



ELSEVIER

Available at
www.ComputerScienceWeb.com
POWERED BY SCIENCE @ DIRECT®

The Journal of Systems and Software xxx (2003) xxx-xxx

 **The Journal of
Systems and
Software**

www.elsevier.com/locate/jss

Temporal logic properties of Java objects

Radu Iosif *, Riccardo Sisto

Dipartimento di Automatica Politecnico di Torino, corso Duca degli Abruzzi 24, 10129 Torino, Italy

Received 23 December 2002; accepted 27 December 2002

6 Abstract

7 Applying finite-state verification techniques to software systems looks attractive because they are capable of detecting very subtle
8 defects in the logic design of these systems. Nevertheless, the integration of existing formal verification tools within programming
9 environments is not yet easy, mainly because of the semantic gap between widely used programming languages and the languages
10 used to describe system requirements. In this paper, we propose a formal requirement specification notation based on linear
11 temporal logic, with regard to object oriented program elements, such as classes and interfaces. The specification is inherently object
12 oriented and is meant for the verification of concurrent and distributed software systems.

13 © 2003 Published by Elsevier Science Inc.

14 *Keywords:* Source code verification; Linear temporal logic; Object orientation; Java

15 1. Introduction

16 Thanks to the recent advances in tool support, finite-
17 state verification (FSV) techniques such as model
18 checking (Holzmann, n.d.) can now be applied with
19 interesting results to the verification of concurrent
20 software systems. Nevertheless, much work is still needed
21 to enable the transition of these techniques from
22 research to actual practice. On one hand, verification
23 techniques are generally difficult to use and not yet well
24 integrated in common programming environments
25 programmers are used to. On the other hand, most of
26 these techniques adhere to a monolithic and rather static
27 model of software, which is no longer adequate to the
28 new programming paradigms in use today. It is a matter
29 of fact that object-oriented (OO) languages and mid-
30 dleware like for example Java and CORBA, providing
31 concurrent, distributed and even mobile objects, are
32 becoming one of the most common tools for building
33 applications. Despite this fact, the new features of this
34 kind of software, mainly dynamicity and object-orientation,
35 are not well tackled by the existing verification
36 tools.

In this paper we focus attention on model checking 37
techniques for concurrent and distributed object oriented 38
source programs, and address the problem of 39
specifying temporal logic properties related to this kind 40
of software. Our specific objective is defining a formal 41
but user-friendly specification technique for expressing 42
properties which follows the object oriented approach, is 43
well integrated in the source code the programmer is 44
familiar with, and can easily express what is typically 45
needed. The Java language is taken as a reference for 46
developing the proposed specification technique, even 47
though the method deals with common OO ideas, which 48
makes it suitable for other similar languages, like C++ 49
and CORBA IDL. 50

As we are considering OO software, it is of crucial 51
importance to be able to associate properties with the 52
language elements used by the programmer, i.e. classes 53
and interfaces, and to exploit the mechanisms of object 54
orientation such as inheritance as much as possible and 55
consistently with the common programming practice. 56
OO programs are actually collections of classes, possibly 57
grouped into packages, and, in the OO philosophy, all 58
these are reusable software modules. This implies that it 59
is important to be able to assess and verify not only 60
properties related to an application as a whole, but also 61
properties that each single reusable class or package 62
should satisfy, and these are typically required to hold 63

* Corresponding author.

E-mail addresses: iosif@athena.polito.it (R. Iosif), sisto@polito.it (R. Sisto).

64 somewhat independently of the way the classes or
65 packages will actually be used. If we consider distributed
66 OO software systems, it is even possible that the whole
67 program does not yet exist when verification is to be
68 done, because in such systems server objects are gener-
69 ally made available to unknown clients. Since clients
70 may eventually be developed later on, the verification of
71 servers has to be done without having a complete ap-
72 plication at hand, but only having the source code of
73 some classes.

74 Expressing properties associated with classes, and not
75 with elements of the global program state as in the
76 classical approach, opens new problems. As long as only
77 static, i.e. global, variables are involved in the proper-
78 ties, the meaning of temporal logic formulae is exactly
79 the same as with other non-object-oriented programs,
80 because the lifetime of static variables coincides with the
81 program lifetime. Instead, if for example formulae are
82 associated with classes, the meaning is different, because
83 in the program lifetime each class can be instantiated
84 many times and not necessarily at program startup.
85 Moreover, the inheritance of properties must be con-
86 veniently defined. Of course, these new kinds of speci-
87 fications generally have some intuitive meaning, related
88 with the common understanding of OO concepts, which
89 programmers can easily learn. Nevertheless, a formal
90 definition of their semantics is needed. This paper ad-
91 dresses the above problems and proposes a consistent
92 solution in the Java environment.

