

ReAssert: Suggesting Repairs for Broken Unit Tests

Brett Daniel, Vilas Jagannath, Danny Dig, and Darko Marinov

University of Illinois at Urbana-Champaign
Urbana, IL 61801, USA

Email: {bdaniel3,vbangal2,dig,marinov}@illinois.edu

Abstract—Developers often change software in ways that cause tests to fail. When this occurs, developers must determine whether failures are caused by errors in the code under test or in the test code itself. In the latter case, developers must repair failing tests or remove them from the test suite. Repairing tests is time consuming but beneficial, since removing tests reduces a test suite’s ability to detect regressions. Fortunately, simple program transformations can repair many failing tests automatically.

We present ReAssert, a novel technique and tool that suggests repairs to failing tests’ code which cause the tests to pass. Examples include replacing literal values in tests, changing assertion methods, or replacing one assertion with several. If the developer chooses to apply the repairs, ReAssert modifies the code automatically. Our experiments show that ReAssert can repair many common test failures and that its suggested repairs correspond to developers’ expectations.

I. INTRODUCTION

Unit testing is becoming an important and widely practiced activity in software development. For example, reports from Microsoft show that 79% of developers use unit tests [1], and the code for unit tests is often larger than the application code under test [2]. Developers manually write (or automatically generate) unit tests for their application code and frequently run them while changing the code.

When changes cause existing unit tests to fail, developers should inspect the failures. There are two possible outcomes. First, if failures are caused by regressions, then developers must revise the application code to make the tests pass. Second, if failures are caused by tests that no longer reflect the intended behavior of the software, then the tests are *broken*—developers must repair the broken tests or remove them from the test suite. (In some cases, developers need to change both the application code and the test code.) Repairing broken tests is time consuming but often preferable to removing tests since removing tests reduces a test suite’s ability to detect regressions that appear in later versions of the software.

The scenario described above assumes that developers first change application code, then repair any broken tests. In contrast, test-driven development advocates changing the tests (to reflect the new requirements) before changing the application code (to implement these requirements) [3], [4]. However, even with advances in change-impact analysis [5]–

[7], developers do not always know *a priori* all tests which will be affected by a particular change, much less how a failure will manifest. Therefore, *a posteriori* repair remains the most common practice.

Unfortunately, developers may not take the time to repair broken unit tests. The problem becomes particularly acute when developers make a “deep” change that breaks many tests or when developers have a large test suite produced by an automatic test-generation tool [8]–[13]. Automatically generated tests tend to be more fragile than manually written tests, e.g., dozens of automatically generated tests can fail on a seemingly simple code change. Developers are often reluctant to manually repair a large number of broken tests—particularly if they were automatically generated—opting instead to remove them.

We present *ReAssert*, a novel technique and tool that suggests repairs for failing unit tests while retaining their power to detect regressions. When tests (manually written or automatically generated) fail, ReAssert can suggest changes to test code that cause the tests to pass. If the suggested repairs match the developer’s intentions, then ReAssert can repair the tests with one mouse click rather than a tedious editing process.

One can compare ReAssert to automated refactoring tools commonly included in modern integrated development environments. Both are code transformation tools that involve well-defined sequences of structural changes [14]. Developers could perform these transformations manually, but doing so is tedious and error-prone. Just as one can perform a refactoring by selecting an item from a menu, ReAssert allows one to repair a failing test at the push of a button.

Both automated refactorings and ReAssert perform non-trivial analyses prior to transformation. ReAssert differs from refactoring tools in that its transformations require analysis of a test’s runtime execution (i.e., the values that caused a test to fail) in addition to the static structure of the code. Also, ReAssert suggests repairs that change the behavior of test code (from failing to passing), while refactorings perform structural code changes but preserve code behavior.

The key challenge in repairing tests is to retain as much of the original test logic as possible. One could trivially “repair” a failing test by removing all of its code so that it

passes but reveals nothing about the correct (or incorrect) behavior of the application code. Our design of ReAssert follows these criteria:

- Make minimal changes: Retain as much of the test code as possible and leave application code unchanged.
- Only change if needed: If no changes cause a test to pass, then leave the test code unchanged. ReAssert may not repair all test failures, but those it does will pass.
- Require developer approval: Allow a developer to inspect, modify, and approve the suggested repairs.
- Produce understandable test code: Produce code that a developer can understand and could write manually; use normal method calls and assertions similar to any other unit test.

This paper makes the following contributions.

Idea: We propose automatic repair of failed unit tests as a means toward reducing the manual effort required to update test code.

Technique: We develop a technique that can automatically suggest repairs to the test code that make failing tests pass. While we present the technique for Java, the concepts generalize to other languages. The technique combines analysis of a test’s dynamic execution (to recover failing values and control flow) with analysis and transformation of the static structure of test code.

Tool: We implemented our technique in a tool called *ReAssert* that can repair Java unit tests written using the JUnit framework (<http://junit.org>). We also built a plugin for the Eclipse IDE (<http://eclipse.org>) that integrates seamlessly with the built-in test runner. It allows the developer to easily compare the failing and repaired test code and to update the code with one mouse click. ReAssert can repair many common types of test failures and is also extensible, allowing developers to write project-specific repair strategies.

Evaluation: We evaluate ReAssert’s effectiveness in three ways. First, we describe two case studies in which researchers used ReAssert to repair failures in their evolving software. Second, we perform a controlled user study to evaluate whether ReAssert’s suggested repairs match developers’ expectations. Third, we assess ReAssert’s ability to suggest repairs for failures in open-source projects, considering both manually written and automatically generated test suites.

Our tool and experimental results are publicly available at <http://mir.cs.illinois.edu/reassert>.

II. EXAMPLE

We illustrate the use of ReAssert through a simplified example based on the shopping cart application from the user study described in detail in Section V-B. The example considers a common code evolution scenario: a developer starts

```
1 public void testRedPenCoupon() {
2     Cart cart = new Cart();
3     cart.addProduct(new RedPen());
4     cart.addProduct(new RedPen());
5     cart.addCoupon(new AnniversaryCoupon());
6     assertEquals(3.0, cart.getTotalPrice());
7     assertEquals(
8         "Discount: -$3.00, Total: $3.00",
9         cart.getPrintedBill());
10 }
```

Figure 1. Example test

from application code and passing tests, receives a change request due to a change in requirements, and has to update the code and/or tests to match the new requirements. We discuss how ReAssert helps in repairing broken (manually written or automatically generated) tests after code changes.

In our example, the application logic is in the classes *Cart* (which represents a shopping cart that can contain products and coupons), *Product* (whose subclasses represent products such as red or black pens), and *Coupon* (whose subclasses represent certain discounts that may apply). The classes provide usual methods for adding/removing products and coupons to/from the cart, computing prices, printing bills, etc. One example coupon celebrates the store’s anniversary with a buy-one-get-one-free discount for any pen.

