Using Component Metadata to Support the Regression Testing of Component-Based Software

Mary Jean Harrold[*]
harrold@cc.gatech.edu

Alessandro Orso[†]
orso@cc.gatech.edu

David Rosenblum[‡]
dsr@ics.uci.edu

Gregg Rothermel[§]
grother@cs.orst.edu

Mary Lou Soffa[¶]
soffa@cs.pitt.edu

Hyunsook Do[‖]
dohy@cs.orst.edu

Abstract

Interest in component-based software continues to grow with the recognition of its potential in managing the increasing complexity of software systems. However, the use of externally-provided components has serious drawbacks, in most cases due to the lack of information about the components, for a wide range of activities in the engineering of component-based applications. Consider the activity of regression testing, whose high cost has been, and continues to be, a problem. In the case of component-based applications, regression testing can be even more expensive. When a new version of one or more components is integrated into an application, the lack of information about such externally-developed components makes it difficult to effectively determine the test cases that should be rerun on the resulting application. In previous work, we proposed the use of metadata, which are additional data provided with a component, to support software engineering tasks. In this paper, we present two new metadata-based techniques that address the problem of regression test selection for component-based applications: a code-based approach and a specification-based approach. First, using an example, we illustrate the two techniques. Then, we present a case study that applies the code-based technique to a real component-based system. The results of the study indicate that, on average, 26% of the overall testing effort can be saved over seven releases of the component-based system studied, with a maximum savings of 99% of the testing effort for one version. This reduction demonstrates that metadata can produce benefits in regression testing by reducing the costs related to this activity.

[*]College of Computing, Georgia Institute of Technology
[†]College of Computing, Georgia Institute of Technology
[‡]Information and Computer Science, University of California, Irvine
[§]Department of Computer Science, Oregon State University
[¶]Department of Computer Science, University of Pittsburgh
[‖]Department of Computer Science, Oregon State University

1 Introduction

Interest in component-based software continues to grow with the recognition of its potential in managing the growing complexity of software systems [15, 22, 24]. With increasing frequency, software engineers are building systems by integrating externally-developed software components with application-specific code [4, 15, 24]. Although the use of components provides many advantages, serious drawbacks in that use are becoming apparent. These drawbacks impinge on a wide range of software engineering activities. For example, component usage threatens our ability to validate software [26] and assess reliability [4], complicates maintenance [24], causes problems in program understanding [9], and introduces threats to security [14].

In many cases, the drawbacks of component-based software technologies arise because of the lack of information about externally provided components. Providing the component's source code is not a viable solution to this problem because of intellectual property issues: Component developers often will be unwilling to provide the required information unless they can do so without revealing source code or other sensitive details. Furthermore, not all the required information can be (easily and efficiently) derived from the code alone: Some information, such as data dependences or complexity metrics, may be expensive to compute; other information, such as documentation or changes with respect to previous versions, must be provided in addition to the code.

Existing component standards, including DCOM [5] and Enterprise JavaBeans [10], already supply some additional information about a component through the use of metadata that are packaged with the component. In general, metadata can be *data about components* and *metamethods* for calculating or retrieving those data. The metadata available so far are typically limited to information that is useful for compile-time and run-time type-checking (e.g., the name of the component's class, the names of its functions, and the types of the functions' parameters), and for design-time customization (e.g., the shape or color of a graphical

user interface component, and the maximum size of the internal buffer of a data storage component). Researchers have proposed extending the use of such information for specific tasks [6, 12, 18, 23, 25, 28], but the varieties of metadata currently supported address only a limited range of software engineering problems, such as providing deployment descriptions of components [6, 18] or enhancing self-documentation [28]. In particular, none of the metadata proposed has addressed the important software engineering task of regression testing of component-based software.