93 The paper is organized as follows: Section 2 empha-
94 sizes the distinction between interface and implementa-
95 tion code, along with its implication regarding object
96 oriented property specification, Section 3 introduces the
97 formal execution model used as the basis for the se-
98 mantics of our temporal logic formulae, Section 4 de-
99 scribes the notation used in property specification,
100 Section 5 addresses some problems related to the veri-
101 fication of the specified properties, and Section 7 con-
102 cludes.

103 2. Interface and implementation properties

104 According to the common object oriented under-
105 standing, an interface is an abstract specification of
106 functionalities, without going through implementation
107 details. Its purpose is to transfer information from the
108 class developer to the class user. These two roles entail
109 two different viewpoints. The user of an already existing
110 class tends to see the class instances as black boxes that
111 can be accessed via a particular interface and implement
112 a particular functionality in a way that in principle can
113 be ignored. By contrast, the developer of a class works
114 on the class internals and sees how the class function-
115 alities are implemented. In this paper we use the pure
116 OO concept of interface, i.e. the interface of a class is

represented by public methods that can be invoked, with
their prototypes, whereas the implementation is all the
rest. This means that all the attributes are considered
encapsulated. Such assumption simplifies our work, but
is not a real restriction, because non-encapsulated at-
tributes can always be represented by means of appro-
priate get and set methods. As an example, let us
consider the Java code in Fig. 1. The C class represents a
possible implementation of the AbstractContainer
interface. Under the assumption that the interface has
been written separately from the class, one might need
to specify abstract properties related to the interface
functionality, disregarding the way it can be actually
achieved. In our example, such a requirement may be
that the read() method, always when called, returns a
positive value. The implementor of the C class must
ensure that this formula holds for the specific imple-
mentation, and may impose a sufficient condition, re-
quiring that, in every object state, _mod has a value
greater than zero. It can be noted at first sight that the C
implementation of the interface respects both require-
ments, the first one referring to the value returned by an
abstract method, while the second one involving also a
class field, defining an instance variable.

We can now divide the properties that can be ex-
pressed about a class into two distinct subsets, according
to the point of view under which they are formulated:

1. *interface properties* expressed by the class user point
of view and involving only interface elements (i.e., ab-
stract methods);
2. *implementation properties* expressed by the class de-
veloper point of view and involving at least an imple-

```

interface AbstractContainer {
    void set(int data);
    int get();
    int update();
    int read();
}
class C implements AbstractContainer {
    int _data;
    int _mod;
    void set(int data) { _data = data; }
    int get() { return _data; }
    int update() {
        if (_data > 0) _mod = _data;
        else _mod = - _data;
    }
    int read() { return _mod; }
}

```

Fig. 1. Interface and implementation.

149 mentation element of the class (i.e., encapsulated
150 class attributes and method statements).

151 A class user who only knows the class interface can
152 only specify interface properties, and expect that every
153 class implementation will satisfy them. A class developer
154 can instead specify both kinds of properties and verify
155 them. In particular, the developer can specify additional
156 implementation properties, for example to express some
157 internal consistency requirements, and can verify that
158 the resulting implementation satisfies all interface and
159 implementation properties before delivering it.

160 3. Behavioral semantics

161 As we specify requirements using linear temporal
162 logic (LTL) formulae, we need to define their formal
163 semantics with respect to a sequential model of com-
164 putation. Formally, such a model can be represented as
165 a labeled transition system $LTS = \langle \Sigma, S, \rho, s_0 \rangle$ where:

- 166 • Σ is the *alphabet* (a finite set of *symbols* representing
167 computation events),
- 168 • S is a set of *states*,
- 169 • $\rho : S \times \Sigma \rightarrow S$ is the *transition mapping* giving, for each
170 state-symbol pair, the next state reached after the oc-
171 currence of the corresponding event.
- 172 • $s_0 \in S$ is the *initial state*.

173 Given the alphabet Σ , an *infinite word* is an infinite
174 sequence of symbols of Σ . An execution of the LTS on
175 an infinite word $w = a_0a_1 \dots$ is an infinite sequence of
176 states $\pi = s_0s_1 \dots$ with the following properties:

- 177 • s_0 is the initial state of the LTS,
- 178 • $s_i = \rho(s_{i-1}, a_{i-1})$ for every $i \geq 1$ that is, every state of
179 the sequence is obtained from the previous one in
180 agreement with the transition mapping.

181 LTL (Manna and Pnueli, 1992) is a language for
182 reasoning about sequences of states a program goes
183 through during its execution. This language is that of
184 propositional calculus augmented with the following
185 four symbols representing temporal operators. The in-
186 terested reader is referred to (Manna and Pnueli, 1992)
187 for a formal definition of these operators' semantics.