Figure 1 shows an example, manually written JUnit test for the anniversary coupon. In object-oriented code, each test creates one or more objects, invokes a sequence of methods on these objects, and *asserts* that certain properties of these objects hold. This test checks that the discount correctly applies when the cart contains two red pens, which in the example cost \$3 each without the coupon. JUnit’s `assertEquals` method compares the expected (first argument) value and the actual (second argument) value. This test is one of the many tests for coupons and products.

All the tests pass for the original code, but the requirements change: the anniversary coupon should now apply to black pens only. When the developer correctly changes the *AnniversaryCoupon* class, this test, as well as several others, fail because they no longer reflect the correct application logic.

ReAssert adds a “Fix Failing Tests” option to Eclipse’s JUnit runner. Choosing the option brings up a compare dialog such as that shown in Figure 2. The top part lists test files with suggested repairs, and the bottom right part shows the repairs. The user can inspect the repairs and modify the code if necessary. Pressing the “Confirm” button applies the repairs to the test code.

In this example, ReAssert suggests to repair the assertions at lines 6 and 7 by changing the expected values in the assertions. Note that ReAssert suggests repairs for both assertions at once, although the failure at line 6 would normally “hide” the failure at line 7. To determine these repairs, ReAssert re-runs the test to record the actual values

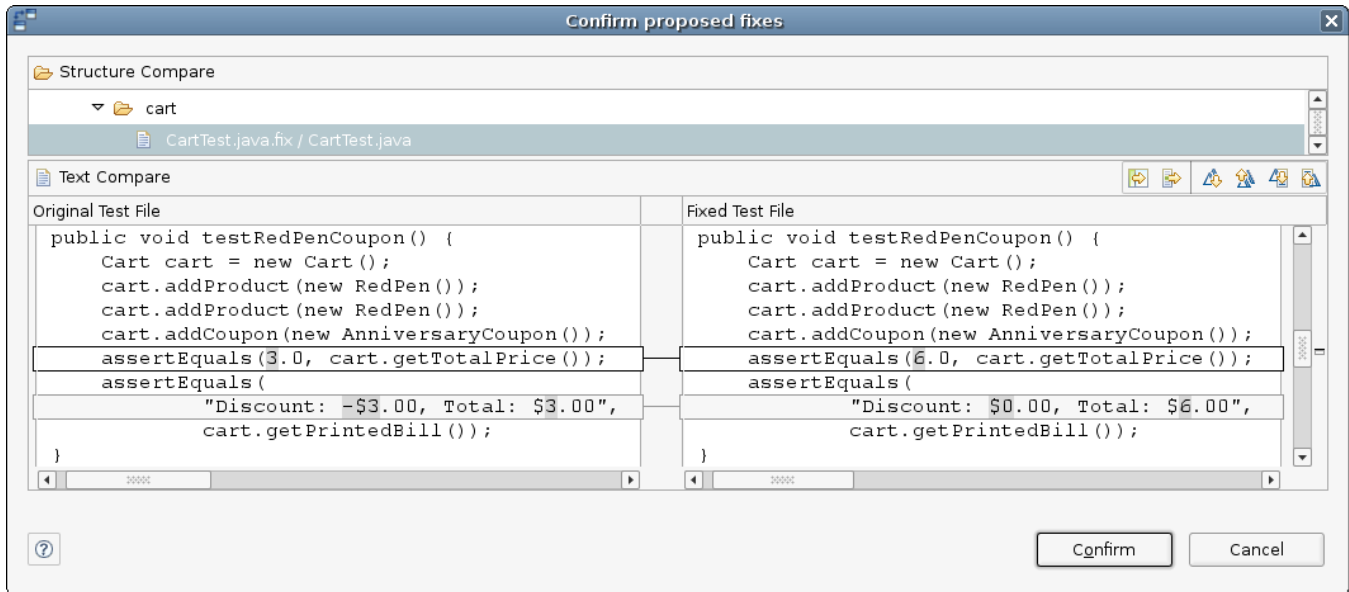


Figure 2. The suggested repair in Eclipse

```

1 public void testRedPenCoupon() {
2     Cart cart = new Cart();
3     cart.addProduct(new RedPen());
4     cart.addProduct(new RedPen());
5     cart.addCoupon(new AnniversaryCoupon());
6     checkCart(cart, 3.0,
7         "Discount: -$3.00, Total: $3.00");
8 }
9 protected void checkCart(
10     Cart cart, double total, String bill) {
11     assertEquals(total, cart.getTotalPrice());
12     assertEquals(bill, cart.getPrintedBill());
13 }

```

Figure 3. Example test with a helper method

in the assertions and then writes the recorded values back into the code.

Internally, ReAssert applied one of its repair strategies called *Replace Literal in Assertion*. This is just one of several strategies described in Section IV. It is a simple but effective strategy that illustrates two important characteristics of all repair strategies. First, the strategy requires analysis of the dynamic execution of the test. In this case, the actual values are a double and a string, which can be written directly as literals into the code. Second, each strategy is tailored to a particular code structure. The *Replace Literal in Assertion* strategy only applies to an `assertEquals` invocation located in the test method (not a helper method). If these conditions are not met, ReAssert chooses a different strategy.

For example, Figure 3 shows another test for the same functionality, but which uses a helper method to check the price and bill. Developers commonly write such helper meth-

ods to reuse several assertions across many tests. (One of the participants in our user study wrote such a helper method for several tests.) The *Replace Literal in Assertion* strategy does not apply because changing the failing assertions directly in the helper method would break other tests that call the same helper method.

ReAssert can still suggest a repair for this test but applies another strategy, *Trace Declaration-Use Path*. This strategy traces a variable's use from an assertion back to its definition (across the dynamic call chain from the test execution) and then writes the values recorded from the test run into the definition. This strategy exploits the fact that test helpers rarely modify the expected value, so it suffices to track declarations rather than control flow. In this example, the strategy changes the arguments to `checkCart` on line 6.

```

6 checkCart(cart, 6.0,
7     "Discount: $0.00, Total: $6.00");

```

While the previous two examples have shown the use of ReAssert to repair manually written tests, ReAssert is also useful for repairing and maintaining automatically generated test suites. For example, we applied the Randoop tool [13] on the shopping cart application, and Randoop generated a regression test suite with 1,598 tests (for the default time limit for generation). The assertions in those tests are similar to the assertions in Figure 1, but tests generated by Randoop are usually longer than manually written tests. After changing the shopping cart code (Section V-B), 834 of these tests failed. We applied ReAssert, and it repaired all tests such that they passed.

