathbf{self} \neq \mathbf{null}$ 
 $\longrightarrow \models \text{inv}_{\text{Node}}@pre(\mathbf{self}) \iff \models \mathbf{self.next@pre} \doteq \mathbf{null}$ 
 $\vee (\models \mathbf{self.next@pre} \neq \mathbf{null} \wedge \models \text{inv}_{\text{Node}}@pre(\mathbf{self.next@pre})$ 
 $\wedge \models \mathbf{self.i@pre} > \mathbf{self.next@pre.i@pre})$ 

```

3.3 Translating Contracts into HOL

Given the fact that $\models (\mathbf{true})_{pre}$ just collapses to \mathbf{true} , the formulae that we add to Γ is the straight-forward simplification of the semantics rule Fact 2:

```

 $\forall \mathbf{self}. \Delta(\mathbf{self}) \longrightarrow \models \mathbf{self.contents}() \triangleq$ 
 $\quad \text{if } \mathbf{self.next} \doteq \mathbf{null} \text{ then } \text{Set}\{i\}$ 
 $\quad \text{else } \mathbf{self.next.contents}() \text{->including}(i)$ 
 $\wedge \neg \Delta(\mathbf{self}) \longrightarrow \models \mathbf{self.contents}() \triangleq \perp$ 

```

where $\Delta(\mathbf{self})$ is a short-cut for $\models \partial \mathbf{self} \wedge \models \mathbf{self} \neq \mathbf{null}$. The variant for $\text{contents@pre}()$ looks as follows:

```

 $\forall \mathbf{self}. \Delta(\mathbf{self}) \longrightarrow \models \mathbf{self.contents@pre}() \triangleq$ 
 $\quad \text{if } \mathbf{self.next@pre} \doteq \mathbf{null} \text{ then } \text{Set}\{i\}$ 
 $\quad \text{else } \mathbf{self.next@pre.contents@pre}() \text{->including}(i)$ 
 $\wedge \neg \Delta(\mathbf{self}) \longrightarrow \models \mathbf{self.contents@pre}() \triangleq \perp$ 

```

4 Test Generation

In the specification-based testing, we are interested in testing the formula ϕ —called *test specification*—in the statement:

$$\Gamma \vdash \phi$$

instead of proving it (the \vdash is interpreted as implication). We are interested in test specifications which contain calls to operations $s.\text{op}(a_1, \dots, a_n)$; the core of the technique consists in selecting arguments consistent with the specification Γ

and the semantic rules for the operations of OCL, executing the implementation of `op` and checking if the result validates ϕ . We follow the classical approach of transforming the test specification into a disjunctive normal form (DNF), extended by invariant-handling and the treatment of recursive definitions, which corresponds to partitioning the input space of the operation(s).

Another class of case distinctions arises from *aliasing*; i. e., the fact that two object references can designate the same object, i. e., `s.next.next` is in fact identical to `s` due to a cycle in the object graph. Aliasing is a crucial phenomenon in object-oriented systems. It is likely that a system behaves differently depending on the aliasing relationships among the objects it handles. Therefore we will add further case distinctions to the specification under analysis that distinguish different aliasing relationships. We will refer to this transformation as *alias closure*.

4.1 Test Specifications: Getting Started

Depending on the specific test purposes, there are various ways to test a system: a test could be concerned with the normal behavior of operations, which will be the default considered here, or with exceptional behavior (what happens if the precondition is not satisfied?), or with operation sequences, e. g., $(\sigma, \sigma') \models \phi \wedge (\sigma', \sigma'') \models \psi$. It is even conceivable to express in test specifications the sharing of pre-state and post-state or different parameters; in our specification, the implementation of `insert` can have a copying semantics (all object-contents in the list were copied to freshly generated objects) as well as a sharing semantics.