- 188 (1) \circ which is read 'at the next time';
- 189 (2) \square which is read 'always in the future';
- 190 (3) \diamond which is read 'eventually in the future';
- 191 (4) \mathcal{U} which is read 'until'.
- 192 (5) \mathcal{W} which is read 'weak until'.

193 The notation $\pi \models_i A$ is read ' A is true for the execu-
194 tion sequence π starting with its i th state'. One says that

a temporal formula A is true for a sequence π , and one
writes $\pi \models A$, if $\pi \models_0 A$ (A is true in the initial state of π).

Thanks to the modularity of the object model, when
reasoning about the execution of an object-oriented
program such as a Java program it is possible to con-
sider the execution of each object separately. The exe-
cution of an object can be modeled at two different
abstraction levels (i.e., interface and implementation),
according to which point of view (class user or devel-
oper) is considered. The intuition is that interface
properties, being more abstract, could be interpreted
only considering sequences of method invocation and
method return events. On the other hand, implementa-
tion properties involve program variables, whose actual
values need to be explicitly represented in the model.
Formally, we separate the LTS model needed to describe
interface behaviors from the one concerning strict im-
plementation details that is, actual object states. This
separation is very important in order to be able to define
the semantics of interface properties independently of
how interfaces are implemented. In this way, formal
reasoning about interface properties is possible even if
implementations are not known, according to the object
oriented paradigm.

Let us consider first the interface-level execution
model. Since the internals of the object are not known to
the class users, state information clients can be aware of
is represented at most by the sequence of method call
and return events that have occurred since the object
creation. In other words, a user cannot distinguish two
objects in which the same sequence of interface events
has taken place, but the two objects could well be in two
different states from the developer point of view. Based
on this consideration, we define the interface-level state
of an object as the ordered sequence of interface-level
events that have occurred in its past. Formally, the in-
terface-level execution model is a labeled transition
system $LTS^n = \langle \Sigma^n, S^n, \rho^n, s_0^n \rangle$ where:

- Σ^n is the set of all possible method call and return
events (including the constructor call and returns).
- S^n is the set of interface-level states, which includes all
the finite sequences of events $s = \langle e_1, e_2, \dots, e_k \rangle$, such
that $e_{1,2,\dots,k} \in \Sigma^n$. The empty sequence is denoted by ϵ .
- $\rho^n(s, e) = s.e$ where $s.e$ denotes the concatenation of
symbol e to sequence s .
- $s_0^n = \epsilon$.

Let us now consider the implementation-level object
execution model. It can be defined as $LTS^m =$
 $\langle \Sigma^m, S^m, \rho^m, s_0^m \rangle$ where:

- Σ^m is the set of implementation-level events. In gen-
eral, they represent computation actions correspond-
ing to the execution of statements, and include also
method call and return events.

- 248 • S^m is the set of implementation-specific object states;
- 249 an object state includes the state of each instance
- 250 variable, and any other state information related to
- 251 the object.
- 252 • $\rho^m(s, e)$, where $s \in S^m$ and $e \in \Sigma^m$ is the new state
- 253 reached after the occurrence of event e .
- 254 • s_0^m is the initial object state, representing the state im-
- 255 mediately following the object creation event.

256 How this kind of model can actually be extracted
257 from the Java source code is outside the scope of this
258 paper. A possible solution is presented for example in
259 (Iosif and Sisto, 2000).

260 For technical reasons, we introduce a modified ver-
261 sion of this LTS, which can be conveniently used as a
262 formal basis for reasoning from the class developer
263 point of view, and for verifying properties. This modi-
264 fied LTS is obtained joining state information of the two
265 previously defined models. In practice, to stress the fact
266 that the implementation-level model is a refinement of
267 the interface-level one, the object state is defined as a
268 pair of state components $s = \langle s_n, s_m \rangle$, where $s_n \in S^n$ co-
269 incides with the interface-level state, whereas $s_m \in S^m$
270 encompasses all the additional state information such as
271 the current state of object attributes and the current
272 state of the method calls that are in progress. Formally,
273 the joint LTS is defined as $LTS^i = \langle \Sigma^i, S^i, \rho^i, s_0^i \rangle$ where:

- 274 • $\Sigma^i = \Sigma^m$ is the set of implementation-level events,
- 275 which is a superset of interface-level events i.e.,
- 276 $\Sigma^n \subseteq \Sigma^m$.
- 277 • $S^i \subseteq S^n \times S^m$ is the set of implementation-level states.
- 278 • $\rho^i(\langle s_n, s_m \rangle, e) = \begin{cases} \langle \rho^n(s_n, e), \rho^m(s_m, e) \rangle & \text{if } e \in \Sigma^n, \\ \langle s_n, \rho^m(s_m, e) \rangle & \text{otherwise.} \end{cases}$
- 279 • $s_0^i = \langle \epsilon, s_0^m \rangle$.

