

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
23 September 2010 (23.09.2010)

PCT

(10) International Publication Number
WO 2010/107671 A2

(51) International Patent Classification:
B24B 7/18 (2006.01) **B24B 41/047** (2006.01)

(21) International Application Number:
PCT/US2010/027139

(22) International Filing Date:
12 March 2010 (12.03.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/160,464 16 March 2009 (16.03.2009) US
12/722,404 11 March 2010 (11.03.2010) US

(71) Applicant (for all designated States except US): **NTT DOCOMO, INC.** [JP/JP]; Sanno Park Tower, 11-1, Nagata-cho, 2-chome, Chiyoda-ku (JP).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **DHURJATI, Dinakar** [IN/US]; 1220 North Fair Oaks Avenue, Apt. 3208, Sunnyvale, CA 94089 (US). **CHANDER, Ajay** [IN/US]; 3333 17th Street, Apt. AA, San Francisco, CA 94110 (US). **INAMURA, Hiroshi** [JP/US]; 1-35-2 Nouk-endai Kanazawaku, Apt. A507, Yokohama, Kanagawa, 236-0057 (JP).

(74) Agents: **VINCENT, Lester, J.** et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 1279 Oakmead Parkway, Sunnyvale, CA 94085-4040 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: DIRECTED TESTING FOR PROPERTY VIOLATIONS

(57) Abstract: A method and apparatus is disclosed herein for automated testing of an application. A processing system executes the application using test input values that are generated for testing the application with respect to a property of interest. During execution, constraints among symbolic variables of the application are collected. Property state is collected and is used to determine whether a branch in the application is relevant to the property. Based on the collected constraints and branch relevancy, new test input values are generated. The process is repeated until all paths in the application corresponding to relevant branches have been explored.



WO 2010/107671 A2

DIRECTED TESTING FOR PROPERTY VIOLATIONS

RELATED APPLICATION

[0001] The present patent application claims priority to and incorporates by reference the corresponding currently pending provisional patent application serial no. 61/160,464, titled, "Directed Testing for Property Violations", filed on March 16, 2009.

FIELD OF THE INVENTION

[0002] The present invention relates to the field of testing computer programs; more particularly, the present invention relates to automatically generating test cases that can detect property violations in computer programs.

BACKGROUND OF THE INVENTION

[0003] Much of the resources for software development in industry are spent on testing and validating software. Industry adoption of automated validation techniques and/or automated test generation techniques has been limited to only some special software domains like automotive/control software. The main technical reason for this limited use of automated techniques is that of scalability. For instance, for moderately large code bases, some automated test generation techniques tend to generate millions of test cases, making them impractical. The predominant way that industry tests software currently is by manually generating test cases and then carrying out the tests. Manual testing, on the other hand, being resource intensive and incomplete, does not scale for increasingly complex modern software.

[0004] One approach to software testing is automatic test generation. There are several bodies of research on using automated testing techniques both to improve coverage and reduce cost. The automatic test generation can be classified into three different types: random testing, symbolic test generation and directed testing. Each of these types is discussed below.

[0005] A more successful industry use of automatic test generation techniques might be the use of random testing. However, the limitations of random

testing are well known. For example, some branches of the program (e.g. if (x == 42) ... where x is an input) have a very little chance of being executed using traditional randomly generated inputs.

[0006] One of the first automatic test generation techniques proposed is to symbolically execute the program and solve the symbolic constraints on each path for each potential error inducing statement. Since then, there have been several improvements proposed including a recent state of the art technique called KLEE. KLEE improves traditional symbolic test generation techniques in many ways including modeling the environment to figure out the errors due to environment, using a powerful constraint solver, and parallel processing to efficiently solve many constraints. KLEE has shown that it is indeed possible to scale symbolic execution to thousands of lines of code. In the case of property testing, however, KLEE is still a brute force solution trying to look for problems in all the areas of the code.

[0007] Directed testing techniques use random testing to avoid intractability and use symbolic analysis to improve coverage. An example of an existing directed testing technique is Concolic Unit Testing Engine (CUTE). “Concolic” is a portmanteau of concrete and symbolic. Concolic testing, such as CUTE, is a hybrid software verification technique that interleaves concrete execution (testing on particular inputs) with symbolic execution, a classical technique that treats program variables as symbolic variables. Symbolic execution is used in conjunction with an automated theorem prover to generate new concrete inputs (test cases) with the aim of maximizing code coverage. Its main focus is finding bugs in real-world software, rather than demonstrating program correctness. CUTE starts off by randomly generating inputs as in the case with pure random testing. While executing the program on the random input, CUTE also collects symbolic constraints on the program statements and branches along the current execution path. After finishing the execution, CUTE generates new inputs for subsequent execution by analyzing the constraints collected in the previous executions. New inputs are generated by iteratively negating the symbolic constraints collected on the branches of the previous iteration and then solving them. If a solution exists, then the new inputs are certain to drive the program on the alternative paths, thus increasing coverage. In cases where it is hard to solve the constraints (e.g. due to limitations of the

constraint solver) or no symbolic constraints exist because of external functions, CUTE falls back on using the concrete values of the previous iteration. While directed techniques like CUTE solve the practical problems of external calls (at the expense of some fault detection capability), they still suffer from the problem of path explosion. For example, a branch inside a loop which is executed N times for one input, could potentially generate 2^N new inputs.

[0008] Another approach, referred to as “property checking,” is based on partial program verification and/or static analysis techniques. Recent research that adopts this approach has focused on partial verification or validation of software. These techniques check whether the software obeys certain properties of interest. The properties are typically expressed as type state properties. Some of these techniques have shown to scale to hundreds of thousands of lines of code. However, the biggest limitation of these techniques is the large number of false positives they report. This is mainly due to limitations on the analysis capability of static techniques in the presence of heap data structures and calls to external libraries.

[0009] Property-based testing describes a method for property testing that first slices the program with respect to the property of interest. It then performs unit testing with respect to the reduced program. The main problem with the conventional property-based approach is the reliance on static slicing algorithms. It is known that static slicing algorithms do not work well in the presence of aliasing and heap data structures.

SUMMARY OF THE INVENTION

[0010] A method and apparatus are disclosed herein for automated testing of an application. A processing system executes the application using test input values that are generated for testing the application with respect to a property of interest. During execution, constraints among symbolic variables of the application are collected. Property state is collected and is used to determine whether a branch in the application is relevant to the property. Based on the collected constraints and branch relevancy, new test input values are generated. The test is repeated until all of the branches in the application are covered.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The present invention will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the invention, which, however, should not be taken to limit the invention to the specific embodiments, but are for explanation and understanding only.

[0012] **Figure 1** illustrates an example of a locking property.

[0013] **Figure 2** illustrates a processing system in which embodiments of the present invention may operation.