In previous work, we introduced a general framework for producing and consuming metadata whose goals are (1) to support the broad range of software engineering tasks that depend on and can benefit from information about external components, and (2) to accommodate component providers' intellectual-property concerns [16]. This paper explores the application of our metadata framework to the problem of regression test selection for component-based software. That is, given (1) an application that uses a set of externally-developed components and has been tested with a test suite, and (2) a new version of this set of which test cases from the test suite must be rerun to exercise the changes in the external components? Test cases in the suite that are not impacted by the changes need not be rerun.

To determine the usefulness of metadata for regression test selection, there are a number of problems that must be resolved. First, we must identify the metadata needed to perform the chosen task. To be able to perform regression test selection, we need information about the coverage achieved by the test suite on the original version of the software. We also need information about the changes made to the set of components. Second, we must determine how to adapt existing regression test selection techniques to incorporate the use of metadata. Different types of techniques must be considered to demonstrate the general applicability of the approach. Finally, we must demonstrate the value of using metadata for regression test selection on component-based software. That is, we have to study how much we can save, in terms of the number of test cases that must be rerun, when using the regression test selection technique proposed.

In this paper, we present metadata and techniques that use metadata for regression test selection on component-based software for two different types of approaches: code-based regression test selection, based on statement-level and method-level regression test selection algorithms [3, 8, 20], and specification-based regression test selection, based on the category-partition method [17]. For both types of approaches, we (1) identify the metadata necessary to determine the test cases to rerun and (2) present techniques for using the metadata for regression test selection. We also describe the results of a study, performed on a real component-based system, that compares the costs of regression test selection with and without metadata. The empirical results

demonstrate that there can be significant savings in the number of test cases that must be rerun for regression testing when component metadata are available, and thus indicate the usefulness of metadata for regression testing.

The main contributions of the paper are the following:

1. identification and use of metadata for regression test selection using a code-based technique;
2. identification and use of metadata for regression test selection using a specification-based technique; and
3. demonstration of the usefulness of metadata in regression test selection for a real program.

In the next section, we present the use of metadata for regression test selection for the two testing techniques. In Section 3, we describe our empirical study and present our results. Section 4 presents a summary and conclusion of our work, and discusses some future work.

2 Component Metadata for Regression Test Selection

This section presents two approaches for providing and using component metadata for regression test selection: code-based and specification-based. The approaches address the general case of an application A that uses a set C of components. To illustrate these two approaches, we use an example consisting of a component and an application that uses it. (For space reasons, the example is limited to a single component.) The component, Dispenser, and the application, VendingMachine, are presented in Figure 1; note that whereas we show the source code of Dispenser in the figure, for the sake of this example we assume that the source code of Dispenser is unavailable to the developer of VendingMachine.

The application models a typical vending machine that is able to dispense specific items to a user. In particular, a user can (1) insert coins into the machine, (2) ask the machine to return the coins inserted and not consumed, and (3) ask the machine to vend a specific item. If the requested item is not available, if the credit is insufficient, or if the selection is invalid, the machine prints an error message and does not dispense the item.

We developed a test suite for VendingMachine (shown in Table 1). Each test case is a sequence of method calls.[1] The test cases are grouped into three sets (1–16, 16–20, 21–25) based on the value of parameter selection that is passed to method VendingMachine.vend. The table indicates whether the test case passed or failed. Test cases 4 and 14 failed because of an error in method Dispense.dispense: If an available item is selected and the credit is insufficient, but greater than zero,