Since OCL expressions cannot have side-effects, properties of non-query operations like `insert` cannot be expressed inside OCL. We therefore suggest to present the test specification directly in HOL. Thus, following Fact 2, we have for the case of a “normal behavior unit test”:

$$\begin{array}{l} \Delta(\mathbf{s}, \mathbf{x}) \longrightarrow \models \mathbf{s}.\text{contents}() \doteq \mathbf{s}.\text{contents@pre}() \rightarrow \text{including}(\mathbf{x}) \\ \wedge \neg\Delta(\mathbf{s}, \mathbf{x}) \longrightarrow \models \mathbf{s}.\text{insert}(\mathbf{x}) \triangleq \perp \end{array}$$

where \mathbf{s} is a free variable for which we look for solutions that meet all possible constraints (arising from the context Γ , but also locally in ϕ). Since $(\Delta \longrightarrow A) \wedge (\neg\Delta \longrightarrow B)$ is equivalent to $(\Delta \wedge A) \vee (\neg\Delta \wedge B)$ and the latter is closer to a DNF, we rewrite our test specification and have:

$$\begin{array}{l} \Delta(\mathbf{s}, \mathbf{x}) \wedge \models \mathbf{s}.\text{contents}() \doteq \mathbf{s}.\text{contents@pre}() \rightarrow \text{including}(\mathbf{x}) \\ \vee \neg\Delta(\mathbf{s}, \mathbf{x}) \wedge \models \mathbf{s}.\text{insert}(\mathbf{x}) \triangleq \perp \end{array}$$

which boils down to:

$$\begin{array}{l} \Delta(\mathbf{s}, \mathbf{x}) \wedge \models \mathbf{s}.\text{contents}() \doteq \mathbf{s}.\text{contents@pre}() \rightarrow \text{including}(\mathbf{x}) \\ \vee \models \mathbf{s} \doteq \text{null} \wedge \models \mathbf{s}.\text{insert}(\mathbf{x}) \triangleq \perp \\ \vee \models \mathbf{x} \doteq \perp \wedge \models \mathbf{s}.\text{insert}(\mathbf{x}) \triangleq \perp \end{array}$$

Here, we can already apply unit propagation in clauses and extract the two test cases: $\models \text{null}.\text{insert}(\mathbf{x}) \triangleq \perp$ and $\models \mathbf{s}.\text{insert}(\perp) \triangleq \perp$ which essentially test the corner-cases imposed by the semantics of OCL.

4.2 Test Hypotheses

The test cases $\models \text{null.insert}(x) \triangleq \perp$ and $\models \text{s.insert}(\perp) \triangleq \perp$ give also some deeper insight into testing. Test cases are *classes* of concrete tests, i. e., the ground instances (e. g., $\models \text{null.insert}(0) \triangleq \perp$ or $\models \text{null.insert}(1) \triangleq \perp$) of these formulae. Overall, a *test case* for an operation op in a DNF is a conjoint:

$$\models \phi_1(x_1, \dots, x_n) \wedge \dots \wedge \models \phi_n(x_1, \dots, x_n)$$

where at least one ϕ_i depends on op . In test cases, we can partition the clauses into two groups: in *oracles* O , i. e., those $\models \phi_i(x_1, \dots, x_n)$ that depend on op , and in *constraints* C , i. e., all others. Constructing a test boils down to finding a solution, i. e., a ground substitution for x_1, \dots, x_n , that satisfies all constraints in C , while the test proceeds by executing the implementation of op for this solution and check if the oracles evaluate to true.

Logically, this means that we made the assumption “if there is an input vector (x_1, \dots, x_n) satisfying all constraints, and if this input passes the oracle execution, the oracles will pass for all inputs satisfying the constraints.” This type of assumption underlying a test is called a *uniformity hypothesis* written:

$$\begin{aligned} (\exists x_1, \dots, x_n. C(x_1, \dots, x_n) \wedge O(x_1, \dots, x_n)) \\ \longrightarrow (\forall x_1, \dots, x_n. C(x_1, \dots, x_n) \longrightarrow O(x_1, \dots, x_n)) \end{aligned}$$

While this is the most fundamental testing hypothesis, there are other useful ones that help to establish case distinctions used in specification-based tests. A notable other well-known form of a testing hypothesis is the *regularity hypothesis*:

$$\begin{aligned} (\forall x_1, \dots, x_n. |x_1, \dots, x_n| < k \wedge C(x_1, \dots, x_n) \longrightarrow O(x_1, \dots, x_n)) \\ \longrightarrow (\forall x_1, \dots, x_n. C(x_1, \dots, x_n) \longrightarrow O(x_1, \dots, x_n)) \end{aligned}$$

or in other words: whenever we tested all data up to a given complexity measure (like size of collections) k , we assume that the execution of op satisfies the specification. In the context of OCL testing, there is a similar form of regularity hypothesis: here, we will implicitly argue over the bound k on the number of different objects in a state that has been used for the tests.