[0014] **Figure 3** illustrates an example of a code segment on which embodiments of the present invention may apply.

[0015] **Figure 4** illustrates an example of a loop.

[0016] **Figure 5** illustrates an example showing a limitation on fault detection.

[0017] **Figure 6** is a flow diagram illustrating an embodiment of a property testing technique.

[0018] **Figure 7** illustrates an example of an abstract syntax for a core language.

[0019] **Figure 8** illustrates an example of concolic execution semantics.

[0020] **Figure 9** illustrates an example of a concolic testing algorithm.

[0021] **Figure 10** illustrates an example of property testing semantics.

[0022] **Figure 11** illustrates an example of a property testing algorithm.

[0023] **Figure 12** is a block diagram of one embodiment of a computer system.

DETAILED DESCRIPTION OF THE PRESENT INVENTION

[0024] A new approach to software testing is described herein. The approach adapts directed testing to generate tests that expose violations of specific properties that a code base is expected to obey. A heuristic is used to prune exploration of paths that are typically not relevant to the property of interest. This focused exploration of the program paths can generate test inputs that detect property violations much more effectively than previous directed testing techniques.

[0025] In the following description, the new approach is referred to as the property testing technique or property testing algorithm. The properties under test can be used in partial verification of software, such as type state properties. The property testing technique solves the problem of path explosion in directed testing by using a new heuristic that focuses on testing for property violations. The property testing technique is distinct from the conventional property-based approach in that the new approach does not rely on static slicing algorithms. Instead, the property testing technique described herein uses a new heuristic that is based on dynamic runtime information to prune possible execution paths, thus side stepping the problem with pure static analysis based approach.

[0026] Figure 1 illustrates an example of a property, namely locking property 100. The locking property 100 is described by a lock state 110, an unlock state 120 and an error state 130. State transitions can be achieved by calls to lock() and unlock(). The locking property 100 may be violated in a number of scenarios, e.g., when a test input causes two consecutive calls to lock() without an unlock() between them. In one embodiment, the property testing technique uses directed testing to identify test inputs which lead to violation of a property, e.g., test inputs that lead the property state to the error state 130. During execution (e.g., at runtime), the property testing technique detects branches in an application (also referred to as a program or code) that do not influence the property of interest (e.g., the locking property 100, a set_timer property, or any other properties of interest to a program developer or a user). These branches are referred to as “non-relevant branches.” As a non-relevant branch does not influence the property of interest, it is often not necessary to repeatedly explore both arms of the branch during testing for property violations. The property testing technique detects non-relevant branches and explores just one arm of such branches, thus drastically eliminating the cost of exponential path exploration. The property testing technique dramatically reduces the number of inputs that need to be explored to detect violations for some properties of interest.

[0027] Embodiments of the present invention provide a scalable automatic test generation technique that focuses on detecting important classes of errors. Given limited testing resources, the property testing technique focuses on critical

errors, such as errors that may encountered by a user during runtime. For example, the property testing technique may detect a runtime exception when a user attempts to add a contact to this address book. By focusing the test generation on critical errors, the demand on testing time and other resources may be greatly reduced.

[0028] In the following description, numerous details are set forth to provide a more thorough explanation of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form, rather than in detail, in order to avoid obscuring the present invention.

[0029] Some portions of the detailed descriptions which follow are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of steps leading to a desired result. The steps are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[0030] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "executing" or "collecting" or "determining" or "generating" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly

represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[0031] The present invention also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, CD-ROMs, and magnetic-optical disks, read-only memories (ROMs), random access memories (RAMs), EPROMs, EEPROMs, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[0032] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language. It will be appreciated that a variety of programming languages may be used to implement the teachings of the invention as described herein.

[0033] A machine-readable medium includes any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory ("ROM"); random access memory ("RAM"); magnetic disk storage media; optical storage media; flash memory devices; etc.

Overview

[0034] Figure 2 is a block diagram that illustrates an embodiment of a processing system 200 in which embodiments of the present invention may operate. The processing system 200 includes a compiler module 210, which receives a source application 213 and a property of interest 215, and compiles a source application

213 into a compiled application 214. For example, the source application 213 may be a program in the C programming language, and the compiled application 214 may be the corresponding object code. The compiler module 210 consults a rule module 220 for compiling the source application 213. The rule module 220 includes a collection of semantic rules for property testing (which will be described in greater detail with reference to Figure 10). According to the semantic rules, the compiler module 210 inserts code into the compiled application 214, such that during runtime, the compiled application 214 generates information necessary for property testing and collects the information in a data structure (e.g., stacks). Therefore, the compiled application 214 is also referred to as instrumented code; that is, code instrumented with the semantic rules. In one embodiment, the information includes a set of non-relevant branches with respect to a property of interest, as well as path constraints imposed by the symbolic variables of the compiled application 214. The processing system 200 further includes a test engine 240, which generates test inputs for the compiled application 214 based on the output from a previous iteration of the compiled application 214. The test engine 240 may include an equation solver 245 that solves the path constraints generated from the previous iteration. The test engine 240 repeatedly generates test inputs until all paths corresponding to relevant branches in the application 214 have been explored and no new path constraints are generated.

[0035] In an alternative embodiment, the compiler module 210 and the test engine 240 may reside in different processing systems. The compiler module 210 and the test engine 240 may be implemented by software, hardware, or a combination of both. The processing system 100 may be a server, a workstation, a personal computer (PC), a mobile phone, a palm-sized computing device, a personal digital assistant (PDA), etc.

[0036] The main idea of the property testing technique can be explained with a code segment 300 shown in Figure 3. Referring to Figure 3, the code segment 300 is to be tested with respect to the locking property 100 of Figure 1. It is assumed that “then func1”, “then func2”, “else func1” and “else func2” do not have any property changing methods (i.e. do not contain any calls to lock() or unlock()), and that p, q, r are identified inputs to the function.

[0037] In one embodiment, a processing system executes the code segment 300 both concretely on an input and also symbolically at the same time. The symbolic execution keeps track of symbolic memory and constraints on branches along the path of execution.

[0038] To illustrate the distinct features of the property testing technique, it is useful to understand how a generic concolic execution (such as CUTE) would be performed on the code segment 300. A generic concolic execution starts off by using randomly generated values for each of the inputs. Assume the initial random values for p, q, r are true, 15681, 20745, respectively. So, the code segment 300 executes the then branches at L1, L2, L4, and the else branch at L3. During execution, a symbolic path constraint is collected by the instrumented code, namely $(q_0 = \text{true}, p_0! = 20, r_0! = 100)$. To force the code segment 300 through a different branch in the next iteration, the processing system calculates a solution to the path constraint $(q_0 = \text{true}, p_0! = 20, r_0 = 100)$ obtained by negating the last predicate of the current path constraint. The solution (illustratively, $(q_0 = \text{true}, p_0 = 345345, r_0 = 100)$) is used as the input for the next run of the code segment 300. A generic concolic technique (such as CUTE) will repeat the above operations until all paths have been explored. The CUTE approach will explore paths corresponding to all combinations of the branches at L1, L2, L3 being true or false (8 paths in total) before concluding that the program does not have any errors with respect to the property of interest.