---

[1] For the sake of brevity, in the call sequences we do not show the initial call to the constructor of class VendingMachine, which is implicitly invoked when the class is instantiated.

```
1.  public class VendingMachine {
2.
3.   final private int COIN = 25;
4.   final private int VALUE = 50;
5.   private int totValue;
6.   private int currValue;
7.   private Dispenser d;
8.
9.   public VendingMachine() {
10.    totValue = 0;
11.    currValue = 0;
12.    d = new Dispenser();
13.  }
14.
15.  public void insert() {
16.    currValue += COIN;
17.    System.out.println("Current value = " + currValue);
18.  }
19.
20.  public void return() {
21.    if ( currValue == 0 )
22.      System.err.println( "no coins to return" );
23.    else {
24.      System.out.println("Take your coins");
25.      currValue = 0;}
26.  }
27.
28.  public void vend( int selection ) {
29.    int expense;
30.    expense = d.dispense( currValue, selection );
31.    totValue += expense;
32.    currValue -= expense;
33.    System.out.println( "Current value = " + currValue );
34.  }
35. } // class VendingMachine
36.
37. public class Dispenser {
38.
39.  final private int MAXSEL = 20;
40.  final private int VAL = 50;
41.  private int[] availSelectionVals = {2,3,13};
42.
43.  public int dispense( int credit, int sel ) {
44.    int val=0;
45.    if ( credit == 0 )
46.      System.err.println("No coins inserted" );
47.    else if ( sel > MAXSEL )
48.      System.err.println("Wrong selection "+sel);
49.    else if ( !available( sel ) )
50.      System.err.println("Selection "+sel+" unavailable");
51.    else {
52.      val = VAL;
53.      if ( credit < val )
54.        System.err.println("Enter "+(val-credit)+" coins");
55.      else
56.        System.err.println("Take selection"); }
57.    return val;
58.  }
59.
60.  private boolean available( int sel ) {
61.    for (int i = 0; i<availSelectionVals.length; i++)
62.      if (availSelectionVals[i] == sel) return true;
63.    return false;
64.  }
65. } // class Dispenser
```

Figure 1. Application VendingMachine and component Dispenser.

Table 1. Test suite used to test VendingMachine.

Test Case #	Test Case	Result
Value passed to vend: 3 (i.e., valid selection, available item)		
1	return	Passed
2	vend	Passed
3	insert, return	Passed
4	insert, vend	Failed
5	insert, insert, return	Passed
6	insert, insert, vend	Passed
7	insert, insert, insert, return	Passed
8	insert, insert, insert, vend	Passed
9	insert, insert, insert, insert, return	Passed
10	insert, insert, insert, insert, vend	Passed
11	insert, insert, return, vend	Passed
12	insert, insert, vend, vend	Passed
13	insert, insert, insert, return, vend	Passed
14	insert, insert, insert, vend, vend	Failed
15	insert, insert, insert, insert, return, vend	Passed
16	insert, insert, insert, insert, vend, vend	Passed
Value passed to vend: 9 (i.e., valid selection, unavailable item)		
17	vend	Passed
18	insert, vend	Passed
19	insert, return, vend	Passed
20	insert, vend, vend	Passed
Value passed to vend: 35 (i.e., invalid selection)		
21	vend	Passed
22	insert, vend	Passed
23	insert, insert, vend	Passed
24	insert, insert, insert, vend	Passed
25	insert, insert, insert, insert, vend	Passed

then variable value (set to VAL at line 52) is not reset to zero; consequently, when control returns from Dispenser.dispense to VendingMachine.vend, currValue is erroneously decremented.

Suppose that the component developer finds and fixes this error in Dispenser by adding statement "val = 0;" after statement 54, and releases a new version Dispenser ' of the component. When we integrate Dispenser ' into VendingMachine, we want to regression test the resulting application. For efficiency, we want to rerun only those test cases in our test suite that exercise modifications from Dispenser to Dispenser '. However, without information about the modifications to Dispenser and how they relate to our test suite, we will be forced to run all or most of the test cases in the test suite. At a minimum, we will need to select all test cases in the test suite that exercise the component (20 of the 25 test cases).

2.1 Metadata for Code-based Selection

In this section, we illustrate a metadata-based technique for regression test selection defined for code-based approaches. Code-based testing techniques select test cases based on a coverage goal expressed in terms of some aspect of the code. There are many entities that can be selected for coverage, such as statements, branches, decisions, paths, methods, or classes [1]. Such coverage is usually used as an adequacy criterion for a test suite: the higher the coverage achieved, the higher the adequacy of the test suite.