4.3 Unfolding

For generating a set of test cases, we start with the test specification given above, restricted to the part where \mathbf{s} is defined and not null.

$$\Delta(\mathbf{s}, \mathbf{x}) \wedge \models \mathbf{s.contents}() \doteq \mathbf{s.contents@pre}() \rightarrow \text{including}(\mathbf{x})$$

This test specification does not show any explicit case distinctions. Rather, the case distinctions are hidden in the recursive specification of `contents()`.

The invariants over the different arguments of the operation (including \mathbf{s}) must be taken into account for the generation of relevant test cases. In our example, only ordered lists can occur in pre-states and post-states of the `insert` operation.

Adding these invariants as constraints over the pre-states or post-states reduces the number of test cases derived from the test specification by removing as many non-satisfiable clauses as possible before the test data selection. Because of the facts contained in T , we obtain:

$$\forall \text{ self. } \models \partial \text{ self} \wedge \models \text{self} \neq \text{null} \longrightarrow \models \text{inv}_{\text{Node}}(\text{self})$$

These invariants can be inserted at any time during the unfolding process.

For instance, we can already insert the invariant for the pre-states and post-states of the `insert` operation, knowing that `s` is defined and not `null`:

$$\begin{aligned} & \Delta(s, x) \wedge \models \text{inv}_{\text{Node}}@\text{pre}(s) \wedge \models \text{inv}_{\text{Node}}(s) \\ & \wedge \models \text{s.contents}() \doteq \text{s.contents}@pre() \rightarrow \text{including}(x) \end{aligned}$$

To enrich this condition with explicit case distinctions, we unfold the operation calls and invariants by replacing them with their specification: an operation call will be replaced with its contract and an invariant with its definition, which is allowed here since we have $\Delta(s, x)$. For the sake of readability, we do not replace the `contents` operation calls directly with their contract but rather conjoin the contract with the existing formulae. We obtain the following conditions:

$$\begin{aligned} & \Delta(s, x) \\ & \wedge (\models \text{s.next}@pre \doteq \text{null} \\ & \quad \vee (\models \text{s.next}@pre \neq \text{null} \\ & \quad \quad \wedge \models \text{s.i}@pre > \text{s.next}@pre.i@pre \wedge \models \text{inv}_{\text{Node}}@\text{pre}(\text{s.next}@pre))) \\ & \wedge (\models \text{s.next} \doteq \text{null} \\ & \quad \vee (\models \text{s.next} \neq \text{null} \wedge \models \text{s.i} > \text{s.next.i} \wedge \models \text{inv}_{\text{Node}}(\text{s.next}))) \\ & \wedge \models \text{s.contents}() \doteq \text{s.contents}@pre() \rightarrow \text{including}(x) \\ & \wedge \models \text{s.contents}() \triangleq \text{if } \text{s.next} \doteq \text{null} \text{ then } \text{Set}\{\text{s.i}\} \\ & \quad \quad \quad \text{else } \text{s.next.contents}() \rightarrow \text{including}(\text{s.i}) \\ & \wedge \models \text{s.contents}@pre() \triangleq \\ & \quad \text{if } \text{s.next}@pre \doteq \text{null} \text{ then } \text{Set}\{\text{s.i}@pre\} \\ & \quad \quad \text{else } \text{s.next}@pre.contents@pre() \rightarrow \text{including}(\text{s.i}@pre) \end{aligned}$$

A second refinement step could be performed by unfolding the invariants and the operation calls a second time: we could insert the invariant definitions again and instantiate the operation contract for the `contents` operation with `s.next` (correspondingly for the pre-state).

The unfolding process and invariant insertion can be stopped at any time, once the refinement is sufficient according to the tester's needs. Then, the DNF of the obtained formula is generated to enumerate the different test cases coming from case distinction. The DNF obtained for the previous formula is the following, leading to four clauses distinguishing whether `s.next` and `s.next@pre` are `null`.