[0039] Embodiments of the present invention take note that only the value of q in Figure 3 affects the state of the program with respect to the property of interest, while values of p and r do not. A processing system (e.g., the processing system 200 of Figure 2) detects this behavior and exploits it by restricting subsequent iterations to just one arm of the branches at both L2 and L3. The processing system does this by also recording property states in addition to the symbolic constraints at each program point (e.g., at a branch point and a merge point of a branch). At every merge point of a branch (where the two arms of the branch join), the processing system checks if either arms of the branch (if they have been executed already) affect the property of interest. If both of the arms do not change the property of interest, then the branch is marked as a non-relevant branch.

[0040] Referring to the example of Figure 3, in the first iteration, the processing system according to one embodiment of the present invention uses random input values for p, q, r. During execution of the instrumented code, the processing system also keeps track of the property state along with the symbolic state and constraints. The path constraint ($q_0 = \text{true}, p_0! = 20, r_0! = 100$) is recorded along with a property state map $L1 \rightarrow \text{uninit}, L2 \rightarrow \text{Locked}, L3 \rightarrow \text{Locked}, L4 \rightarrow \text{Locked}$. It then proceeds to negate the last predicate of the path constraint and solves the constraints to obtain a new input (illustratively, ($q_0 = \text{true}, p_0 = 345709, r_0 = 100$), for the second iteration. After the second iteration, the processing system using a property testing heuristic to take note that both arms of the branch at L3 have been explored, the initial state at L3 and the state at the merge point L4 for both the branches is the same Locked. It then marks the branch at L3 when reached with state Locked as non relevant (that is, the path constraint for the branch needs not be negated in future iterations). In the third iteration, the solution for path constraint ($q_0 = \text{true}, p_0 = 20$) is explored (illustratively, ($q_0 = \text{true}, p_0 = 20, r_0 = 19746$)). At the end of the third iteration, the path constraint is ($q_0 = \text{true}, p_0 = 20, r_0! = 100$). This is where the property testing technique diverges from other existing concolic approaches (such as CUTE). Here, a processing system implementing the property testing technique takes note that the property state at L3 is Locked in the third iteration, and determines that negating the predicate at L3 is not necessary since the branch is unaffected when the property state is Locked. The processing system then proceeds to directly explore the solution of path constraint obtained by negating ($q_0 = \text{true}$) in the fourth iteration. Consequently, for the example above, the property testing technique checks only these six paths : (L1-then, L2-then, L3-else) (L1-then, L2-then, L3-then) (L1-then, L2-else, L3-else) (L1-else, L2-then, L3-else) (L1-else, L2-then, L3-then) (L1-else, L2-else, L3-else). As the number of non-relevant branches increases, the property testing technique achieves a greater benefit when compared to the CUTE approach.

Loop Handling

[0041] Figure 4 illustrates a loop example in which embodiments of the present invention have an order of magnitude benefit over the CUTE approach,

especially for branches inside of loops, while still retaining much of the ability to quickly detect common property violations.

[0042] One of the main constructs in C programs that results in exponential behavior in the CUTE algorithm is loops. Thus, the property testing technique in comparison to CUTE when dealing with loops deserves special mention. If a loop executes n iterations in a run and contains k branches, there will be $n*k$ branches that can possibly be flipped in the subsequent iterations. In the worst case, 2^{n*k} number of inputs will need to be generated in the subsequent iterations to completely cover all the possibilities. Due to this exponential nature of the CUTE algorithm, any sufficiently large loop iteration can quickly make the problem intractable. This is illustrated in a code segment 400 shown in Figure 4. Referring to Figure 4, in each iteration of the loop, there is a branch point resulting in 100 branches in the unrolled initial run of the program. In order to fully explore all these branches, CUTE generates 2^{100} different values for x . Depending on the implementation, CUTE may never proceed to analyze the branches either before the loop or after them.

[0043] In particular, if `do_some_thing` does not affect the property state, then the entire loop does not affect the property state and may never need to be explored to detect property violations.

[0044] In one embodiment, the property testing technique provides a solution to loop handling by computing branch relevancy based on the branch label. For all of the 100 branches, the branch label is the same in every iteration of the loop. So once it is detected that the branch in the loop does not affect the property state (e.g., within two input sequences), the branch is never flipped again (that is, it will not be necessary to explore the alternative arm of the branch), thus covering the entire loop within two iterations. Therefore, the property testing technique is much more efficient than existing concolic testing. As paths are efficiently pruned, the property testing technique can get to the critical portions (e.g., error-prone portions) of the program quickly.

Fault Detection

[0045] Figure 5 is a code segment 500 that illustrates a limitation of the property testing technique. Typically, programmers use guard variables to keep track of the property state within the program. In a scenario where a programmer incorrectly updates guard variables without any corresponding change in the property state, the property testing technique described herein may fail to detect some faults. Nevertheless, the property testing technique is able to detect a converse scenario in which a programmer updates the state but forgets to update the guard variable.

[0046] Referring to Figure 5, assume that a non-zero value for p and r is chosen as the initial test input. In this case, the property testing technique will detect that r is not a relevant branch after two iterations since there is no state change operation. In the third iteration when p is chosen to be false, the property testing technique does not explore both the arms of the second branch since it deems that branch is irrelevant. Consequently, in one scenario, it may choose to only explore the branch in which $r > 0$ and miss the error that occurs when p is false. This situation can be remedied by providing a rule that any changes to the guard variables count as a property state change either manually or automatically (e.g., by incorporating dynamic slicing or conservative static slicing).

Property Testing Method

[0047] Figure 6 is a flow diagram illustrating one embodiment of a method 600 for the property testing technique described herein. The method 600 may be performed by processing logic that may comprise hardware (circuitry, dedicated logic, etc.), software (such as that run on a general purpose computer system or a dedicated machine), or a combination of both. In one embodiment, the method 600 may be performed by the processing system 200 of Figure 2.