In particular, for branch coverage techniques, the program is instrumented so that, when it executes, it records the branches traversed by each test case in T, the test suite. With this information, it is possible to associate a subset of

Table 2. Branch coverage for VendingMachine.

Test Case #	Branches Covered
1	(9,10), (20,21), (21,22)
2	(9,10), (28,29)
3	(9.10), (15,16), (20,21), (21,23)
4	(9,10), (15,16), (28,29)
5	(9,10), (15,16), (20,21), (21,23)
6	(9,10), (15,16), (28,29)
7	(9,10), (15,16), (20,21), (21,23)
8	(9,10), (15,16), (28,29)
9	(9,10), (15,16), (20,21), (21,23)
10	(9,10), (15,16), (28,29)
11	(9,10), (15,16), (20,21), (21,23), (28,29)
12	(9,10), (15,16), (28,29)
13	(9,10), (15,16), (20,21), (21,23), (28,29)
14	(9,10), (15,16), (28,29)
15	(9,10), (15,16), (20,21), (21,23), (28,29)
16	(9,10), (15,16), (28,29)
17	(9,10), (28,29)
18	(9,10), (15,16), (28,29)
19	(9,10), (15,16), (20,21), (21,23), (28,29)
20	(9,10), (15,16), (28,29)
21	(9,10), (28,29)
22	(9,10), (15,16), (28,29)
23	(9,10), (15,16), (28,29)
24	(9,10), (15,16), (28,29)
25	(9,10), (15,16), (28,29)

the branches in P, the program under test, with each test case in T. The branches of interest in VendingMachine are (1) entries into each method (i.e., branching because of a method call) and (2) branches from decision statements: Branches (9,10), (15,16), (20,21), and (28,29) represent the former; branches (21,22) and (21,23) represent the latter. Because there are no decision statements in constructor VendingMachine, branch (9,10) represents the call to Dispenser in statement 12, and because there are no decision statements in method vend, branch (28,29) represents the call to dispense in statement 30. Table 2 shows the branches in VendingMachine exercised by each of our test cases.

Code-based regression test selection techniques (e.g., [8, 11, 19, 21, 27]) construct some representation, such as a control-flow graph, a call graph, or a class-hierarchy graph, for a program P and record the coverage achieved by the original test suite T with respect to some entities in that representation. When a modified version P' of P becomes available, these techniques construct the same representation for P' that they constructed for P. The algorithms then use the representations for P and P' and compare them to select the test cases from T for use in testing P', based on (1) differences between the representation for P and P', with respect to the entities considered, and (2) information about which test cases cover the modified entities.[2]

We will use DejaVu as a representative of these code-based approaches [19]. In particular, the DejaVu approach uses a control-flow graph as the representation, and the entities are the branches in the graph. To select test cases to be rerun, DejaVu performs a synchronous traversal of the control-flow graph (CFG) for P and the control-flow graph (CFG') for P', identifies branches modified from CFG to CFG', and selects the test cases that cover such branches as the test cases to be rerun.

For example, to perform regression test selection on application VendingMachine when component Dispenser is changed to Dispenser', DejaVu constructs a control-flow graph CFG' for VendingMachine'. However, because the code for Dispenser is unavailable to the developer of VendingMachine, DejaVu cannot construct control-flow graphs for any of the methods in Dispenser. Therefore, DejaVu can only select test cases based on the analysis of CFG and CFG' for VendingMachine by conservatively considering each branch that represents a call to component Dispenser to be modified. In this case, when DejaVu performs its synchronous traversal of CFG and CFG', it finds that branch (28,29) is affected by the change, and the algorithm selects all test cases that exercise this branch—{2,4,6,8,10–25}.