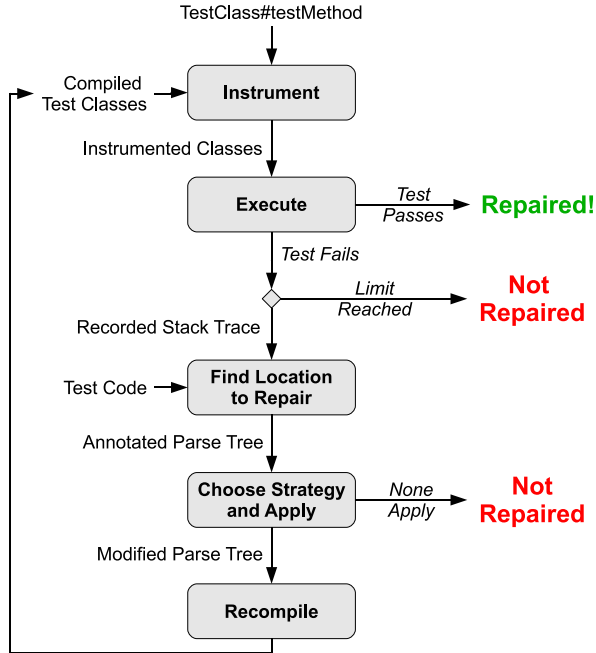


Figure 4. ReAssert’s repair process

III. REPAIR PROCESS

ReAssert’s repair process starts when the user chooses a set of failing tests to repair. For each test (which can have more than one failing assertion), ReAssert iteratively attempts repairs until the test passes, no strategies apply, or the iteration limit is reached.

To repair a single failure in a test, ReAssert follows the five-step process shown graphically in Figure 4. ReAssert first instruments the test classes to record values of method arguments for failing assertions (Section III-A). It then re-executes the test and catches the failure exception that contains both the stack trace and recorded values (Section III-B). It next traverses the stack trace to find the code to repair (Section III-C) and examines the structure of the code and the recorded values to choose a repair strategy to apply to change the code (Section III-D). It finally recompiles the code changes and repeats these steps if the test has another failure (Section III-E).

A. Instrumenting Assertions

ReAssert instruments some methods to dynamically record the runtime values that the repair strategy will use when repairing the test code. ReAssert does not instrument all methods; only those that are assertion methods and are declared in classes referenced (perhaps indirectly) by the current test. By default, JUnit’s standard assertions (such as `assertEquals`) are instrumented as assertion methods, and ReAssert provides a general mechanism for instrumenting additional methods.

```

public static void assertEquals(
    Object expected,
    Object actual) {
    try {
        // ... original code of assertEquals
    } catch (Error e) {
        throw new RecordedAssertFailure(
            e, expected, actual);
    }
}

```

Figure 5. Instrumentation of an assertion method

The instrumentation records values by wrapping the exception thrown by a failing assertion with an exception holding the values of the assertion’s arguments. Recalling the example from Figure 1, ReAssert instruments the `assertEquals` method to record the expected and actual values. Normally, `assertEquals` throws a `java.lang.AssertionError`. ReAssert instruments the method such that it catches any exception and re-throws ReAssert’s own `RecordedAssertFailure`. This exception holds the original exception and the expected and actual values. Figure 5 shows conceptually how ReAssert instruments the method.

This form of instrumentation has several good characteristics. First, it re-throws an exception only on the failing assertions that ReAssert cares about. There is no global state and no runtime overhead for non-failing calls. Second, ReAssert instruments bytecode dynamically using Java’s class loaders, allowing one to instrument arbitrary methods without requiring recompilation or source code (though ReAssert necessarily requires the source of *calls* to the method so it can repair them). Finally, one can record values from multiple assertions at different points on the call stack. These last two benefits enable custom repair strategies described in Section IV-H.

B. Re-Executing Tests

ReAssert uses JUnit to re-execute failing (instrumented) tests. JUnit catches the exception thrown by a failing test (likely a `RecordedAssertFailure`) and returns it to ReAssert. Even though ReAssert delegates to JUnit in our current implementation, ReAssert is effectively decoupled from a particular test framework. It could just as easily repair tests written using TestNG (<http://testng.org/>) or another framework. ReAssert simply requires the exception thrown by the test. This exception contains recorded values of method arguments and information about the cause of the test failure, including the stack trace.

C. Finding the Code to Repair

Repair strategies explore the stack trace of a failing test to find the static code location to repair. Java’s stack traces include source code line numbers as shown in Figure 6. ReAssert parses the test code into an *annotated parse*

```

edu.illinois.reassert.RecordedAssertFailure:
org.junit.AssertionFailedError:
expected:<3.0> but was:<6.0>
  at org.junit.Assert.assertEquals(Assert.java:116)
  at CartTest.checkCart(CartTest.java:11)
  at CartTest.testRedPenCoupon(CartTest.java:6)
  ...
  at org.junit...JUnitCore.run(JUnitCore.java:109)
Caused by: java.lang.AssertionError:
expected:<3.0> but was:<6.0>
  at org.junit.Assert.fail(Assert.java:71)
  at org.junit.Assert.failNotEquals(Assert.java:451)
  at org.junit.Assert.assertEquals(Assert.java:99)
  at org.junit.Assert.assertEquals(Assert.java:116)
  ...

```

Figure 6. Stack trace for the failure shown in Figure 3

tree [15] that contains all line numbers and type information. Thus, given a location from the stack trace, a particular strategy can easily find the appropriate code fragment.

Most repair strategies simply change the failing assertion’s call site one stack frame below the instrumented method. Some strategies require a more complex analysis. For instance, the *Trace Declaration-Use Path* strategy follows the stack trace to find the location in the source code of the literal value that flowed into the expected side of a failing `assertEquals` call (if such a value exists). Figure 3 shows one test that uses this strategy, and Figure 6 shows the stack trace produced by the failure of the assertion at line 11.

The *Trace Declaration-Use Path* strategy examines the code at the failing assertion’s call site (line 11) and finds that the expected side is a variable. The variable is declared as the second formal argument of the helper method `checkCart`, so the strategy retrieves the code one stack frame lower at line 6. The tracing stops because the second actual argument of the call is the literal value `3.0`, and the strategy can replace it with the recorded value. Had the argument been another variable, tracing would have continued, following the dynamic call chain from the stack trace and looking for declaration of variables. (Since test code rarely changes method arguments, it suffices to follow declarations rather than control flow.) The repair may occur many frames lower in the stack trace than the failed assertion.

D. Repairing Test Code

ReAssert chooses a repair strategy by looping over all strategies to determine which one can apply to the given failure. ReAssert passes the exception, the parse tree, and the recorded values to each strategy. Each strategy checks a set of *preconditions* to determine whether it can repair the failure. The first strategy whose preconditions are satisfied updates the parse tree. If no strategies apply, then ReAssert cannot repair the failure. The next section describes individual strategies in more detail. We give here only one example.