[0048] Referring to Figures 2 and 6, in one embodiment, the method 600 begins when the test engine 240 randomly generates values as test inputs to a compiled application (block 610). The processing system 200 executes the compiled application with the test inputs (block 620). During execution, the processing system 200 also executes the compiled application symbolically and

collects constraints among the symbolic variables (block 630). Concurrently, the processing system 200 also maintains a property state at every execution step (block 640). In an alternative embodiment, the property state is recorded at the branch point (where the arms of the branch fork) and the merge point (where the arms of a branch join) of every branch. The property state is recorded in a data structure, such as a stack. During execution, the processing system 600 compares the property state at the merge point of a not-yet-covered branch with the property state at its branch point (block 650). The processing system 600 determines whether the property state is the same at these two points, that is, whether the property state is unchanged by the corresponding branch operation (block 655). If the property state at these two points of the branch is the same, the arm of the branch taken by the execution is marked as non-relevant (block 660). Otherwise, the arm of the branch taken by the execution is marked as relevant (block 670). A branch is non-relevant if both of its arms are marked non-relevant.

[0049] Continuing from blocks 660 and 670, in one embodiment, the processing system 200 negates a constraint (e.g., the last constraint in the execution) only if its corresponding branch has not been determined as non-relevant (block 680). That is, potentially non-relevant branches are executed at least once before recognizing they are non-relevant. Negating a constraint means that the predicate controlling the direction of the branch is given a value that forces the execution in the next iteration into the alternative arm of the branch. For example, if the “then” arm is taken in the current iteration, negating the constraint means imposing constraints on the program variables such that the “else” arm of the branch will be taken in the next iteration.

[0050] Based on the constraints collected at block 630 and the negated constraint produced at block 680 (only if its corresponding branch has not yet been determined as non-relevant), the test engine 240 solves the constraints to generate new test inputs for a subsequent iteration (block 690). In the subsequent iterations, property state for the not-yet-covered branches is compared to update branch relevancy. The processing system 200 repeats the operations at blocks 620-690 until all paths corresponding to relevant branches have been explored and no new path constraints are generated. In one embodiment, the processing system 200 maintains

a data structure (e.g., a coverage stack) that keeps track of the branches covered by the iterations. Once the processing system 200 determines that all relevant branches of paths have been explored and no new path constraints are generated, the method 600 may terminate and generate no more test inputs.

Detailed Implementation

[0051] Before describing the detailed implementation of the property testing algorithm, it is useful to explain the syntax and semantics on which the algorithm is based. Figure 7 illustrates an abstract syntax of a core language used by the property testing algorithm, according to one embodiment of the present invention. Figure 8 illustrates the semantic rules used by a generic concolic testing algorithm (e.g., the CUTE algorithm) and Figure 9 illustrates the concolic testing algorithm. Figure 10 illustrates the semantic rules used by the property testing algorithm and Figure 11 illustrates the property testing algorithm, according to one embodiment of the present invention.

Syntax

[0052] Figure 7 illustrates an example of a simple imperative programming language (also referred to as the core language). A statement in this language is one of the following: (1) the abort statement signifying program error, (2) an input statement $x := \text{input}()$ that gets an external input into a variable x , (3) an assignment $x := E$, and (4) a labeled conditional statement $\text{if}(E) \text{ then } S1 \text{ else } S2$, where E is an expression and $S1, S2$ are statements in the program.

[0053] In Figure 7, $VEnv$, the variable environment, is a partial map between local variables and their values. IL is the list of inputs from which the program reads the inputs. The rest of the terms, including, $stateChange$, $popS$, $SymVarEnv$, $PathConstraints$, $CoverageStack$, $State$, $PropStack$, $NonRelevantBranches$, and $Branch-PropStateMap$ are not part of the core language or the concrete semantics, but are necessary for describing the CUTE semantics and the semantics for property testing to be discussed later. The details of the term $SymExp$ are omitted here. A description of $SymExp$ can be found in J. C. King,

“Symbolic execution and program testing,” *Communications of the ACM*, 19(7):385-394, 1976.

Concrete Semantics for Expressions

[0054] A function `eval_concrete` implements the concrete semantics for evaluating expressions in a program. It takes as input a `VarEnv` (a mapping from variables to their values) and an expression to evaluate, and returns a value.

$$\text{eval_concrete} : \text{VarEnv} \times \text{Expressions} \rightarrow \text{Value} \quad (1)$$

Symbolic Semantics for Expressions

[0055] In symbolic semantics, a mapping is maintained from variables to symbolic expressions along with a regular environment. This mapping is called a symbolic environment and is referred to as `SymVarEnv`. A function `eval_symbolic` implements the symbolic semantics for evaluating expressions in a program. The function takes as input a `VarEnv`, `SymVarEnv` and an expression, and returns a symbolic expression, `SymExp`.

$$\text{eval_symbolic} : \text{VarEnv} \text{ SymVarEnv} \text{ Expressions} \rightarrow \text{SymExp} \quad (2)$$

[0056] Concrete values of the variables (available from the `VarEnv`) can be used to simplify `SymExp` by substituting concrete values for symbolic expressions. The substitution may be performed whenever the symbolic expressions go beyond the complexity that can be handled by symbolic decision procedures.

CUTE Semantics and Testing Algorithm

[0057] Figure 8 illustrates the semantic rules for generic concolic execution of a program, such as CUTE. The semantic rules are discussed herein to illustrate the differences from the property testing semantics, to be described with reference to Figure 10. Referring to Figure 8, the symbolic constraint ζ maintains the current path constraint when executing the program. At every conditional statement `!if(e) S1 else S2`, if the concrete execution takes the then branch, the symbolic constraint ζ is updated to $\zeta \wedge (1, \text{eval_symbolic}(\text{VEnv}, \text{SEnv}, e) > 0)$ as shown in **R5**, **R5'**. If the

execution takes the else branch, then the path constraint ζ is updated to $\zeta < l$, $\text{neg}(\text{eval_symbolic}(\text{SEnv}, \text{VEnv}, e)) >$ as shown in **R4** and **R4'**.

[0058] During the generic concolic execution, a processing system maintains a coverage stack CS which is primarily used for bookkeeping and used in a test generation algorithm (see Figure 9). CS essentially maintains the coverage information from previous iterations of the testing algorithm useful for judging whether a branch has both its arms covered for a path prefix. $\langle l, < 0, c \rangle$ indicates that the else branch at l is covered and $\langle l, < 0, nc \rangle$ indicates that the else branch at l is not covered. Similarly, $\langle l, < 1, c \rangle$ indicates that the then branch at l is covered and $\langle l, < 1, nc \rangle$ indicates that the then branch at l is not covered. pos is a position marker on the coverage stack; $(CS\ pos\ \langle l, b, d \rangle\ CS1)$ indicates the next branch that might be executed will be at l . pos at the top of the coverage stack ($CS\ pos$) indicates that there is no coverage information for branches from this point on, i.e. this is the first time these branches are being explored with this path prefix. In **R4** and **R5** the top of the coverage stack is uncovered and is updated to be covered. These uncovered branches on top of the coverage stack denote the branches where the testing algorithm in the previous iteration executed one arm of the branch and in this iteration has forced the execution to proceed along the other arm of the branch. Specifically in **R4**, the previous execution took the then branch and the current execution is taking the else branch. As a result, the branch can be marked as covered. In **R4'** and **R5'** the next branch to be executed matches the predicted branch on the coverage stack, and the processing system proceeds by just moving the pos marker to the next branch. It is possible that the predicted branch on top of the coverage stack does not match the current branch, in which case the term is “stuck.” When “stuck” occurs, the processing system can detect it and restart the testing algorithm with a different set of random inputs. **R4''** and **R5''** denote the cases where the branches are being covered for the first time with the current path prefix, and coverage stack (CS) is updated by noting that both arms of these branches are not covered.