To achieve better regression test selection when the source code of the component is unavailable, we can use component metadata. To support test selection for code-based regression testing, we need three types of metadata for each component. First, we need to know the branch coverage achieved by the test suite with respect to the component so that we can associate test cases with branches. Second, we need to know the component version. Third, we need a way to query the component for the branches affected by changes in the component between two given versions. The component developer can provide this information in the form of metadata and metamethods, and package them with the component.

We could then construct, for example, a metadata-aware version DejaVu$_{MA}$ of DejaVu. This tool would build the matrix "test cases"–"branches covered" by gathering the component coverage data for each test case. According to the framework presented in Reference [16], a possible interaction of DejaVu$_{MA}$ with component c for incrementally populating the matrix "test cases"–"branches covered" could consist of the following steps:[3]

1. Get the list of types of coverage metadata provided by the component:
   List lmd = c.getMetadata("analysis/dynamic/coverage")
2. Check whether lmd contains the metadata needed (i.e., "analysis/dynamic/coverage/branch"); assume that it does.

---

[2]Regression test selection selects test cases from the original test suite for use in testing the modified program P'. However, modified or new code from P to P' may not be exercised by test cases in T. In this case, the test suite must be augmented by developing new test cases that cover these unexercised parts of the program.

[3]Note that this scenario assumes the existence of some hierarchical scheme for naming and accessing available metadata, as described in [16].

3. Get information on how to access the metadata through metamethods:
MetadataUsage mu = c.getMetadataUsage("analysis/dynamic/coverage/branch")

4. Based on information in mu, fetch the coverage metadata by first enabling the built-in coverage facilities:
c.enableCoverage("analysis/dynamic/coverage/branch")

5. At this point, the built-in coverage facilities provided with component c are enabled, so we can start producing coverage information; for each test case t in the test suite:

- Reset the built-in coverage to get the coverage for t:
c.resetCoverage("analysis/dynamic/coverage/branch")
- Run test case t
- Get the coverage for t:
Metadatum md = getCoverage("analysis/dynamic/coverage/branch").

Now, when $Dispenser$ ′ is acquired, DejaVu$_{MA}$ (1) retrieves from $Dispenser$ its version, (2) using this information queries $Dispenser$ ′ about which branches are affected by the changes between it and $Dispenser$, and (3) selects the test cases to rerun, based on the affected branches and the matrix. In this case, the differences between $Dispenser$ and $Dispenser$ ′ affect only branch (53,54), which is exercised only by test cases 4 and 14. Therefore, only test cases 4 and 14 are selected to be rerun, which is a substantial savings over the approach that does not use metadata.

The technique for code-based regression testing that we have just illustrated is defined at the branch level. When the size of the code increases, the statement-level approach may become impractical. However, the technique can be defined at different levels of granularity. In particular, possible alternatives are to define the technique at the method level, at the class level, or at the subsystem level. In such cases, both the coverage and change information provided through metamethods would be defined at the method, class, or subsystem levels, respectively, rather than at the branch level. In our experiments, we used the method-level approach, as described in Section 3.

2.2 Metadata for Specification-based Selection

In this section, we illustrate a second metadata-based technique for regression test selection, defined for a specification-based approach. Specification-based testing techniques develop test cases based on a functional description of the system. One such technique, the *category-partition method* [17], produces *test frames* that represent a test specification for the functional units in the system. The method is composed of several phases. In the first phase, the tester analyzes the specification to identify the individual functional units in the system; for each unit, the tester identifies parameters (inputs to the unit) and environment factors (elements outside of the code that affect the behavior of the unit). In the second phase, the tester partitions each

functionality **dispense**

- Params:

credit
- zero   [if Available]
- insufficient   [if Available]
- sufficient   [if Available]
- over   [if Available]

selection
- correct   [property Correct]
- incorrect   [error]

- Environment

availability
- available   [if Correct] [property Available]
- unavailable [if Correct] [error]

Figure 2. A possible set of categories, choices, and constraints for the component $Dispenser$.

parameter and environment entity into mutually exclusive choices. In the third phase, the tester identifies constraints among choices, based on their mutual interactions. Finally, in the fourth phase, the tester develops a set of test frames for each unit by computing the cross-product of the different choices; in this phase, the constraints among choices are used to eliminate meaningless or contradictory combinations and to reduce the number of frames, possibly through re-iteration of the third phase.