In the test shown in Figure 1, the *Replace Literal in Assertion* strategy applies. Figure 7 shows this strategy’s simplified pseudocode. Its preconditions require that the

```

public class ReplaceLiteralStrategy
  implements RepairStrategy {
  public void fix(
    Error failure,
    MethodInvocation assertion,
    Object[] recordedArgValues) {
    if (failure instanceof AssertionError
        && "assertEquals".equals(assertion.getName())
        && isLiteral(recordedArgValues[1])) {
      assertion.setArgument(0, recordedArgValues[1]);
    }
  }
}

```

Figure 7. Pseudocode for the *Replace Literal in Assertion* strategy

failure be an instance of `java.lang.AssertionError`, the failing method be `assertEquals`, and the recorded actual value can be written directly into the code. (There is also the implicit precondition that the expected side is not a variable, since if it was, the *Trace Declaration-Use Path* strategy would apply.) The test code in Figure 1 meets these conditions, so the strategy writes the recorded actual value into the expected side of the `assertEquals` as shown in Figure 2.

E. Iteratively Repairing Failures

Once a repair strategy completes, ReAssert recompiles the changed test code and repeats the entire repair process until the test passes, no more strategies apply, or a limit on the maximum number of repair attempts is reached. In this way, ReAssert is able to repair multiple assertions in a single test: it repairs the first, recompiles, repairs the second, and so on.

IV. REPAIR STRATEGIES

ReAssert’s library of repair strategies determines its ability to repair broken tests in meaningful and useful ways. ReAssert currently implements seven general repair strategies for common types of assertions and failures. Additionally, ReAssert is extensible so that developers can add new strategies tailored to a particular project or test suite.

Due to space limitations, this section briefly lists all the strategies. See [16] for more information, including the types of failures that can be repaired, how each strategy modifies test code, and applications of each strategy on examples encountered in the evaluation described in Section V.

A. Replace Assertion Method

This strategy replaces a failing assertion with a similar assertion that passes or that can be repaired by another strategy. For example, if the argument to `assertTrue` is a call to `equals` or the identity operator (`==`), then the strategy creates an `assertEquals` invocation, which other strategies can repair when the test is re-executed. This strategy often serves as a “preprocessing step” for other strategies.

B. Invert Relational Operator

This strategy inverts a relational operator in the argument to an `assertTrue` or `assertFalse` assertion. This approach works particularly well for assertions against objects that implement the `java.lang.Comparable` API.

C. Replace Literal in Assertion

This strategy replaces the expected (left) side of an `assertEquals` assertion with the literal value computed by the actual (right) side. This strategy is applicable to primitive types and certain reference types such as `String` and `Class` that can be written directly into code. The running example in sections II and III describes this strategy in more detail.

D. Replace with Related Method

This strategy applies when the argument to `assertTrue` or `assertFalse` is a call to a common library method that is closely related to another. There are many examples of such related methods in common APIs, including `java.util.Date`'s `after` and `before`, `String`'s `contains` and `indexOf`, and `java.util.Collection`'s `isEmpty` and `size`. In each case, the assertion fails on a boolean accessor (e.g., `isEmpty`) and can be repaired by asserting against the value returned from the related method.

In the following example, a call to `isEmpty` on a non-empty collection object causes the `assertTrue` to fail.

```
assertTrue(collection.isEmpty());
```

Rather than trivially changing `assertTrue` to `assertFalse`, this strategy asserts against the actual size of the nonempty collection:

```
assertEquals(5, collection.size());
```

E. Trace Declaration-Use Path

It is common for developers to write helper methods that bundle several assertions for reuse across many tests. Repairing such helper methods requires tracing an argument used in a failing assertion back to its definition (following dynamic call chain and static declaration-use paths) and replacing the value there. Section III-C explains this process in detail.

F. Accessor Expansion

All of the previous strategies operate on primitive values; this strategy differs in that it repairs failures related to object inequality. It replaces a failing `assertEquals` whose arguments are reference types with a list of assertions that test values returned from both arguments' accessor methods.

An `assertEquals` likely fails because a small number of the actual object's fields are not equal to the expected object's. This strategy makes the differing values explicit by matching the actual side's accessors with the expected side's. If the two differ, then the strategy asserts against a literal value or, if the accessors' values are reference

types, expands further. This process is similar to how some automatic test generation [10] or test augmentation tools [12] produce assertions.

G. Surround with Try-Catch

Sometimes developers change code to throw an exception rather than return an error value or silently fail. In these cases, the exception is expected, and a test should verify that it was thrown. This repair strategy surrounds a failing method call with a try-catch block that asserts that a particular exception is caught.

H. Custom Repair Strategies

ReAssert provides an extension API which can be used to instrument arbitrary methods and to define custom repair strategies. This capability allows one to repair application-specific failures or tests written in a custom test framework. An extension needs to provide a class that implements the repair strategy interface and give ReAssert the names of methods to instrument. Section V-A2 describes a case study that required custom repair strategies to update the content of external files.

V. EVALUATION

To evaluate ReAssert's effectiveness, we answer the following research questions:

- Q1 How often can ReAssert suggest repairs for failing unit tests?
- Q2 Are the repairs suggested by ReAssert usable for developers? Do the developers accept the changes suggested by ReAssert?
- Q3 Does ReAssert give false confidence to the developers such that they mask regression errors by repairing failing tests?

We describe two case studies in which ReAssert was used to repair failures in evolving research software, which provided some quantitative data for Q1 and Q2, and some qualitative data for Q3. We also performed a controlled user study that addressed all three questions directly. For Q1, we assess ReAssert's ability to repair failures in open-source applications. Finally, we discuss reasons why ReAssert is unable to repair certain failures.

A. Case Studies

We asked two teams of researchers to try ReAssert when changes to their evolving research software caused unit tests to fail. ReAssert successfully (Q1) and usefully (Q2) repaired the failures and revealed regressions in the software (Q3). The following sections describe these case studies.

Project	Tests	Failures	Accepted Repairs
Basset	-	17	15 (88%)
DPJ compiler	86	6	5 (83%)
DPJizer	78	14	9 (64%)

Figure 8. Using ReAssert to repair tests in Basset, DPJ, and DPJizer projects.

```

1 public class TestReductionModes
2 extends TestActor {
3     private static String language = ...
4     ...
5     @Test public void clientserver_jpfcomp() {
6         options = ...
7         runAndCheck(language, Driver.class, options,
8             null, 5, 1, 24, 2, 5, 0, 0);
9     }
10    ...
11 }

```

Figure 9. Example Basset unit test

1) *Case Study: Basset*: Java PathFinder (JPF) [17] is a system for verifying Java programs using an explicit-state model checker that acts on Java bytecode. JPF offers a powerful extension mechanism that allows one to verify programs in many domains. Basset [18] is one such extension currently under development that verifies programs that utilize the actor model of concurrency. This extension has a large suite of JUnit tests that often break as both JPF and Basset improve.