[0059] Semantic rule **R2** states: at every statement $x := \text{input}()$, a mapping $x \rightarrow \alpha_x$ is introduced from the variable x to a fresh symbolic value α_x . Semantic rule **R4** states: at every assignment $x := e$, the symbolic environment is updated with

mapping of x to the symbolic expression obtained by evaluating e in the current symbolic memory.

[0060] Assume `eval_concolic` is an implementation of the concolic semantics given above in Figure 8. It has following signature:

`eval_concolic : S X IL X CoverageStack -> PathConstraints X CoverageStack` (3)

`eval_concolic` initializes the `PathConstraints` to `nil`, `VEnv` to $\varnothing$ and `SEnv` to $\varnothing$ and uses the semantic rules above to implement the rest.

[0061] Given the concolic program execution semantics, the generic concolic testing generates a new test as follows. It selects a conditional `I: if(e) then S1 else S2` along the path that was executed such that (1) the current execution took the “then” (respectively, “else”) branch of the conditional, and (2) the “else” (respectively, “then”) branch of this conditional is uncovered. Let P' be the symbolic constraint just before executing this instruction and P be the constraint generated by the execution of conditional instruction e . Using a decision procedure, concolic testing finds a satisfying assignment for the constraint $P' \wedge \neg(P)$. The property of a satisfying assignment is that if these inputs are provided at each input statement, then the new execution will follow the old execution up to the beginning of the branch, but then takes the conditional branch opposite to the one taken by the old execution, thus ensuring that the other branch gets covered. Figure 9 shows an example of a function `solve_path_constraints` 900 that uses the satisfying assignment to define a new input list for the next run of the program. Figure 9 also shows an example of a concolic testing algorithm 910 that performs the generic concolic testing using the input list generated by the function `solve_path_constraints` 900.

Property Testing Semantics and Property Testing Algorithm

[0062] Embodiments of the present invention extend the generic concolic testing described above in Figures 8 and 9. According to one embodiment of the present invention, three additional pieces of information are collected by a processing system during each execution (i.e., iteration) of a program:

- (1) Property state of the program at every execution step during runtime.

This property state is changed only by the `stateChange` operation. In some

embodiments, the property state is collected at a branch point and a merge point of every branch.

- (2) A set of non-relevant branches. If an arm of a branch does not change the property state, then the set is updated.
- (3) A map from each branch to the corresponding property state in the execution.

[0063] Figure 10 shows an embodiment of the full semantics for property testing. The terms in bold show the extra terms for property testing (as compared to the concolic execution semantics in Figure 8). With respect to (1), a property state (referred to as **P**) of the system is introduced in the semantics. This property state, an integer in this formulation, can be modified only by the **stateChange** operation. In each execution of the program, the current property state of the program is maintained at every execution step. With respect to (2), a set of non-relevant branches (referred to as **NRB** in the syntax) is maintained and updated at every execution step. To compute relevancy, a processing system determines whether an arm of a branch changes the property state or not. The processing system makes the determination based on the knowledge of the property state at the beginning of the branch execution. The property state at the beginning of the branch execution can be obtained by maintaining a property state stack (referred to as **PS** in the syntax) at every execution step.

[0064] In one embodiment, a processing system starts the execution of a branch by pushing the current property state onto the stack (see **R4, R4', R4'', R5, R5', R5''** in Figure 10) and adding the **pop** operation to the body of the branch at the end. When the processing system executes the **pop** operation, it checks whether the current property state is the same as the top of the stack. If the property state stays the same, then the processing system adds the current branch to the set of non-relevant branches for that property state (see **R6, R7**). These non-relevant branches are intended to be updated and reused over many executions of the program in the testing algorithm. With respect to (3), a mapping is introduced from branches to property states (referred to as **BPM**) to maintain the property state for a branch in an execution. Overall, the different terms in the semantics and their meanings are as follows:

VEnv	Variable environment
SEnv	Symbolic Var Environment
ζ	Path constraint
CS	coverage stack
PS	property state stack for figuring out the relevancy of a branch
BPM	mapping from branches to property state in the current execution
NRB	records not relevant branches over all previous executions
P	current property state
IL	input list
S	statement in the program

[0065] The property testing semantics described above with reference to Figure 10 can be implemented by a function `property_test_semantics` with the following signature:

$S \times IL \times CoverageStack \times \mathbf{NRB} \rightarrow PathConstraints \times CoverageStack \times \mathbf{BPM} \times \mathbf{NRB},$ (4)

where `property_test_semantics` initializes the `PathConstraints` to `nil`, `VEnv` to φ , `SEnv` to φ , `PS` to `pos`, `BPM` to φ , `P` to 0 and uses the semantic rules in Figure 10 to implement the rest.

[0066] Figure 11 illustrates an example of a function `solve_path_constraints` 1100. A main difference between the function `solve_path_constraints` 900 (in Figure 9) and `solve_path_constraints` 1100 is in the latter function's use of non-relevant branches (NRB). The function `solve_path_constraints` 1100 not only looks for an uncovered branch (whose negation has not been explored) but also checks that both arms of the uncovered branch are relevant branches. Figure 11 also shows an example of a property testing algorithm 1120 that performs the property testing using the input list generated by the function `solve_path_constraints` 1100.

An Example of a Computer System

[0067] Figure 12 is a block diagram of an exemplary computer system that may perform one or more of the operations described herein. Referring to Figure 12, computer system 1200 may comprise an exemplary client or server computer system. Computer system 1200 comprises a communication mechanism or bus

1211 for communicating information, and a processor 1212 coupled with bus 1211 for processing information. Processor 1212 includes a microprocessor, but is not limited to a microprocessor, such as, for example, Pentium™, PowerPC™, Alpha™, etc.

[0068] System 1200 further comprises a random access memory (RAM), or other dynamic storage device 1204 (referred to as main memory) coupled to bus 1211 for storing information and instructions to be executed by processor 1212. Main memory 1204 also may be used for storing temporary variables or other intermediate information during execution of instructions by processor 1212.