$$\begin{aligned} & (\Delta(s, x) \\ & \wedge \models \text{s.next} \doteq \text{null} \\ & \wedge \models \text{s.next}@pre \doteq \text{null} \\ & \wedge \models \text{s.contents}() \doteq \text{s.contents}@pre() \rightarrow \text{including}(x) \\ & \wedge \models \text{s.contents}() \triangleq \text{Set}\{\text{s.i}\} \\ & \wedge \models \text{s.contents}@pre() \triangleq \text{Set}\{\text{s.i}@pre\}) \end{aligned}$$

```

∨ (Δ(s, x)
  ∧ ⊨ s.next ≠ null ∧ ⊨ s.i > s.next.i ∧ ⊨ invNode(s.next)
  ∧ ⊨ s.next@pre ≐ null
  ∧ ⊨ s.contents() ≐ s.contents@pre()->including(x)
  ∧ ⊨ s.contents() ≐ s.next.contents()->including(s.i)
  ∧ ⊨ s.contents@pre() ≐ Set{s.i@pre})
∨ (Δ(s, x)
  ∧ ⊨ s.next ≐ null
  ∧ ⊨ s.next@pre ≠ null ∧ ⊨ s.i@pre > s.next@pre.i@pre
  ∧ ⊨ invNode@pre(s.next@pre)
  ∧ ⊨ s.contents() ≐ s.contents@pre()->including(x)
  ∧ ⊨ s.contents() ≐ Set{s.i}
  ∧ ⊨ s.contents@pre() ≐ s.next@pre.contents@pre()
  ->including(s.i@pre))
∨ (Δ(s, x)
  ∧ ⊨ s.next ≠ null ∧ ⊨ s.i > s.next.i ∧ ⊨ invNode(s.next)
  ∧ ⊨ s.next@pre ≠ null ∧ ⊨ s.i@pre > s.next@pre.i@pre
  ∧ ⊨ invNode@pre(s.next@pre)
  ∧ ⊨ s.contents() ≐ s.contents@pre()->including(x)
  ∧ ⊨ s.contents() ≐ s.next.contents()->including(s.i)
  ∧ ⊨ s.contents@pre() ≐ s.next@pre.contents@pre()
  ->including(s.i@pre))

```

The first case boils down (due to constant propagation and set reasoning) to:

```

Δ(s, x) ∧ ⊨ s.next ≐ null ∧ ⊨ s.next@pre ≐ null
∧ ⊨ s.i ≐ s.i@pre ∧ ⊨ s.i ≐ x

```

All other cases are not yet “ground” enough and contain application redexes like $\text{invNode}(s.\text{next})$ for further invariant unfolding. The derivation

```

Δ(s, x)
∧ ⊨ s.next ≠ null ∧ ⊨ s.i > s.next.i ∧ ⊨ invNode(s.next)
∧ ⊨ s.next@pre ≐ null
∧ ⊨ s.next.contents()->including(s.i) ≐ Set{s.i@pre}
  ->including(x)

```

for the second case expands to:

```

(Δ(s, x)
  ∧ ⊨ s.next ≠ null ∧ ⊨ s.i > s.next.i ∧ ⊨ s.next.next ≐ null
  ∧ ⊨ s.next@pre ≐ null
  ∧ ⊨ s.next.contents()->including(s.i) ≐ Set{s.i@pre}
  ->including(x))
∨ (Δ(s, x)
  ∧ ⊨ s.next ≠ null ∧ ⊨ s.i > s.next.i
  ∧ ⊨ s.next.i > s.next.next.i ∧ ⊨ s.next.next ≠ null
  ∧ ⊨ invNode(s.next.next)
  ∧ ⊨ s.next@pre ≐ null
  ∧ ⊨ s.next.contents()->including(s.i) ≐ Set{s.i@pre}
  ->including(x))

```

While the second sub-case is unsatisfiable since it asserts that the insertion increases the list length by two, the first sub-case reduces to:

$$\begin{aligned}
 & \Delta(s, x) \\
 & \wedge \models s.\text{next} \neq \text{null} \wedge \models s.i > s.\text{next}.i \wedge \models s.\text{next}.\text{next} \doteq \text{null} \\
 & \wedge \models s.\text{next@pre} \doteq \text{null} \\
 & \wedge \models \text{Set}\{s.\text{next}.i\} \rightarrow \text{including}(s.i) \doteq \text{Set}\{s.i@pre\} \\
 & \hspace{15em} \rightarrow \text{including}(x)
 \end{aligned}$$

which, due to set reasoning, corresponds to a test case in which the inserted element x is not already in the list. The test cases still containing an occurrence of the invariance predicate correspond to the class of “yet to be tested” test cases.