Figure 9 shows a simplified example of a Basset unit test. Each test does nothing except pass several arguments to a helper method called `runAndCheck` declared in the superclass `TestActor`. This method executes Basset using the first four arguments. After Basset completes, `runAndCheck` verifies that the latter seven arguments (which we will call the “postcondition numbers”) match numbers derived from Basset’s internal state.

These tests are exceptionally fragile and difficult to repair manually. First, the postcondition numbers often change as Basset improves. Second, The postcondition numbers are difficult to calculate prior to running the test (yet reasonably easy to check afterward). Third, each test takes tens of seconds to run due to JPF’s high overhead. Therefore, when tests failed due to changes in Basset, the developers were forced to spend many minutes manually re-running each failing test to calculate each of the seven postcondition numbers.

ReAssert removes much of this manual labor since it automates re-execution and can fix multiple failures. Internally, ReAssert applies the *Trace Declaration-Use Path* strategy to replace postcondition numbers.

On one failing test run, shown in Figure 8, the Basset developers accepted 15 of ReAssert’s 17 suggested repairs (Q1 and Q2). This corresponds to replacing 105 postcondition numbers. The remaining two failures were

due to null pointer dereferences, which, while repairable using the *Surround with Try-Catch* strategy, were caused by regressions (Q3).

2) *Case Study: Deterministic Parallel Java*: Deterministic Parallel Java (DPJ) [19] is an extension to the Java language. DPJ gives static guarantees that a program that type-checks with the DPJ compiler will be deterministic (giving the same output for the same input) when executed by any number of parallel tasks. DPJ incorporates a sophisticated effect system to verify noninterference of parallel tasks. DPJizer [20] is a tool that infers effect annotations in DPJ programs.

Because the DPJ language is an evolving research language, the DPJ compiler and the DPJizer tool must also evolve. This evolution frequently causes tests to break and requires a substantial manual effort to repair them.

Both DPJ compiler’s and DPJizer’s unit tests verify that the output from the tool matches the expected content in a prepared file. We implemented a custom repair strategy that is conceptually similar to the *Replace Literal in Assertion* strategy: it must record the expected content and write it to the expected file.

This strategy instruments two “levels” of assertions. The highest level is a helper method called `assertFilesAreTheSame` that takes the expected and actual file names. The expected file contains the prepared output, and the actual file contains the output from DPJ/DPJizer. The helper method reads the content of both files to strings that it passes to JUnit’s `assertEquals`. Thus, the instrumented `assertFilesAreTheSame` records the expected file name, and `assertEquals` records the actual content.

Figure 8 shows how many tests were repaired with ReAssert and this custom strategy. The repairs were correct—in that they would have caused the test to pass (Q1)—in all cases, and the developers accepted 14 out of 20 of the suggested repairs (Q2). The remainder revealed regressions in the system under test (Q3).

The developers expect to continue using ReAssert as DPJ and DPJizer evolve. They commented that ReAssert was useful not only for updating existing tests but also for writing new tests: they left the expected content blank, and ReAssert filled it in with the actual content.

B. Controlled User Study

We performed a controlled user study to evaluate how ReAssert assists developers in writing unit tests. We quantitatively answered all three of our research questions and qualitatively assessed ReAssert’s usefulness using a simple but realistic development project.

1) *Participants*: Our study had 18 participants: 13 graduate students, 3 undergraduate students, and 2 industry professionals. They had an average of 5.1 years of experience with Java, 2.4 years with JUnit, and 2.7 years with Eclipse.

Control Group			
Test Suite	Failures	Repairs	Matches
User written	26	26 (100%)	25 (96%)
Provided	47	47 (100%)	42 (89%)
ReAssert Group			
Test Suite	Failures	Repairs	Matches
User written	19	15 (79%)	12 (80%)
Provided	43	43 (100%)	34 (79%)

Figure 10. User study results.

No rewards were offered for participating in the study. Invitations to the study were sent out through class and departmental mailing lists as well as individual emails.

We randomly split the participants into two groups: a control group who were asked to perform the tasks described below using the Eclipse IDE but without ReAssert, and a ReAssert group who were asked to perform the tasks using the Eclipse IDE with ReAssert. In the end, nine members of the control group and nine members of the ReAssert group finished the study. One participant from the ReAssert group did not finish all tasks because the participant used an older version of ReAssert that did not have the *Trace Declaration-Use Path* strategy.

2) *Tasks*: The participants were given the application code and passing unit tests for the shopping cart application introduced in Section II. They were asked to first familiarize themselves with the code/tests and then to perform the following five tasks:

Task 1: Write some unit tests of their own to test previously untested functionality.

Task 2: Implement a requirement change which could potentially cause some of *their tests* to fail.

Task 3: Repair all failing tests.

Task 4: Implement another requirement change which would cause some of the initially *provided tests* to fail.

Task 5: Repair all failing tests.

The participants were asked to perform the tasks in order and click on a button at the end of each task to indicate completion. This helped us record the state of the code and tests at the end of each task.

3) *Quantitative Results*: The quantitative results of the user study provide some data for all three of our research questions.

Q1: The first two columns of Figure 10 summarize the number of failures caused by the participants' code changes in tasks 2 and 4 and the number of those failures that ReAssert could repair such that they pass. Note that the participants from the control group did not actually run ReAssert, so we evaluated post-mortem whether ReAssert would have suggested a repair that made their failing tests pass.

In total, ReAssert could repair 98% (131 of 135) of failures caused by the participants' code changes. Three of the failures were not repairable because a participant was using an older version of ReAssert which did not implement the *Trace Declaration-Use Path* strategy; those failures can now be repaired with the latest version of ReAssert. The remaining unrepaired failure was caused by the nondeterministic iteration order of `java.util.HashMap`; common practice dictates that tests should be deterministic, and thus this test should have been written differently.

Q2: The final column of Figure 10 summarizes the number of ReAssert's suggested repairs that *exactly matched* the changes made by the participants to repair the failing tests in tasks 3 and 5. We determined this match by structurally comparing tests recorded at the end of task 3 or 5 with the tests obtained by applying ReAssert's suggested repairs on the code written by the participants in the preceding task 2 or 4, respectively. Our structural comparison ignored white spaces, comments, variable names, and other code differences that don't affect code behavior.

For task 3, in which participants repaired tests they had written themselves, 90% (37 of 41) of ReAssert's suggested repairs matched the participants' repairs. For task 5, in which participants repaired tests provided to them, 84% (76 of 90) of ReAssert's repairs matched.