[0069] Computer system 1200 also comprises a read only memory (ROM) and/or other static storage device 1206 coupled to bus 1211 for storing static information and instructions for processor 1212, and a data storage device 1207, such as a magnetic disk or optical disk and its corresponding disk drive. Data storage device 1207 is coupled to bus 1211 for storing information and instructions.

[0070] Computer system 1200 may further be coupled to a display device 1221, such as a cathode ray tube (CRT) or liquid crystal display (LCD), coupled to bus 1211 for displaying information to a computer user. An alphanumeric input device 1222, including alphanumeric and other keys, may also be coupled to bus 1211 for communicating information and command selections to processor 1212. An additional user input device is cursor control 1223, such as a mouse, trackball, trackpad, stylus, or cursor direction keys, coupled to bus 1211 for communicating direction information and command selections to processor 1212, and for controlling cursor movement on display 1221.

[0071] Another device that may be coupled to bus 1211 is hard copy device 1224, which may be used for marking information on a medium such as paper, film, or similar types of media. Another device that may be coupled to bus 1211 is a wired/wireless communication capability 1225 to communication to a phone or handheld palm device. System 1200 may also include an external network interface 1220 for providing network connections to external devices.

[0072] Note that any or all of the components of system 1200 and associated hardware may be used in the present invention. However, it can be appreciated that other configurations of the computer system may include some or all of the devices.

[0073] Whereas many alterations and modifications of the present invention will no doubt become apparent to a person of ordinary skill in the art after having read the foregoing description, it is to be understood that any particular embodiment shown and described by way of illustration is in no way intended to be considered limiting. Therefore, references to details of various embodiments are not intended to limit the scope of the claims which in themselves recite only those features regarded as essential to the invention.

CLAIMS

We claim:

1. A method for generating test inputs for testing an application, the method comprising:
 - executing, by a processing system, an application using test input values that are generated for testing the application with respect to a property of interest;
 - collecting constraints among symbolic variables of the application during execution;
 - determining whether a branch in the application is relevant to the property;
 - and
 - generating new test input values based on the constraints and branch relevancy.
2. An article of manufacture having one or more computer readable storage media storing instructions therein which, when executed by a system, causes the system to perform a method comprising:
 - executing, by a processing system, an application using test input values that are generated for testing the application with respect to a property;
 - collecting constraints among symbolic variables of the application during execution;
 - determining whether a branch in the application is relevant to the property;
 - and
 - generating new test input values based on the constraints and branch relevancy.
3. A method for testing an application, the method comprising:
 - generating, by a processing system, new test inputs for a next iteration based on test values of a current iteration and symbolic constraints collected during the current iteration; and
 - pruning execution paths using runtime generated information that includes relevancy of a plurality of branches in the application with respect to a property of interest.




Figure 1: Locking property




Figure 2 Processing System

300
↙

```
function f {  
L1: if (q) {  
    lock;  
    q = true;  
}  
  
L2: if (p != 20)  
    then_func1();  
    else  
        else_func1();  
  
L3:  if (r == 100)  
    then_func2();  
    else  
        else_func2();  
  
L4:  if (q) {  
    unlock;  
    q = false;  
}  
}
```


Figure 3 Running Example



A diagram consisting of the number '400' with an arrow pointing down and to the left towards the code block.

```
for (int i = 0; i < 100; ++i) {  
    CREST_int(x); //denote x as input  
    if (x == 40345) {  
        do_some_thing;  
    }  
}
```


Figure 4 Loop Example



A diagram consisting of the number '500' with an arrow pointing down and to the left towards the code block.

```
if (p) {  
    ... //doesnt change state or value of p  
} else {  
    ... //doesnt change state or value of p  
}  
  
if (r) {  
    p = 3 //forgets to update the state  
} else {  
    ... //doesnt change state or value of p  
}  
  
if (p) {  
    ... //doesnt change state or value of p  
} else {  
    state_change_error();  
}
```

Figure 5 An Example Showing a Limitation

**Figure 6**

Int n, m, i, j		
Vars Var		$x, y, \dots$
SymVars $SVar$		Obs, α_k
Expressions E	$::=$	$E \text{ op } E$ Var n
Label l		
Statements S	$::=$	abort — $x := \text{input}()$ — $x := E$ — $l: \text{if } (E) \text{ then } S \text{ else } S$ — $S ; S$ — $\text{stateChange}(m)$ — popS
VarEnv $VEnv$	$::=$	$VEnv[Var \longrightarrow n]$ ϕ
InputList IL	$::=$	$n ; IL$ ϕ
Terms for Concolic Semantics		
SymExp	SymExp	$\dots$ (omitted,)
SymVarEnv $SEnv$	$::=$	$SEnv[Var \longrightarrow \text{SymExp}]$ ϕ
Branch b	$::=$	$0 \mid 1$
Coverage d	$::=$	$nc \mid c$
PathConstraints ζ	$::=$	$\text{true} \mid \zeta \langle \text{Label}, \text{SymExp} \rangle,$
CoverageStack CS	$::=$	$CS \ CS \mid \langle \text{Label}, \langle \text{Branch}, \text{Coverage} \rangle \rangle \mid \text{pop}$
Terms for Property Testing Semantics		
State P	$::=$	n
PropStack PS	$::=$	$\phi \mid m ; PS$
NonRelevantBranches NRB	$::=$	$\phi \mid NRB \langle (\text{Label}, \text{Branch}), \text{State} \rangle$
BranchPropStateMap BPM	$::=$	$\phi \mid BPM \langle (\text{Label}, \text{Branch}), \text{State} \rangle$