280 In LTS^i , a transition is fired either by a method call/
281 return event, in which case both state components
282 change, or by other implementation-level events, with a
283 change in the implementation-level state component
284 only. Fig. 2 shows a graphical representation of an in-
285 terface-level execution sequence and a corresponding
286 implementation-level sequence, related to the sample
287 Java code in Fig. 1, where call(set,data) represents the
288 event corresponding to the issue of a call of method set
289 with argument data, whereas ret(set) is the event corre-
290 sponding to a return from method set. It can be noticed
291 how interface-level states are mapped onto sequences of
292 implementation-level states. This correspondence is
293 formally defined in Section 5.

294 4. Property specification

295 In this paper we consider only verification of prop-
296 erties associated with the program source code (source

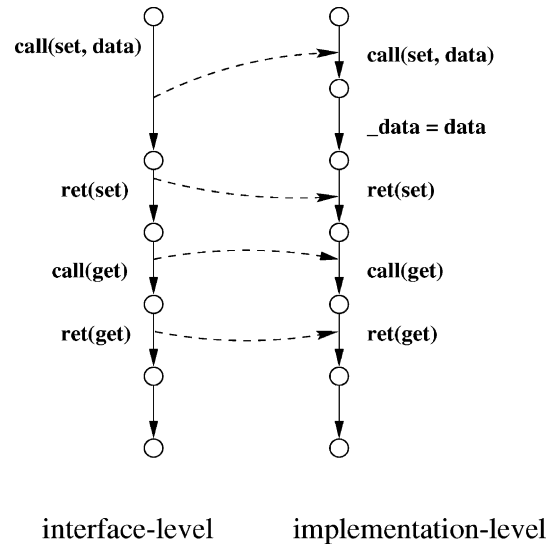


Fig. 2. Sample execution path.

code verification) and not verification of already com- 297
298 piled code (binary code verification). Source code veri-
299 fication can be integrated into the software development
300 process more easily and facilitates the programmer in
301 specifying correctness requirements, since such require-
302 ments can be associated directly with the program ele-
303 ments manipulated by the programmer, such as
304 packages, classes and methods.

305 Before presenting the notation, some general princi-
306 ples that have been followed to define it are illustrated.
307 As interface and implementation properties play differ-
308 ent roles, they are treated differently. First of all, im-
309 plementation properties can only be associated with
310 class implementation definitions (i.e. Java class envi-
311 ronments), and not with interface definitions (such as
312 Java interface environments). Moreover, given that each
313 instance of a class that implements an interface can be
314 seen as an instance of the interface itself, and that each
315 instance of a class is also an instance of all its super-
316 classes, it is required that both interface and imple-
317 mentation properties hold in all the classes derived by
318 inheritance or implementation. An equivalent form of
319 the requirement is that interface and implementation
320 properties are inherited by the derived classes. However,
321 following the general assumption that classes derived by
322 inheritance can override implementation details but not
323 interface characteristics, we admit overriding of imple-
324 mentation properties only. In this way, it is guaranteed
325 that inheritance preserves the class interface, along with
326 all the associated properties, and clients can rely on the
327 fact that interface properties are satisfied by all the de-
328 rived classes. At the same time, the developer of a de-
329 rived class is free to override not only part of the class
330 implementation, but also some of the implementation
331 properties, the only firm requirement being the preser-
332 vation of interface properties.

333 4.1. Syntax and semantics

334 In order to define the syntax of our property speci-
335 fication language, we first define a set of atomic prop-
336 ositions that are composed into property expressions
337 (formulae). The set of formulae is the core of the spec-
338 ification language. Users can rely directly on it for
339 writing new properties or apply already written patterns
340 (Dwyer et al., 1999) that come as a standard library.
341 Atomic propositions in property formulae can be:

- 342 (1) Java boolean expressions; they are evaluated atomi-
343 cally, yielding information which regards the object
344 state.
345 (2) Special atomic propositions, usually yielding infor-
346 mation concerning the flow of control.

347 Let us now consider special atomic propositions. The
348 specification of interface properties makes use of the
349 following two special atomic propositions:

- 350 • `calling(m [, argument_list])`, which is true
351 in all object states where some call to method `m` with
352 actual arguments `argument_list` is being executed
353 (i.e. is pending). Square brackets indicate optionality:
354 if method `m` has no arguments, the argument list is
355 void. In practice, this atomic proposition becomes
356 true whenever a call to method `m` with actual argu-
357 ments `argument_list` is issued, and it remains true
358 until the corresponding method execution terminates.
359 Of course, in a concurrent environment it is possible
360 to have time-overlapping executions of `m`, in which
361 case the predicate remains true until the current num-
362 ber of concurrent executions of `m` with actual argu-
363 ments `argument_list` becomes 0.
364 • `returns(m [, argument_list] [, x])` is true
365 in a certain object state provided that the last inter-
366 face-level event occurred in the object is a return of
367 value `x` from a call of method `m` with actual argu-
368 ments `argument_list`. As with the previous pred-
369 icate, the argument list and the return value can be
370 missing.