Analogous to code-based regression test selection techniques, specification-based techniques record the coverage of the original test suite T with respect to entities in the functional representation. In the case of the category-partition method, we can consider the test frames as the entities to be covered. A test case in T covers a test frame *tf* if (1) the parameters of calls to single functionalities match the corresponding choice in *tf*, and (2) the state of the component matches the environment characteristics in *tf*. To compute the coverage of the component achieved by a given test case in terms of test frames, the code must be instrumented according to the identified test frames. In this way, for each test case in T we can identify the test frames that it covers. Therefore, we are able to associate a subset of T with each test frame. This information can be used when performing regression testing of component P'. If we know which frames are affected by the changes, then we can rerun only the test cases associated with such frames. Each test frame identifies a family of test cases that satisfy it. Such test cases, in turn, identify a family of paths within the component—the paths traversed by the execution of the test cases. These paths can therefore be associated with the frame. We say that a change *affects* a test frame *tf* if at least one of the paths associated with *tf* traverses a statement either changed or eliminated in the new version of the component. The component developer can, based both on analysis of the component and on his/her knowledge, identify which frames are affected by the changes between two versions of a given component.

functionality **dispense**

  1 selection: incorrect, availability: X, credit: X
  2 selection: correct, availability: unavailable, credit: X
  3 selection: correct, availability: available, credit: zero
  4 selection: correct, availability: available, credit: insufficient
  5 selection: correct, availability: available, credit: sufficient
  6 selection: correct, availability: available, credit: over

Figure 3. Test frames for component Dispenser (value "X" indicates "don't care" values).

Figure 2 illustrates, for method dispense, a possible set of categories, choices, and constraints on the choices derived by applying the category-partition method to the component Dispenser. Figure 3 shows a set of test frames derived from the test specifications in Figure 2.

For the specification-based approach, to perform regression test selection when Dispenser is modified, we need (1) to know the test frames for the component, (2) to have a way of computing which test cases for application VendingMachine cover which test frames of component Dispenser, and (3) to have information about the test frames affected by the changes in the component. In the case in which a specification for the component is available, we can define test frames for the component. However, because the code for Dispenser is unavailable, we have no way of computing the coverage information (because we need access to the state of the component to check which environmental conditions are satisfied by each test) and no way of identifying which test frames are affected by the changes in Dispenser.

Therefore, to support test selection for specification-based regression testing we need three types of metadata. First, we need to know the coverage achieved by the test suite with respect to the test frames for the component, so that we can associate test cases with frames. Second, we need to know the component version. Third, we need a way to query the component about the test frames affected by the changes in the component between two versions. Again, the component developer will provide this information in the form of metadata and metamethods, packaged with the component.

We could now construct a metadata-aware tool, analogous to the $\text{DejaVu}_{\text{MA}}$ tool of Section 2.1, that would build the matrix "test cases"–"test frames covered" by enabling the coverage computation and gathering the data for each test case. Like $\text{DejaVu}_{\text{MA}}$, this tool would be based on the framework presented in Reference [16], and could consists of the following steps:

1. Get the list of types of coverage metadata provided by the component:
   List lmd = c.getMetadata("analysis/dynamic/coverage")
2. Check whether lmd contains the metadata needed (i.e., "analysis/dynamic/coverage/testframes"); assume that it does.

Table 3. Test frames for the Dispenser component covered by the test cases for the vending machine.