Figure 7 Abstract Syntax for Core Language

- R1** $\frac{\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{IL}, \text{S1} \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}', \text{SEnv}', \zeta', \text{CS}', \text{IL}', e \rangle}{\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{IL}, \text{S1}; \text{S2} \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}', \text{SEnv}', \zeta', \text{CS}', \text{IL}', \text{S2} \rangle}$
- R2** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{IL}, x = E \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv} [x \leftarrow \text{eval_concrete}(\text{VEnv}, E)], \text{SEnv} [x \leftarrow \text{eval_symbolic}(\text{SEnv}, \text{VEnv}, E)], \zeta, \text{CS}, \text{IL}, e \rangle$
- R3** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS}, (m; \text{IL}), x = \text{input}() \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}[x \leftarrow m], \text{SEnv} [x \leftarrow \alpha_k], \zeta, \text{CS}, \text{IL}, e \rangle$ where $\alpha_k \notin \text{Range}(\text{SEnv})$
- R3'** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS}, (), x = \text{input}() \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}[x \leftarrow m], \text{SEnv} [x \leftarrow \alpha_k], \zeta, \text{CS}, \text{IL}, e \rangle$ where $\alpha_k \notin \text{Range}(\text{SEnv})$ and m is random.
- R4** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS pos } \langle l, \langle 1, \text{nc} \rangle \rangle, \text{IL}, l: \text{if } E \text{ then } S \text{ else } S' \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}, \text{SEnv}, \zeta \langle l, \text{neg}(\text{eval_symbolic}(\text{SEnv}, \text{VEnv}, E)) \rangle, \text{CS } \langle l, \langle 0, c \rangle \text{ pos, IL}, S' \rangle$ if $\text{eval_concrete}(\text{VEnv}, E) = 0$
- R4'** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS pos } \langle l, \langle 0, d \rangle \rangle \text{CS}', \text{IL}, l: \text{if } E \text{ then } S \text{ else } S' \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}, \text{SEnv}, \zeta \langle l, \text{neg}(\text{eval_symbolic}(\text{SEnv}, \text{VEnv}, E)) \rangle, \text{CS } \langle l, \langle 0, d \rangle \rangle \text{ pos CS}', \text{IL}, S' \rangle$ if $\text{eval_concrete}(\text{VEnv}, E) = 0$
- R4''** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS pos, IL}, l: \text{if } E \text{ then } S \text{ else } S' \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}, \text{SEnv}, \zeta \langle l, \text{neg}(\text{eval_symbolic}(\text{SEnv}, \text{VEnv}, E)) \rangle, \text{CS } \langle l, \langle 0, \text{nc} \rangle \rangle \text{ pos, IL}, S' \rangle$ if $\text{eval_concrete}(\text{VEnv}, E) = 0$
- R5** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS pos } \langle l, \langle 0, \text{nc} \rangle \rangle, \text{IL}, l: \text{if } E \text{ then } S \text{ else } S' \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}, \text{SEnv}, \zeta \langle l, \text{eval_symbolic}(\text{SEnv}, \text{VEnv}, E) \rangle, \text{CS } \langle l, \langle 1, c \rangle \rangle \text{ pos, IL}, S' \rangle$ if $\text{eval_concrete}(\text{VEnv}, E) \neq 0$
- R5'** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS pos } \langle l, \langle 1, d \rangle \rangle \text{CS}' \text{ end, IL}, l: \text{if } E \text{ then } S \text{ else } S' \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}, \text{SEnv}, \zeta \langle l, \text{eval_symbolic}(\text{SEnv}, \text{VEnv}, E) \rangle, \text{CS } \langle l, \langle 1, d \rangle \rangle \text{ pos CS}' \text{ end, IL}, S' \rangle$ if $\text{eval_concrete}(\text{VEnv}, E) \neq 0$
- R5''** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS pos, IL}, l: \text{if } E \text{ then } S \text{ else } S' \rangle \longrightarrow_{\text{stmt}} \langle \text{VEnv}, \text{SEnv}, \zeta \langle l, \text{eval_symbolic}(\text{SEnv}, \text{VEnv}, E) \rangle, \text{CS } \langle l, \langle 1, \text{nc} \rangle \rangle \text{ pos, IL}, S' \rangle$ if $\text{eval_concrete}(\text{VEnv}, E) \neq 0$
- R6** $\langle \text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{IL}, \text{abort} \rangle \longrightarrow_{\text{stmt}} \text{Error}$

Figure 8 Concolic Execution Semantics

900
↙

```

solve_path_constraints(CS,  $\zeta$ ) {
  find l, the top most branch in CS, such that it is not covered
  i.e.  $CS = CS1 \langle l, \langle b, nc \rangle \rangle CS2$  s.t. for all  $\langle l', \langle b, d \rangle \rangle \in CS2, d \neq c$ .
  if l does not exist then return (1, -); //we are done
  let  $(\zeta_1 \langle l, SE \rangle \zeta_2) = \zeta$ 
  if path constraints  $\zeta_1 \langle l, SE \rangle$  has a solution IL'
    return (0, update(IL, IL'), CS1  $\langle l, \langle b, nc \rangle \rangle$ )
  else
    return solve_path_constraints(CS1,  $\zeta_1$ )
}

```

910
↙

```

ConcolicTesting(S) {
  IL := empty; //empty input list
  CS := e; //empty coverage stack
  while (1)
    try ( $\zeta, CS'$ ) = eval_concolic(S, IL, CS)
    (done, IL, CS) = solve_path_constraint( $CS'$ ,  $\zeta$ )
    if (done) return;
  catch abortException
  print "bug found"
  return;
}