371 Although omitted here for brevity, the above prop-
372 ositions can be formally defined by giving their truth
373 value as a function of the interface-level state defined in
374 Section 3.

375 This is the minimum core feature to express interface
376 properties. However, to facilitate the task of specifying
377 properties, it is useful to extend the language with some
378 more propositions. For example, in some circumstances
379 it is useful to refer to the number of pending calls to a
380 method `m` with actual arguments `argument_list` in a
381 certain program state, and this is denoted as `#calls(m`
382 `[, argument_list])`.

Property specifications take the syntax given in Fig. 3. 383
For brevity, we present here only the top-most grammar 384
rules, the rest of them being described informally. The 385
upper-case symbols denote terminals, while the lower- 386
case ones are non-terminal symbols. The symbols en- 387
closed in square braces are optional. Properties can be 388
parameterized with respect to a number of free variables 389
introduced by the `formal_parameter_list` symbol, 390
that takes the same syntactical form of a Java method 391
parameter list. A parameterized property is denoted also 392
as an *open* property. All other properties are denoted 393
also as *closed* properties. Open properties are not meant 394
for actual verification, rather they are introduced as 395
patterns for further specialization or simply for being re- 396
used. Indeed, in many instances, specializing already 397
written properties proves to be a useful feature. The 398
substitution of formal parameters with actual arguments 399
in open properties is literal. Closed properties can be 400
quantified over Java types. The `expression` symbol 401
denotes a Java boolean expression used to restrict the 402
quantification domain. An `ltl_formula` is usually 403
obtained from any number of basic propositions con- 404
nected with the standard LTL operators. In Fig. 3, the 405

```

closed_property:
  IDENTIFIER '=' [quantifier_expression]
                ltl_formula

open_property:
  IDENTIFIER '(' [formal_parameter_list] ')'
                '=' ltl_formula

quantifier_expression:
  quantifier formal_parameter_list
                ['(' expression ')']

quantifier: one of
  'forall' 'exists'

ltl_formula:
  atomic_proposition
  | property_reference
  | '(' ltl_formula ')'
  | ltl_formula binary_operator ltl_formula
  | unary_operator ltl_formula

binary_operator: one of
  'U' 'W' '->' '<->' '||' '&&'

unary_operator: one of
  '[' '<>' 'X' '!'

```

Fig. 3. Properties syntactic grammar.

406 binary_operators U and W denote the temporal \mathcal{U}
407 and \mathcal{W} operators respectively, while the unary_op-
408 erators [], <> and X denote the temporal operators
409 \square , \diamond and \circ respectively.

410 Properties can be referenced by means of their name.
411 The property reference is literally substituted with the
412 referenced property LTL formula. Of course, it is an
413 error to specify mutually recursive properties. More-
414 over, there is another well-formedness condition that
415 involves referencing quantified properties. The meaning
416 of a quantifier which occurs other than at the beginning
417 of a formula is undefined. The syntax of property
418 specification ensures that this requirement is respected,
419 imposing that the LTL formulas describing properties
420 can only refer to open properties, which as said, cannot
421 be quantified. In this way, a property reference can be
422 used also as an actual argument to specialize an open
423 property. Let us consider the following open property
424 which expresses an overall truth:

425 `Always(boolean P) = [](P)`

426 It can be specialized with respect to any boolean ex-
427 pression, including a reference to a property:

428 `myTruth() = (var == 0) -> <>(var > 0)`
429 `myAlways() = Always(myTruth())`

430 First, the reference to myTruth() literally substitutes
431 each occurrence of the formal parameter P. Then each
432 reference is literally substituted by its LTL formula. In
433 the end, the meaning is exactly the same as if myAl-
434 ways() were defined `[]((var == 0) -> <>(var >`
435 `0))`.

436 Open properties parameterized by formulae intro-
437 duce patterns, used to cover a broad range of require-
438 ments for real systems, in terms of parameters that must
439 be filled with descriptions of specific system states or
440 events. These descriptions can be more complex than
441 just a boolean proposition or event e.g., they can be also
442 properties given in terms of LTL formulae. Users can be
443 provided with a specification pattern (Dwyer et al.,
444 1999) library written as a Java interface declaring a
445 collection of open properties. Fig. 4 shows part of such a
446 library. Informally, the GlobalAbsence, Before-
447 Absence, AfterAbsence and BetweenAbsence
448 open properties express the absence of an event P
449 overall, before event R, after event Q and between Q and
450 R, respectively. In a similar way, it is even possible for
451 users to define their own specification patterns. An ex-
452 ample of library use in coding actual properties is given
453 in the next section.

```
public interface SPL {
/*@
GlobalAbsence(P) = [](!P))

BeforeAbsence(P, R) =
<>(R) -> (!P) U (R))

AfterAbsence(P, Q) = []((Q) -> [](!P))

BetweenAbsence(P, Q, R) =
[](((Q) && !(R) && <>(R))
-> (!P) U (R)))
*/
...
}
```

Fig. 4. Specification patterns library.