These results show that the repairs suggested by ReAssert are very useful because they very frequently match the repairs made by the participants. In cases where the repairs suggested by ReAssert did not match, the participants' repairs most often involved changes to the test setup or the addition of new assertions. We suspect that ReAssert could have been useful even in these cases, and thus the exact match is the lower limit for the number of usable repairs.

The participants found the *Replace Literal in Assertion* strategy particularly useful when they changed how the shopping cart printed the final bill. Those without the tool most often copied and pasted the output from JUnit, essentially replicating this strategy manually.

Q3: The participants introduced a total of 20 bugs in their code. We expected some bugs since our requirements were fairly vague, simulating real-life situations. Eleven of the bugs were introduced in task 2, and nine were introduced in task 4. Twelve of the bugs were introduced by participants who used ReAssert, and eight were introduced by participants who did not use ReAssert. One of the reasons why ReAssert users introduced more bugs could be that they become overly reliant on the tool. This can be mitigated by training users to carefully inspect the repairs suggested by ReAssert rather than accepting them blindly.

4) *Qualitative Results*: We obtained qualitative feedback through a simple questionnaire and post-study interviews. We modeled our questionnaire and analysis after a study by Saff and Ernst [21].

Question	Responses
Useful for the study	very 50%, yes 38%, no 12%
Use for own projects	yes 56%, maybe 33%, no 11%
Include in Eclipse	yes 67%, maybe 33%
Recommend to others	yes 67%, maybe 33%

Figure 11. Responses to Qualitative Questions

Project	Versions	Description
Checkstyle	3.0, 3.5	Code style checker
JDepend	2.8, 2.9	Design quality metrics
JFreeChart	1.0.7, 1.0.13	Chart creator
Lucene	2.2.0, 2.4.1	Text search engine
PMD	2.0, 2.3	Java program analysis
XStream	1.2, 1.3.1	XML serialization

Figure 12. Subject applications

Questionnaire Responses: We asked the nine participants in the ReAssert group four questions after they completed their tasks. Figure 11 summarizes their responses. The majority of the participants found ReAssert useful for the user study tasks and thought that it would also be useful for their own development and testing tasks. Most would recommend ReAssert to other people and thought that it should be included as part of the Eclipse IDE.

Positive Feedback: The participants who found ReAssert useful liked the fact that ReAssert repairs all the assertions in a test method at once rather than individually. Some participants mentioned that they would write more tests if they had a tool like ReAssert to help them maintain tests. They also noted that some assertions were very tedious to write (especially those involving large strings) and mentioned that they used ReAssert to build such assertions by leaving assertions “empty” and letting the tool compute the real expected value. This shows that ReAssert could also be helpful when creating new tests.

Cautionary Feedback: Most participants expressed concerns that ReAssert could induce carelessness in developers. These concerns do hold merit as shown by the results for Q3 in Section V-B3. However, we believe that this problem can be mitigated by explicitly informing developers that they should carefully inspect each repair suggested by ReAssert before applying it.

C. Failures in Open-Source Projects

We address Q1 by measuring how many failed tests ReAssert can successfully change such that they pass. Doing so requires many examples of failing tests in real applications. We consider the open-source projects listed in Figure 12. All are widely used and actively developed Java applications whose source code is available online. Several have been used in previous studies [22]–[24]. More importantly, they each had a large, manually written test suite and evolved in ways that would cause some tests to fail.

```
public void testclasses25 () {
    ...
    int var25 = 0;
    ArrayDeque var26 = new ArrayDeque (var25);
    ...
    short var28 = 1;
    var26.addLast (var28);
    Object var30 = var26.peekLast ();
    Object var31 = var26.pop ();
    ...
    assertEquals (1, var30);
    assertEquals (1, var31);
}
```

Figure 13. A simplified test generated by Randoop

We executed both manually written and automatically generated test suites. Manually written tests came from the projects’ source code repositories. For a subset of the projects, we generated many regression tests using Randoop [13], a feedback-directed random test generator.

We obtained failing tests by applying tests (manually written or automatically generated) from an early version of the software to a newer version of the system under test. This procedure mimics the evolution of the system under test, but differs in that the tests would have also evolved. We removed any tests that did not compile, ran the tests, and recorded the number of failures. ReAssert attempted to repair the failures, and we recorded the number of successful repairs as well as the repair strategies used.

1) *Generating Tests with Randoop:* Randoop is one example of an automatic test generator. Given a system to test, it produces a regression test suite with minimal user input. It does so using a process called *feedback-directed test generation*. This process starts from short method sequences that it extends by adding randomly selected method calls. It runs the resulting sequences and feeds the results back into the generation process. The resulting regression tests assert against some observed values from the generated sequences.

Figure 13 shows a simplified test generated by Randoop. Like many automatic test-generation tools, the tests it produces are particularly well-suited to automatic repair. First, many are likely to break on seemingly simple changes. Second, the generated tests tend to have very simple structure and assertions (although they can be quite long).

For our evaluation, we directed Randoop to generate tests for a subset of classes in the subject applications. See [16] for details on our selection criteria and how we executed Randoop.

2) *Results:* Figure 14 shows version difference failures for both manually written and Randoop-generated test suites. ReAssert’s success rate, measured as a ratio of repaired to failing tests, is 45% on manually written tests and 97% (with caveats explained in Section V-D) on tests that we generated using Randoop.

ReAssert’s success rate is more dependent on the structure of the tests than the type of failure. For example, almost

Manually Written			
	Tests	Failures	Repairs
Checkstyle	143	34	9 (26%)
JDepend	53	6	6 (100%)
JFreeChart	1696	18	16 (89%)
Lucene	663	47	12 (25%)
PMD	448	5	5 (100%)
XStream	726	60	28 (47%)

Randoop			
	Tests	Failures	Repairs
JFreeChart	2050	14	9 (64%)
Lucene	1077	278	75+203* (100%)
XStream	692	33	27 (82%)

Figure 14. Tests that fail due to version differences and how many ReAssert repairs.

Strategy	Manual	Randoop
<i>Replace Assertion Method</i>	14	0
<i>Invert Relational Operator</i>	2	0
<i>Replace Literal in Assertion</i>	95	130+203*
<i>Replace with Related Method</i>	0	0
<i>Accessor Expansion</i>	5	0
<i>Trace Declaration-Use Path</i>	11	0
<i>Surround with Try-Catch</i>	12	28

Figure 15. Repair strategies used.

all of the failures in JDepend, JFreeChart, and PMD can be repaired since their tests have simpler control flow than CheckStyle’s and Lucene’s. Similarly, Randoop-generated tests are exceptionally simple, allowing ReAssert to repair nearly all of them.

Figure 15 shows which repair strategies each evaluation scenario used. It is clear that *Replace Literal in Assertion* is a particularly useful strategy despite its simplicity; both scenarios used it heavily. The distribution of other strategies depends on the system under test, the structure of the test code, and the types of failures.