4.4 Alias Closure

Unfolding and invariant insertion represent only a first step of the exploration of the specification by case distinction. There is another implicit case distinction that needs to be considered, since the two references s and $s.\text{next}$ could actually refer to the same object, due to a cycle in the object graph. We should then distinguish the cases where $s.\text{next} \doteq s$ and where $s.\text{next} \not\doteq s$. In the same way, we should distinguish the cases $s.\text{next@pre} \doteq s$ and $s.\text{next@pre} \not\doteq s$.

To handle these four cases in the test case generation, we add the following tautology, called *alias distinction*, to the unfolding of our test specification:

$$\begin{aligned}
 & (\models s.\text{next} \doteq s \vee \models s.\text{next} \not\doteq s) \\
 \wedge & (\models s.\text{next@pre} \doteq s \vee \models s.\text{next@pre} \not\doteq s)
 \end{aligned}$$

In the cases $s.\text{next} \doteq s$ and $s.\text{next@pre} \doteq s$, the invariants evaluate to false due to the strict inequality, thus only the cases $s.\text{next} \not\doteq s$ and $s.\text{next@pre} \not\doteq s$ remain. Computing the DNF in our example leads to almost the same formula as in the previous subsection, where $\models s.\text{next} \not\doteq s \wedge \models s.\text{next@pre} \not\doteq s$ is added to each conjoint.

In the general case, the alias closure of a formula is the conjoint of the tautologies $p \doteq q \vee p \not\doteq q$ for all the references p and q occurring in the formula (all other reference pairs are not relevant for case-splitting; so when we decided to unfold the invariants to a certain depth, we also made a decision on the maximum path-sizes and finally the maximum number of nodes in a state). Formally, let $\text{Path}(\varphi)$ be the set of path-expressions (references) occurring in a formula φ . We define $\text{AliasClosure}(\varphi)$ as the set of formulae

$$\{ p \doteq q \vee p \not\doteq q \mid p, q \in \text{Path}(\varphi) \wedge p \text{ non-identical to } q \}$$

This produces all possible objects graphs, instead of only tree-like structures.

4.5 Generating Test Object-Graphs from Test Cases

Finally, ground instantiations of the underlying object model (i. e., in our example, instances of singly-linked lists) need to be generated. For example, a concrete state-pair $\tau = (\sigma, \sigma')$ that can be given for test case

Table 1. Sample set of resulting test cases

List in pre-state	Inserted element	List in post-state
$\boxed{3}$	3	$\boxed{3}$
$\boxed{5}$	9	$\boxed{9} \rightarrow \boxed{5}$
$\boxed{6} \rightarrow \boxed{1}$	1	$\boxed{6} \rightarrow \boxed{1}$
$\boxed{6} \rightarrow \boxed{3}$	5	$\boxed{6} \rightarrow \boxed{5} \rightarrow \boxed{3}$
$\boxed{8} \rightarrow \boxed{5} \rightarrow \boxed{1}$	5	$\boxed{8} \rightarrow \boxed{5} \rightarrow \boxed{1}$
$\boxed{7} \rightarrow \boxed{6} \rightarrow \boxed{5}$	9	$\boxed{9} \rightarrow \boxed{7} \rightarrow \boxed{6} \rightarrow \boxed{5}$

$\Delta(s, x) \wedge \models \mathbf{s.next} \neq \mathbf{null} \wedge \models \mathbf{s.i} > \mathbf{s.next.i}$
 $\wedge \models \mathbf{s.next.next} \doteq \mathbf{null} \wedge \models \mathbf{s.next@pre} \doteq \mathbf{null}$
 $\wedge \models \mathbf{Set}\{\mathbf{s.next.i}\} \rightarrow \mathbf{including}(\mathbf{s.i}) \doteq \mathbf{Set}\{\mathbf{s.i@pre}\}$
 $\rightarrow \mathbf{including}(x)$

and the ground instance $\mathbf{s.insert}(2)$ of the operation call $\mathbf{insert}()$ is:

$\sigma = \{ oid_0 \mapsto (i = 3, next = null) \}$
 $\sigma' = \{ oid_0 \mapsto (i = 3, next = oid_1), oid_1 \mapsto (i = 2, next = null) \}$

which describes the requirement that inserting 2 into the list that only contains the element 3 should result in the sorted singly-linked list that contains the elements 3 and 2. Table 1 shows a sample set of test cases resulting from an unfolding of the test specification up to a list length of 3. For every list length there are two test cases: one for the case that the inserted element is already in the list and one for the case of an actual insertion.