```

Figure 9 Concolic Testing Algorithm

- $$\begin{array}{l}
 \text{R1} \quad \frac{(\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \text{S1}) \longrightarrow_{\text{stmt}} (\text{VEnv}', \text{SEnv}', \zeta', \text{CS}', \text{PS}', \text{BPM}', \text{NRB}', \text{P}, \text{IL}', \epsilon)}{(\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \text{S1} ; \text{S2}) \longrightarrow_{\text{stmt}} (\text{VEnv}', \text{SEnv}', \zeta', \text{CS}', \text{PS}', \text{BPM}', \text{NRB}', \text{P}, \text{IL}', \text{S2})} \\
 \\
 \text{R2} \quad \frac{(\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \text{x} = \text{E}) \longrightarrow_{\text{stmt}} (\text{VEnv} [\text{x} \leftarrow \text{eval_concrete}(\text{VEnv}, \text{E})], \text{SEnv} [\text{x} \leftarrow \text{eval_symbolic}(\text{SEnv}, \text{VEnv}, \text{E})], \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \epsilon)}{(\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \text{x} = \text{E}) \longrightarrow_{\text{stmt}} (\text{VEnv} [\text{x} \leftarrow \text{eval_concrete}(\text{VEnv}, \text{E})], \text{SEnv} [\text{x} \leftarrow \text{eval_symbolic}(\text{SEnv}, \text{VEnv}, \text{E})], \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \epsilon)} \\
 \\
 \text{R3} \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, (m ; \text{IL}), \text{x} = \text{input}()) \longrightarrow_{\text{stmt}} (\text{VEnv} [\text{x} \leftarrow m], \text{SEnv} [\text{x} \leftarrow \alpha_k], \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \epsilon) \text{ where } \alpha_k \notin \text{Range}(\text{SEnv}) \\
 \\
 \text{R3}' \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \phi, \text{x} = \text{input}()) \longrightarrow_{\text{stmt}} (\text{VEnv} [\text{x} \leftarrow m], \text{SEnv} [\text{x} \leftarrow \alpha_k], \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \epsilon) \text{ where } \alpha_k \notin \text{Range}(\text{SEnv}) \text{ and } m \text{ is random.} \\
 \\
 \text{R4} \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS} \langle 1, \langle 1, 0 \rangle \rangle, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, 1: \text{if E then S else S}') \longrightarrow_{\text{stmt}} (\text{VEnv}, \text{SEnv}, \zeta \langle 1, \text{neg}(\text{eval_symbolic}(\text{SEnv}, \text{VEnv}, \text{E})) \rangle, \text{CS} \langle 1, \langle 0, 1 \rangle \rangle, (\text{PS} ; \text{P}), \text{BPM} [(1,0) \leftarrow \text{P}], \text{NRB}, \text{P}, \text{IL}, \text{S}' ; \text{popS}(1,0)) \text{ if } \text{eval_concrete}(\text{VEnv}, \text{E}) = 0 \\
 \\
 \text{R4}' \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS} \langle 1, \langle 0, d \rangle \rangle \text{CS}', \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, 1: \text{if E then S else S}') \longrightarrow_{\text{stmt}} (\text{VEnv}, \text{SEnv}, \zeta \langle 1, \text{neg}(\text{eval_symbolic}(\text{SEnv}, \text{VEnv}, \text{E})) \rangle, \text{CS} \langle 1, \langle 0, d \rangle \rangle \text{CS}', (\text{PS} ; \text{P}), \text{BPM} [(1,0) \leftarrow \text{P}], \text{NRB}, \text{P}, \text{IL}, \text{S}' ; \text{popS}(1,0)) \text{ if } \text{eval_concrete}(\text{VEnv}, \text{E}) = 0 \\
 \\
 \text{R4}'' \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS} \langle 1, \langle 1, d \rangle \rangle \text{CS}', \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, 1: \text{if E then S else S}') \longrightarrow_{\text{stmt}} (\text{VEnv}, \text{SEnv}, \zeta \langle 1, \text{neg}(\text{eval_symbolic}(\text{SEnv}, \text{VEnv}, \text{E})) \rangle, \text{CS} \langle 1, \langle 0, 1 \rangle \rangle, (\text{PS} ; \text{P}), \text{BPM} [(1,0) \leftarrow \text{P}], \text{NRB}, \text{P}, \text{IL}, \text{S}' ; \text{popS}(1,0)) \text{ if } \text{eval_concrete}(\text{VEnv}, \text{E}) = 0 \\
 \\
 \text{R5} \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS} \langle 1, \langle 0, \text{nc} \rangle \rangle, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, 1: \text{if E then S else S}') \longrightarrow_{\text{stmt}} (\text{VEnv}, \text{SEnv}, \zeta \langle 1, \text{eval_symbolic}(\text{SEnv}, \text{VEnv}, \text{E}) \rangle, \text{CS} \langle 1, \langle 1, c \rangle \rangle, (\text{PS} ; \text{P}), \text{BPM} [(1,1) \leftarrow \text{P}], \text{NRB}, \text{P}, \text{IL}, \text{S} ; \text{popS}(1,1)) \text{ if } \text{eval_concrete}(\text{VEnv}, \text{E}) \neq 0 \\
 \\
 \text{R5}' \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS} \langle 1, \langle 1, d \rangle \rangle \text{CS}', \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, 1: \text{if E then S else S}') \longrightarrow_{\text{stmt}} (\text{VEnv}, \text{SEnv}, \zeta \langle 1, \text{neg}(\text{eval_symbolic}(\text{SEnv}, \text{VEnv}, \text{E})) \rangle, \text{CS} \langle 1, \langle 1, d \rangle \rangle \text{CS}', (\text{PS} ; \text{P}), \text{BPM} [(1,1) \leftarrow \text{P}], \text{NRB}, \text{P}, \text{IL}, \text{S} ; \text{popS}(1,1)) \text{ if } \text{eval_concrete}(\text{VEnv}, \text{E}) \neq 0 \\
 \\
 \text{R5}'' \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS} \langle 1, \langle 0, d \rangle \rangle \text{CS}', \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, 1: \text{if E then S else S}') \longrightarrow_{\text{stmt}} (\text{VEnv}, \text{SEnv}, \zeta \langle 1, \text{neg}(\text{eval_symbolic}(\text{SEnv}, \text{VEnv}, \text{E})) \rangle, \text{CS} \langle 1, \langle 1, c \rangle \rangle, (\text{PS} ; \text{P}), \text{BPM} [(1,1) \leftarrow \text{P}], \text{NRB}, \text{P}, \text{IL}, \text{S} ; \text{popS}(1,1)) \text{ if } \text{eval_concrete}(\text{VEnv}, \text{E}) \neq 0 \\
 \\
 \text{R6} \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \text{abort}) \longrightarrow_{\text{stmt}} \text{Error} \\
 \\
 \text{R7} \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, (\text{PS} ; \text{P}'), \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \text{popS}(1,b)) \longrightarrow_{\text{stmt}} (\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB} [(1,b) \leftarrow \text{P}], \text{P}, \text{IL}, \epsilon) \text{ if } (\text{P} = \text{P}') \\
 \\
 \text{R7}' \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, (\text{PS} ; \text{P}'), \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \text{popS}(1,b)) \longrightarrow_{\text{stmt}} (\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \epsilon) \text{ if } (\text{P} \neq \text{P}') \\
 \\
 \text{R8} \quad (\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, \text{P}, \text{IL}, \text{stateChange}(m)) \longrightarrow_{\text{stmt}} (\text{VEnv}, \text{SEnv}, \zeta, \text{CS}, \text{PS}, \text{BPM}, \text{NRB}, m, \text{IL}, \epsilon)
 \end{array}$$

Figure 10 Property Testing Semantics

1100

↙

```

solve_path_constraints(CS,  $\zeta$ , BPM, NRB) {
    find l, the top most branch in CS such that it is not covered and both the branch and its the
    negation are relevant.
    i.e. CS = CS1 <l, <b, nc> > CS2 s.t. for all <l', <b', d> >  $\in$  CS2, d = c and if <l, not(b)>,
    BPM(l,b) >  $\notin$  NRB and <(l,b), BPM(l,b) >  $\notin$  NRB
    if l does not exist then return (1, --) : //we are done

    let ( $\zeta_1$  <l, SE >  $\zeta_2$ ) =  $\zeta$ 
    if path constraints  $\zeta_1$  <l, SE > has a solution IL'
        return (0, update(IL, IL'), CS1 <l, <b, nc> > )
    else
        return solve_path_constraints(CS1,  $\zeta_1$ )
}

```

1120

↙

```

PropertyTesting(S) {
    IL := empty; //empty input list
    CS := e; //empty coverage stack
    NRB := e; //initially every branch is relevant
    while (1)
        try ( $\zeta$ , CS', BPM, NRB) = property_test_semantics(S, IL, CS)
        (done, IL, CS) = solve_path_constraint(CS',  $\zeta$ , BPM, NRB)
        if (done) return;
    catch abortException
        print "bug found"
        return;
}

```

Figure 11 Property Testing Algorithm




FIG. 12