The dependence on test code is particularly apparent in the different strategies used to repair manually written versus Randoop-generated tests. Since Randoop only produces simple `assertTrue`, `assertFalse`, `assertNull`, and `assertEquals` calls, it has little need for complex strategies like *Accessor Expansion* or *Trace Declaration-Use Path*.

D. Unrepairable Failures

ReAssert cannot repair all tests, nor would we expect any tool to be able to do so in a usable manner. Removing all failing assertions would trivially cause a test to pass but would not reveal anything about the correct (or incorrect) behavior of the system under test. Alternately, some failures require large changes to both test and application code. Since ReAssert makes minimal changes to test code and leaves application code unchanged, it is inevitable that it cannot repair all broken tests. The following are common reasons why this is the case.

Test Framework Restrictions: The starred number in figures 14 and 15 represents repairable assertions that cannot be expressed in the test framework. In particular, Randoop produced many `assertNull` invocations that failed when passed a non-null array of primitive types. The *Replace Literal in Assertion* strategy can repair these failures by changing the `assertNull` to `assertArrayEquals`. However, Randoop’s tests require JUnit version 3 while `assertArrayEquals` is implemented in JUnit version 4.

Shared Logic: If a failure occurs in code that is executed multiple times in different contexts—for example if a common setup or teardown method fails, an assertion fails in a loop, or a helper method cannot be repaired by the *Trace Declaration-Use Path* strategy—then the failure is likely unrepairable since a repair may break subsequent executions.

Too Many Assertions: Exceptionally large tests with many failing assertions may exceed the maximum number of failures that ReAssert can repair. By default the limit is 10, but the number can be increased if needed.

Nondeterminism: Tests should generally execute deterministically, but we encountered some that do not, causing unrepairable failures. Similarly, if the *Accessor Expansion* strategy asserts against a nondeterministic accessor, then the test will probably fail. This is a common problem encountered by many test generation tools, including Randoop.

External Resources: Certain tests depend on external resources such as input files. As discussed in Section V-A2, such failures can be repaired but require custom strategies tailored to the applications.

No Applicable Strategies: Certain failures may not meet the preconditions of any repair strategy. For example, a combination of helper methods, complex control flow, and data operations may prevent both *Trace Declaration-Use Path* and *Replace Literal in Assertion* from applying.

Repair Would Remove Test Logic: In certain cases a repair is possible but would remove rather than update test logic. In these cases, we do not implement a repair strategy. For example, a hypothetical Remove Try-Catch strategy would repair several failures in Randoop-generated tests, but would be equivalent to removing an “assertion” that an exception was thrown.

VI. RELATED WORK

While there has been much prior work in program repair and maintenance, we are unaware of any other tools or techniques that automatically repair general-purpose unit tests. Most existing work focuses on repairing the system under test rather than the test code. Those techniques that act on test code often rely on test (re)generation or are applicable only to a particular domain such as GUI test scripts.

A. Fault Localization and Repair

Many researchers have proposed techniques for locating and fixing faults in the system under test. The first step, finding those statements likely relevant to a particular failure, is referred to as *fault localization* [25]–[27]. Our technique—particularly the *Trace Declaration-Use Path* strategy—locates statements to repair in test code by traversing the failure trace; this is conceptually similar to fault localization using backward dynamic program slices [28]–[31].

The second step, repairing the system under test to remove the fault, is sometimes referred to as *automated debugging* [32]. He and Gupta [33] introduce the concept of *path-based weakest preconditions* and uses this formalism to locate and repair erroneous statements. Arcuri [34], Jeffrey et al. [35], and Weimer et al. [36] all model automated debugging as a search problem solvable using genetic algorithms and machine learning. Finally, Abraham and Erwig [37] describe a semi-automatic debugger for spreadsheet formulas.

Other fault repair techniques focus on repairing program state rather than code. Elkarablieh et al. [38] and Demsky et al. [39] both describe techniques that repair faulty data structures while a program executes.

B. Automatic Test Generation

Automatic test generation tools produce a suite of regression tests (sometimes called *characterization tests*) with minimal user input. We use one such tool, Randoop [13], in our evaluation, but there are many other examples utilizing automatic invariant detection [9], [40], symbolic execution [10], software version differences [41], and static program analysis [8].

When automatically generated tests fail due to changes in the system under test, users rarely take the time to examine the failures in detail or repair the tests. Instead it is much more common to throw away the broken tests and re-run the tool to generate new tests. ReAssert is complementary to automatic test generation in that it provides a way to maintain unit tests without re-generation.

ReAssert most frequently repairs unit tests by modifying or creating assertions. Thus, it is closely related to test augmentation techniques that aim to improve regression test oracles [12], [42].

C. Maintaining GUI Test Scripts

Testing graphical user interfaces (GUIs) requires a great deal of manual effort. Test engineers manually write test scripts that interact with GUI widgets or use a record-and-replay testing tool [43]. These scripts or recordings are very fragile, and several researchers have proposed techniques for maintaining GUI tests [44], [45]. In particular, Memon et al. have done extensive work in generating and repairing GUI tests [46]–[48].

Techniques for GUI test repair are conceptually similar to ReAssert’s *Replace Literal in Assertion* strategy: both must replace the incorrect portions of test code with the correct value(s) derived from a test execution. Unlike GUI test maintenance tools, ReAssert repairs general-purpose unit tests and provides an extension mechanism for domain-specific repair strategies.

VII. CONCLUSION

We presented ReAssert, a novel technique and tool that automates repair of broken unit tests. ReAssert performs a combined dynamic and static analysis to find repairs that developers are likely to accept with a mouse click, relieving them from a tedious editing process. Our experiments show that ReAssert can propose repairs for a large percentage of failing tests and that developers often find the suggested repairs useful. We believe that by using ReAssert, developers can write more unit tests or more effectively use automatic test generation tools, improving their ability to detect regressions.

ACKNOWLEDGMENTS

This work was supported in part by the National Science Foundation under Grant No. CCF-0746856. We would like to thank Rob Bocchino and Mohsen Vakilian for their help with DPJ and DPJizer; Steve Lauterburg and Bobak Hadidi for their help with Basset; the participants of our user study for providing an hour of their time; and Milos Gligoric, Munawar Hafiz, Yun Young Lee, and Samira Tasharofi for their valuable comments.