4.2. Property specification example

The interface-level atomic propositions enable us to
define several interface properties of interest. As an ex-
ample, let us consider the interface of a concurrently
accessible integer element stack object:

```
interface IntegerStack {
void push(int x);
int pop();
/*@
lifo = forall int x, y, z (x != y)
[]((returns(push, x) U (!returns(push, y)
U returns(pop, z))) -> x == z)
*/
}
```

The interface property lifo informally says that if a
push(x) is followed by a pop() with no intermediate
other push(y), then the return value of pop() is x. It
is a way to specify the LIFO (last in first out) behavior
of a stack. A semantically equivalent way to specify this
property is using a library pattern property from the
collection shown in Fig. 4:

```
lifo = forall int x, y, z (x != y)
[] (BetweenAbsence(returns(push, y),
returns(push, x),
returns(pop, z))
-> x == z)
```

Let us now consider a possible implementation of the
Stack interface:

```
class VectorStack implements Stack {
Vector data = new Vector();
```

```

484 int top;
485 public synchronized void push(int info) {
486     data.add(top++, new Integer(info));
487     notifyAll();
488 }
489 // @popPre = [] (calling(pop) -> top >= 0)
490 public synchronized int pop() {
491     while (top == 0)
492         try {
493             wait();
494         } catch (InterruptedException e) {}
495     Object info = data.remove(--top);
496     return ((Integer) info).intValue();
497 }
498 }

```

499 The `lifo` property is automatically inherited by the
500 `VectorStack` class. It can be easily seen that the
501 `VectorStack` implementation respects the `lifo`
502 property because both `push` and `pop` are synchronized,
503 preventing multiple threads to access the stack internal
504 data. As it can be noticed, the code of the `Vector-`
505 `Stack` class is also annotated with an implementation
506 property, named `popPre`. This property expresses a
507 pre-condition of the `pop` method, ensuring that, when-
508 ever `pop` is called, the instance variable `top` has a
509 positive value. In our case, the implementation of the
510 class meets the requirement, because of the `wait-`
511 `notifyAll` protocol used in the synchronized `pop` and
512 `push` methods, respectively.

513 5. Property verification

514 By *verifying* a property in the context of a given
515 program, we mean deciding if it holds in every execution
516 sequence of the program. For class properties, the de-
517 cision resumes to proving that the property holds in
518 every execution path of each possible instance of the
519 class.

520 As said, execution paths are formally described by an
521 LTS. Generally, the decision of LTL formulae in the
522 frame of an LTS is possible algorithmically (Manna and
523 Pnueli, 1992), given that the set of states is finite. This is
524 typically achieved by means of program slicing and
525 abstraction-based specializations (Corbett et al.,
526 2000a,b). Decision of temporal logic formulae in the
527 frame of a finite LTS was made cost effective by the
528 development of model checking techniques (Holzmann,
529 n.d.), i.e. algorithms that attempt to exhaustively ex-
530 plore the state space generated by a specification in or-
531 der to find counterexamples of the requirements. LTL
532 model checking tools generally do not deal directly with
533 quantified temporal logic formulae, because quantifica-
534 tion increases the complexity of verification tasks and, if
535 quantification domains are infinite, verification becomes

undecidable. Nevertheless, we decided to introduce
quantification in our notation, because it makes the
specification of many properties of interest more direct.
Of course, to make model checking of quantified for-
mulas possible and viable, it is necessary to have suffi-
ciently small quantification domains, which can be
achieved by using abstract representations for quantifi-
cation variables.

The verification of object properties ideally follows a
top-down model. The intuition is that an interface
property must be verified for every class that implements
the interface. Moreover, a property associated with a
class must be verified for every possible instance of the
class. In practice, the mechanism whereby interface
properties are inherited by class implementations en-
sures that any interface property is automatically asso-
ciated also with all the classes that implement the
interface. Thus, the verification task regards only
properties directly or indirectly associated with classes,
that must hold for all their instances. In other words,
implementation properties can be seen as implicitly
quantified over the domain of all existing class instances.
In what follows, we explain the condition under which
interface properties can be verified within implementa-
tion frames.