5 Integrating the Technique in HOL-TESTGEN

HOL-TESTGEN [10] is a specification and test case generation environment extending the interactive theorem prover Isabelle/HOL [18]. The HOL-TESTGEN method is two-staged: first, the original formula is partitioned into test cases by transformation into a normal form called test theorem. Second, the test cases are analyzed for ground instances (the test data) satisfying the constraints of the test cases. Particular emphasis is put on the control of explicit test hypotheses. Finally, HOL-TESTGEN supports the generation of test drivers that allow for validating that an implementation fulfills its abstract specification. As such, developing UML/OCL support for HOL-TESTGEN, including its integration into a formal model-driven development toolchain, e.g., [6], enables the validation that an implementation fulfills the test specifications given in UML/OCL.

Extending HOL-TESTGEN with support for OCL creates certain challenges: the unfolding of OCL invariants introduces a new kind of splitting rule that needs to be supported efficiently by the splitter algorithm of HOL-TESTGEN. Moreover, HOL-TESTGEN does not yet support conjoint clauses that have no

reference to the program under test (resulting in a failure during test data form computation). This is motivated by the fact that we cannot generate test drivers for such conjoint clauses. In the future, the generated test drivers could either silently drop such test cases or, following the OCL semantics, test for a deadlock.

6 Related and Future Work

6.1 Related Work

While there are several works that discuss specification-based test case generation based on UML/OCL models, none of them supports the three-valuedness of OCL. The most closely related works [1, 2, 4, 24] are all inspired by the seminal work of Dick and Faivre [13] and, thus, share the idea of using symbolic DNF computation for partitioning the input space. Moreover, there are works using sequence diagrams as an input for test case generation, e. g., [16], or pairwise testing of OCL contracts, e. g., [19]. Finally, Gogolla et al. [14] apply random-testing strategies for analyzing properties of OCL specifications.

For program-based tests, there are two test data generators that apply symbolic techniques: Korat [5] and Java Pathfinder [23]. Korat [5] generates from preconditions and a bound on the number of nodes of data structures, an input partitioning by a combination of symbolic execution and (simple) constraint solving. The idea of integrating a symbolic state deeply inside the execution environment, i. e., inside a Java virtual machine (JVM) as suggested in JPF-SE [3] (a successor of Java Pathfinder [23]), substantially improved the approach and inspired systems such as Pex [21] (a model-based testing tools for the .net).

6.2 Future Work

HOL-TESTGEN's generation strategies are geared towards inductively generated data (such as enumeration types, or lists and sets). In this paper, we have shown how *co*-inductively generated data such as object graphs can be tackled. The described translation is a pre-computation step, but it remains to provide new tactical infra-structure to implement the unfolding strategies and the alias closure. The concrete model-generation for the resulting specification is a standard-game for SMT-based model-construction generators in HOL-TESTGEN.