REFERENCES

- [1] G. Venolia, R. DeLine, and T. LaToza, “Software development at Microsoft observed,” Microsoft Research, TR, 2005.
- [2] N. Tillmann and W. Schulte, “Unit tests reloaded: Parameterized unit testing with symbolic execution,” Microsoft Research, TR, 2005.
- [3] K. Beck, *Test-Driven Development By Example*, 2003.
- [4] S. Fraser, D. Astels, K. Beck, B. W. Boehm, J. D. McGregor, J. Newkirk, and C. Poole, “Discipline and practices of TDD,” in *OOPSLA Companion*, 2003.
- [5] B. Ryder and F. Tip, “Change impact analysis for object-oriented programs,” in *PASTE*, 2001.
- [6] X. Ren, F. Shah, F. Tip, B. Ryder, and O. Chesley, “Chianti: a tool for change impact analysis of java programs,” in *OOPSLA*, 2004.
- [7] T. Apiwattanapong, A. Orso, and M. J. Harrold, “Efficient and precise dynamic impact analysis using execute-after sequences,” in *ICSE*, 2005.
- [8] Parasoft, “Jtest.” [Online]. Available: <http://www.parasoft.com/jsp/products/home.jsp?product=Jtest>

- [9] M. Boshernitsan, R.-K. Doong, and A. Savoia, "From Daikon to Agitator: lessons and challenges in building a commercial tool for developer testing," in *ISSTA*, 2006.
- [10] N. Tillmann and J. de Halleux, "Pex-white box test generation for .NET," in *TAP*, 2008.
- [11] C. Csallner and Y. Smaragdakis, "DSD-Crasher: A hybrid analysis tool for bug finding," in *ISSTA*, 2006.
- [12] T. Xie, "Augmenting automatically generated unit-test suites with regression oracle checking," in *ECOOP*, 2006.
- [13] C. Pacheco, S. K. Lahiri, M. D. Ernst, and T. Ball, "Feedback-directed random test generation," in *ICSE*, 2007.
- [14] M. Fowler, *Refactoring: Improving the Design of Existing Code*. Addison-Wesley Professional, 1999.
- [15] J. Purtilo and J. Callahan, "Parse tree annotations," *CACM*, 1989.
- [16] B. Daniel, V. Jagannath, D. Dig, and D. Marinov, "Re-Assert: Suggesting repairs for broken unit tests," University of Illinois at Urbana-Champaign, Tech. Rep. <http://hdl.handle.net/2142/13628>, 2009.
- [17] W. Visser, K. Havelund, G. Brat, S. Park, and F. Lerda, "Model checking programs," *Automated Software Engineering Journal*, vol. 10, no. 2, 2003.
- [18] S. Lauterburg, M. Dotta, D. Marinov, and G. Agha, "A framework for state-space exploration of Java-based actor programs," in *ASE*, 2009.
- [19] R. Bocchino, V. Adve, D. Dig, S. Adve, S. Heumann, R. Komuravelli, J. Overbey, P. Simmons, H. Sung, and M. Vakilian, "A Type and Effect System for Deterministic Parallel Java," in *OOPSLA*, 2009, (To Appear).
- [20] M. Vakilian, D. Dig, R. Bocchino, J. Overbey, V. Adve, and R. Johnson, "Inferring method effect summaries for nested heap regions," in *ASE*, 2009.
- [21] D. Saff and M. D. Ernst, "An experimental evaluation of continuous testing during development," in *ISSTA*, 2004.
- [22] D. Schuler, V. Dallmeier, and A. Zeller, "Efficient mutation testing by checking invariant violations," Universitaet des Saarlandes, TR, 2009.
- [23] B. Daniel and M. Boshernitsan, "Predicting effectiveness of automatic testing tools," in *ASE*, 2008.
- [24] L. Zhang, S.-S. Hou, C. Guo, T. Xie, and H. Mei, "Time-aware test-case prioritization using integer linear programming," in *ISSTA*, 2009.
- [25] S. Hangal and M. S. Lam, "Tracking down software bugs using automatic anomaly detection," in *ICSE*, 2002.
- [26] J. A. Jones, M. J. Harrold, and J. Stasko, "Visualization of test information to assist fault localization," in *ICSE*, 2002.
- [27] H. Cleve and A. Zeller, "Locating causes of program failures," in *ICSE*, 2005.
- [28] F. Tip, "A survey of program slicing techniques," *Journal of Programming Languages*, vol. 3, no. 3, 1995.
- [29] B. Korel and J. W. Laski, "Dynamic slicing of computer programs," *Journal of Systems and Software*, vol. 13, no. 3, 1990.
- [30] H. Agrawal and J. R. Horgan, "Dynamic program slicing," in *PLDI*, 1990.
- [31] M. Weiser, "Program slicing," in *ICSE*, 1981.
- [32] A. Zeller, "Automated debugging: Are we close," *Computer*, vol. 34, no. 11, 2001.
- [33] H. He and N. Gupta, "Automated debugging using path-based weakest preconditions," in *FASE*, 2004.
- [34] A. Arcuri and X. Yao, "A novel co-evolutionary approach to automatic software bug fixing," in *CEC*, 2008.
- [35] D. Jeffrey, M. Feng, N. Gupta, , and R. Gupta, "Bugfix: A learning-based tool to assist developers in fixing bugs," in *ICPC*, 2009.
- [36] W. Weimer, T. V. Nguyen, C. L. Goues, and S. Forrest, "Automatically finding patches using genetic programming," in *ICSE*, 2009.
- [37] R. Abraham and M. Erwig, "Goal-directed debugging of spreadsheets," in *VL/HCC*, 2005.
- [38] B. Elkarablieh, I. Garcia, Y. L. Suen, and S. Khurshid, "Assertion-based repair of complex data structures," in *ASE*, 2007.
- [39] B. Demsky and M. Rinard, "Automatic detection and repair of errors in data structures," in *OOPSLA*, 2003.
- [40] C. Csallner, Y. Smaragdakis, and T. Xie, "DSD-Crasher: A hybrid analysis tool for bug finding," *TOSEM*, 2008.
- [41] S. Person, M. Dwyer, S. Elbaum, and C. Pasareanu, "Differential symbolic execution," in *ESEC/FSE*, 2008.
- [42] Y. Song, S. Thummalapenta, and T. Xie, "UnitPlus: assisting developer testing in Eclipse," in *OOPSLA workshop on Eclipse Technology eXchange*, 2007.
- [43] J. H. Hicinbothom and W. W. Zachary, "A tool for automatically generating transcripts of human-computer interaction," in *HFES*, 1993.
- [44] M. Grechanik, Q. Xie, and C. Fu, "Maintaining and evolving GUI-directed test scripts," in *ICSE*, 2009.
- [45] C. Kaner, "Improving the maintainability of automated test suites," *Software QA*, vol. 4, no. 4, 1997.
- [46] A. Memon, I. Banerjee, N. Hashmi, and A. Nagarajan, "Dart: A framework for regression testing "nightly/daily builds" of GUI applications," in *ICSM*, 2003.
- [47] A. Memon, "Automatically repairing event sequence-based GUI test suites for regression testing," *TSE*, vol. 18, no. 2, 2008.
- [48] A. Memon and M. L. Soffa, "Regression testing of GUIs," in *ESEC/FSE*, 2003.