As discussed in Section 3, the meaning of an interface
property is defined with respect to an interface-level
model, denoted as LTS^n , but it should be verified con-
sidering the implementation-level object behavior de-
scribed by a more detailed LTS, denoted as LTS^i . In
what follows, we denote by $\mathcal{L}(LTS)$, the language of an
LTS that is, the set of all paths it can generate. Since
 LTS^i was defined as a refinement of LTS^n , it is always
possible to extract, from an implementation-level path
 π^i the corresponding interface-level path. Let
 $h : \mathcal{L}(LTS^i) \rightarrow \mathcal{L}(LTS^n)$, be a function defined as fol-
lows:

$$\begin{aligned}
 h(\langle s_k^n, s_k^m \rangle \langle s_{k+1}^n, s_{k+1}^m \rangle \dots) \\
 = \begin{cases} h(\langle s_{k+1}^n, s_{k+1}^m \rangle \dots) & \text{if } s_k^n = s_{k+1}^n, \quad \forall k \geq 0 \\ s_k^n h(\langle s_{k+1}^n, s_{k+1}^m \rangle \dots) & \text{otherwise} \end{cases}
 \end{aligned}$$

Informally, the h function extracts, from an implemen-
tation-level path generated by a program, the corre-
sponding interface-level path, on which we can interpret
an interface property. It can be proven that the expres-
sion above defines indeed a functional relation on
 $\mathcal{L}(LTS^i) \times \mathcal{L}(LTS^n)$, but we will omit the proof for
brevity reasons. Taking into consideration this relation,
it is now necessary to show under what conditions the
outlined verification procedure for interface properties is
sound.

Specifically, soundness requires that if the property
holds on all the implementation-level execution paths
 $\mathcal{L}(LTS^i)$, then it holds also on the corresponding in-
terface-level paths, which are the image of function h ,

588 here denoted as $h(\mathcal{L}(\text{LTS}^i))$. First of all, let us remind
589 that any interface property can be directly interpreted in
590 the frame of LTS^i , because the state of this LTS includes
591 the interface-level state. Of course, this interpretation is
592 based on the fact that any atomic proposition p defined
593 on the interface-level state can also be defined on the
594 implementation-level state in the obvious way:

$$p(\langle s^n, s^m \rangle) = p(s^n) \quad (1)$$

596 If p is a predicate, and s' is the successor of state s in
597 some execution sequence, state s' is said to be a *p-stuttering*
598 of state s if p has the same truth value in both
599 states. An LTL formula f is said to be *closed under*
600 *stuttering* when, for every predicate p that occurs in f , its
601 truth value remains the same under state sequences that
602 differ only by *p-stuttered* states. In the following, we
603 denote the k th element of a sequence σ by σ_k . We can
604 now express the soundness claim:

605 **Theorem 1.** *Let ϕ be an interface property. Then*

$$\sigma \models \phi \iff h(\sigma) \models \phi, \quad \forall \sigma \in \mathcal{L}(\text{LTS}^i) \quad (2)$$

607 *holds if ϕ is closed under stuttering.*

608 **Proof.** Let $\sigma = s_0, s_1, \dots \in \mathcal{L}(\text{LTS}^i)$. Then for each
609 $k \geq 0$, $s_k = \langle s_k^n, s_k^m \rangle$, and for each atomic proposition p
610 that occurs in ϕ we have $p(s_k) = p(s_k^n)$, from (1). For
611 each $k \geq 0$, we have $s_{k+1} \in \rho^i(s_k, \tau_k)$. If, for some $k \geq 0$,
612 $\tau_k \notin \Sigma$, then we have $s_{k+1} = \langle s_k^n, s_{k+1}^m \rangle$ from the definition
613 of LTS^i , which implies $p(s_{k+1}) = p(s_k)$. Consequently,
614 $h(\sigma)$ differs from σ by at most *p-stuttered* states. As ϕ
615 was supposed to be closed under stuttering, its truth
616 value over σ remains unchanged over $h(\sigma)$. ■

617 This result gives us the decision criterion for interface
618 properties. Since all next-free LTL formulae are closed
619 under stuttering, this limitation preserves a good ex-
620 pressive power, enough to describe many meaningful
621 properties.

622 6. Related work

623 The problem of using an object-oriented approach to
624 the formal specification of temporal logic properties to
625 be model checked on object-oriented source code has
626 not been considered so much up to now. Indeed, the
627 main research projects about source-level model check-
628 ing of object-oriented software (Corbett et al., 2000a,b;
629 Havelund and Skakkebaek, 1999; Young, 1994; De-
630 martini et al., 1999) have focused attention on other
631 problems, such as abstraction and slicing techniques,
632 and have always used classical non-object-oriented
633 techniques to express properties. Instead, object-ori-
634 ented temporal logic techniques have been proposed for
635 behavioral specification of object-oriented concurrent

systems (for example in Denker et al., 1997), which is quite different from source-level property specification.