Test Case #	Test Frames Covered	Test Case #	Test Frames Covered
1		14	4, 6
2	3	15	3
3		16	3, 6
4	4	17	2
5		18	2
6	5	19	2
7		20	2
8	6	21	1
9		22	1
10	6	23	1
11	3	24	1
12	3, 5	25	1
13	3		

3. Get information on how to access the metadata through metamethods:
   MetadataUsage mu = c.getMetadataUsage("analysis/dynamic/coverage/testframes")
4. Based on information in mu, fetch the coverage metadata by first enabling the built-in coverage facilities:
   c.enableCoverage("analysis/dynamic/coverage/testframes")
5. At this point, the built-in coverage facilities provided with component c are enabled, so we can start producing coverage information; for each test case t in the test suite:

   - Reset the built-in coverage facilities to get the coverage for t:
     c.resetCoverage("analysis/dynamic/coverage/testframes")
   - Run test case t
   - Get the coverage for t:
     Metadatum md = getCoverage("analysis/dynamic/coverage/testframes").

When a new version of the component is acquired, the tool (1) gathers the metadatum about the version from the old component, (2) using this information queries the new component for the test frames affected by the changes between its version and the version currently in the system, and (3) selects the test cases to rerun, based on the affected frames and the matrix.

Suppose we apply this technique to the VendingMachine example. First, we run the 25 test cases for the application and gather the test-frame coverage information; in Table 3, we show how the different test frames in Figure 3 are covered by the test cases for the vending machine. Second, when we acquire Dispenser ', we check which test frames are affected by the changes between the two versions of the component; we discover that only test frame 4 is affected. Finally, we use the matrix "test cases"–"test frames covered" to select the test cases to be rerun; according to the information in Figure 3, we select test cases 4 and 14.

As with the code-based approach, the specification-

based approach provides a meaningful reduction in the number of test cases to be rerun for the new version of VendingMachine.

3 Case Study

To investigate whether the use of metadata can benefit regression testing of applications built with external components, we performed a case study. Specifically, we investigated the following research question:

> Let A be a program created by an application developer using a set of externally-developed components C. Let T be a test suite created to test A. Suppose a new version C' of C is created through modifications to one or more of the components in C, and suppose the developer of A then wishes to adopt C' for use in A. If metadata are available with C and C', can the developer reuse T to regression test A more efficiently than if metadata are not available?

In this study, we restricted our attention to the use of metadata for code-based regression test selection techniques, as described in Section 2, and we focused on two specific regression test selection techniques:

No metadata. The developer of A knows only that one or more of the components in C have been modified, but not which. Therefore, to selectively retest A safely, the developer must rerun any test case in T that exercises code in one or more of the components in C. We will refer to this as the NOMETA technique.

Metadata for method-level regression test selection.
The developer of A possesses metadata provided by the developer of C, sufficient to support selection of test cases that exercise methods changed in producing C' from C using the procedure described in Section 2. We will refer to this as the META technique.

3.1 Measures

Regression test selection techniques achieve savings by reducing the effort required to regression test a modified program. Thus, one method used to compare such techniques [3] is to measure and compare the degrees to which the techniques reduce test suite size for given modified versions of a program. We adopt this approach. For each regression test selection technique R that we consider, and for each (version, subsequent-version) pair (P_i, P_{i+1}) of program P, where P_i is tested by test suite T, we measure the percentage of T selected by R to test P_{i+1}.

3.2 Study Subject

As a subject for our study we used several versions of the Java implementation of the SIENA server [7]. SIENA (Scalable Internet Event Notification Architecture) is an Internet-scale event notification middleware for distributed event-based applications deployed over wide-area networks, responsible for selecting notifications that are of interest to clients (as expressed in client subscriptions) and then delivering those notifications to the clients via access points.