References

- [1] van Aertryck, L., Jensen, T.: UML-CASTING: Test synthesis from UML models using constraint resolution. In: Jézéquel, J.M. (ed.) AFADL 2003 (2003)
- [2] Aichernig, B.K., Pari Salas, P.A.: Test case generation by OCL mutation and constraint solving. In: QSIC 2005, pp. 64–71. IEEE Computer Society, Los Alamitos (2005)
- [3] Anand, S., Păsăreanu, C.S., Visser, W.: JPF-SE: A Symbolic Execution Extension to Java PathFinder. In: Grumberg, O., Huth, M. (eds.) TACAS 2007. LNCS, vol. 4424, pp. 134–138. Springer, Heidelberg (2007)
- [4] Benattou, M., Bruel, J.M., Hameurlain, N.: Generating test data from OCL specification. In: WITUML (2002)

- [5] Boyapati, C., Khurshid, S., Marinov, D.: Korat: automated testing based on Java predicates. In: ISSTA, pp. 123–133 (2002)
- [6] Brucker, A.D., Doser, J., Wolff, B.: An MDA framework supporting OCL. *Electronic Communications of the EASST 5* (2006)
- [7] Brucker, A.D., Krieger, M.P., Wolff, B.: Extending OCL with null-references. In: Ghosh, S. (ed.) *MODELS 2009*. LNCS, vol. 6002, pp. 261–275. Springer, Heidelberg (2010)
- [8] Brucker, A.D., Wolff, B.: HOL-OCL: A Formal Proof Environment for UML/OCL. In: Fiadeiro, J.L., Inverardi, P. (eds.) *FASE 2008*. LNCS, vol. 4961, pp. 97–100. Springer, Heidelberg (2008)
- [9] Brucker, A.D., Wolff, B.: An extensible encoding of object-oriented data models in HOL. *Journal of Automated Reasoning* 41, 219–249 (2008)
- [10] Brucker, A.D., Wolff, B.: HOL-TESTGEN: an interactive test-case generation framework. In: Chechik, M., Wirsing, M. (eds.) *FASE 2009*. LNCS, vol. 5503, pp. 417–420. Springer, Heidelberg (2009)
- [11] Brucker, A.D., Wolff, B.: Semantics, calculi, and analysis for object-oriented specifications. *Acta Informatica* 46(4), 255–284 (2009)
- [12] Church, A.: A formulation of the simple theory of types. *Journal of Symbolic Logic* 5(2), 56–68 (1940)
- [13] Dick, J., Faivre, A.: Automating the generation and sequencing of test cases from model-based specifications. In: Larsen, P.G., Woodcock, J.C.P. (eds.) *FME 1993*. LNCS, vol. 670, pp. 268–284. Springer, Heidelberg (1993)
- [14] Gogolla, M., Hamann, L., Kuhlmann, M.: Proving and visualizing OCL invariant independence by automatically generated test cases. In: Fraser, G., Gargantini, A. (eds.) *TAP 2010*. LNCS, vol. 6143, pp. 38–54. Springer, Heidelberg (2010)
- [15] Gogolla, M., Richters, M.: Expressing UML class diagrams properties with OCL. In: Clark, A., Warmer, J. (eds.) *Object Modeling with the OCL*. LNCS, vol. 2263, pp. 85–114. Springer, Heidelberg (2002)
- [16] Li, B.L., shu Li, Z., Qing, L., Chen, Y.H.: Test case automate generation from UML sequence diagram and OCL expression. In: *Computational Intelligence and Security*, pp. 1048–1052. IEEE Computer Society, Los Alamitos (2007)
- [17] de Moura, L., Bjørner, N.S.: Z3: An efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) *TACAS 2008*. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
- [18] Nipkow, T., Paulson, L.C., Wenzel, M.: Isabelle/HOL—A Proof Assistant for Higher-Order Logic. LNCS, vol. 2283. Springer, Heidelberg (2002)
- [19] Noikajana, S., Suwannasart, T.: An improved test case generation method for Web service testing from WSDL-S and OCL with pair-wise testing technique, pp. 115–123. IEEE Computer Society, Los Alamitos (2009)
- [20] Object Management Group: UML 2.2 OCL specification (2010), Available as OMG document formal/2010-02-01
- [21] Tillmann, N., de Halleux, J.: Pex—white box test generation for.NET. In: Beckert, B., Hähnle, R. (eds.) *TAP 2008*. LNCS, vol. 4966, pp. 134–153. Springer, Heidelberg (2008)
- [22] Torlak, E., Jackson, D.: Kodkod: A relational model finder. In: Grumberg, O., Huth, M. (eds.) *TACAS 2007*. LNCS, vol. 4424, pp. 632–647. Springer, Heidelberg (2007)
- [23] Visser, W., Havelund, K., Brat, G.P., Park, S., Lerda, F.: Model checking programs. *Autom. Softw. Eng.* 10(2), 203–232 (2003)
- [24] Weissleder, S., Schlingloff, B.H.: Quality of automatically generated test cases based on OCL expressions. In: *ICST*, pp. 517–520. IEEE Computer Society, Los Alamitos (2008)