The first FSV tools for Java that have appeared so far follow the typical approach of considering only properties related to the global program scope. In the current version of the JPF tool (Havelund and Skakkebaek, 1999), properties can be specified in the source code, but with reference to static variables only, whereas the JCAT tool (Demartini et al., 1999) does not provide property specification at all, because it deals with deadlock detection only. Recently Corbett et al. (2000a,b) have proposed BSL (Bandera specification language) to specify the properties that can be verified by their tool (Corbett et al., 2000a,b). This language was designed to cover a broad range of notations including assertions, pre- and post-conditions for methods, and temporal logic specifications, and makes it possible to associate properties with classes and methods. However, the semantics of such notations is given only informally and an underlying formal model to enable mathematical reasoning is not defined. As their notation is intended for specifying properties of complete applications, rather than independent classes, the problem of specifying behavioral properties of interfaces is not addressed, while expressing class properties can be done by explicit quantification over the domain of all existing class instances.

An object-oriented property specification technique in part related to our one has been proposed in the context of C++ (Cline and Lea, 1990) to annotate classes and methods with expected properties. In this case, however, the annotated properties are not in the temporal logic form, and they are not intended for verification by FSV techniques. They are rather assertions to be checked at run time.

7. Conclusions

A formal specification technique has been introduced to specify properties related to object-oriented source code, and particularly concurrent and distributed code, taking as a reference the Java language. Specifications generated according to the presented approach can be used to drive source code verification tools such as the ones already delivered for Ada and Java, but also other kinds of software validation tools.

Specifications use an intuitive and simple notation, well integrated in the source code, which makes it possible to associate properties with specific program modules (classes, interfaces, packages) and not only with whole programs, thus enabling an easy object-oriented specification of properties related to open or component-based systems. In particular, interface properties make it possible to express the expected behavior of interfaces independently of how they will actually be

636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670

671

672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688

689 implemented and such properties can soundly be verified
690 on the corresponding implementation level models,
691 provided that next-free formulae are used.

692 The specification of class properties is inherently
693 object oriented. They annotate classes as a whole, thus
694 avoiding the live-code dead-data notion common in
695 program verification strategies, but antithetical to the
696 object-oriented programming paradigm. Moreover, the
697 use of specification patterns can easily be incorporated
698 in the specification task by means of inheritance and
699 parameterized (open) properties.

700 In this paper we have presented a 'core language'
701 containing only the essential features. This notation can
702 easily be extended with more elements (e.g. assertions
703 related to implementations, or more kinds of atomic
704 propositions). The practical goal of having a specifica-
705 tion language with a formal semantics is to make proofs
706 automatically possible. It is the authors' intention to
707 incorporate the property specification notation pre-
708 sented here in a future version of the JCAT tool (De-
709 martini et al., 1999).

710 References

- 711 Cline, M.P., Lea, D., 1990. The Behavior of C++ classes. In:
712 Proceedings of Symposium on Object Oriented Programming
713 Emphasizing Practical Applications, Marist College.

- Corbett, J., Dwyer, M., et al., 2000a. Bandera: extracting finite-state
models from Java source code. In: Proceedings of 22nd Interna-
tional Conference on Software Engineering. 714
715
Corbett, J., Dwyer, M., et al., 2000b. A language framework for
expressing checkable properties of dynamic software. In: Pro-
ceedings of 7th SPIN International Workshop on Model Checking
of Software. 716
717
Demartini, C., Iosif, R., Sisto, R., 1999. A deadlock detection tool for
concurrent Java programs. *Software: Practice & Experience* 29
(7), 577-603. 718
719
Denker, G., Ramos, J., et al., 1997. A Linear Temporal Logic
Approach to Objects with Transactions Algebraic Methodology
and Software Technology. In: *Lecture Notes in Computer Science*,
vol. 1349. Springer-Verlag., pp. 170-184. 720
721
Dwyer, M., Avrunin, G., Corbett, J., 1999. Patterns in property
specifications for finite-state verification. In: Proceedings of 21st
International Conference on Software Engineering. ACM Press,
pp. 411-421. 722
723
Havelund, K., Skakkebaek, J., 1999. Applying Model Checking in
JAVA Verification. In: *Lecture Notes in Computer Science*, 1680,
pp. 216-231. 724
725
Holzmann, G., n.d. The model checker SPIN. *IEEE Transactions on*
Software Engineering, SE-23 (5), 279-295. 726
727
Iosif, R., Sisto, R., 2000. A formal execution model for Java programs.
Technical Report, DAI-ARC, Politecnico di Torino. Available
from <<http://www.dai-arc.polito.it/dai-arc/manual/papers/jtl.ps>>. 728
729
Manna, Z., Pnueli, A., 1992. *The Temporal Logic of Reactive and*
Concurrent Systems. Springer-Verlag. 730
731
Young, M., 1994. A concurrency analysis tool suite for Ada. SERC
Technical Report TR-1288-P. 732
733
734
735
736
737
738
739
740
741
742
743