To investigate the effects of using component metadata for regression test selection, we required an application program, constructed using external components that could be provided with metadata. SIENA is logically divided into a set of six components (consisting of nine classes of about 1.5KLOC), which constitute "a set of external components C," and a set of 17 other classes of about 2KLOC, which constitute an application that could be constructed using C.

We obtained the source code for all the different versions of SIENA, from its first to its last release (about 15 different releases), in the form of an RCS repository. We extracted, from the SIENA repository, eight different sequentially released versions of C (versions 1.8 through 1.15), which we refer to as $C_1, C_2, \ldots, C_8$, respectively. Each version provides enhanced functionality or bug fixes over the previous version. The net effect of this act was the provision of eight successive versions of SIENA, $A_1, A_2, \ldots, A_8$, constructed using $C_1, C_2, \ldots, C_8$, respectively. These versions of SIENA represent a succession of versions, each of which the developer of A would want to retest. The pairs of versions (A_k, A_{k+1}), $1 \leq k \leq 7$, formed the (version, modified-version) pairs for our study.

To investigate the impact of metadata on regression test selection we also required a test suite for our base version A_1 of SIENA that could be reused in retesting subsequent versions. Such a test suite did not already exist for the SIENA release we considered, so we created one. To provide test cases in an unbiased manner, one of the authors of this paper, who is involved in defining the requirements and design of SIENA but is unfamiliar with its implementation details, independently created a black-box test suite, based on the functionality of SIENA, that consists of 138 test cases. This set of test cases served as the subject regression test suite for our study.

3.3 Procedure

Because the creation of metamethods and support tools for directly implementing our target techniques would be expensive, our goal was to discover a way to address our research question without creating such infrastructure. We designed a procedure by which we could determine precisely, for a given test suite and (program, modified-version) pair, which test cases would be selected by our two target techniques. For each (program, modified-version) pair

(P_i, P_{i+1}), we used the Unix diff utility and inspection of the code to locate differences between P_i and P_{i+1}, including modified, new, and deleted code. In cases where variable or type declarations differed, we found the methods in which those variables or types were used, and treated those methods as if they had been modified. We used this information to determine the methods in P_i that would be reported changed for the META technique. We instrumented each such method so that, when executed, the method outputs the text "selected", and we then constructed an executable of the application from this instrumented code.

Given this procedure, to determine which test cases in T would be selected by the META technique for (P_i, P_{i+1}) it was sufficient to execute all test cases in T on our instrumented version of P_i, and record which test cases caused P_i to output (one or more times) the text "selected". By construction, these are exactly the test cases that would be selected by an implementation of the META technique.

Determining the test cases that would be selected by the NOMETA technique required a similar, but simpler approach. We instrumented the application developer's portion of the code for P, inserting code that outputs "selected" prior to any invocation of any method in C, and then executed the test cases in T on that instrumented version.

The foregoing procedures require us to execute all test cases in T to determine which test cases would be selected by an actual regression test selection tool; thus, the approaches are of use only for experimentation. However, the approaches let us determine exactly the test cases that would be selected by the techniques.

We applied this approach to each of the seven (program, modified-version) pairs of the SIENA system with our given test suite, and recorded, for each of the two regression test selection techniques, the percentage of the test suite selected by that technique for that (program, modified-version) pair. These percentages served as the data set for our analysis.

3.4 Results and Discussion

Figure 4 contains a graph depicting the test selection results obtained in this study. In the graph, each modified version of SIENA occupies a position along the x-axis, and the test selection data for that version are represented by a vertical bar, black for the NOMETA technique and grey for the META technique. The height of the bar depicts the percentage of tests selected by the technique on that version.

As the figure shows, the NOMETA technique always selected 97% of the test cases. Only 3% of the test cases do not exercise components in C (the set of external components), and thus all others must be re-executed. Also, since the NOMETA technique selects all test cases that executed any components in C, and the test cases in our test suite that encounter C do not vary across versions, the NOMETA technique selected the same test cases for each version.
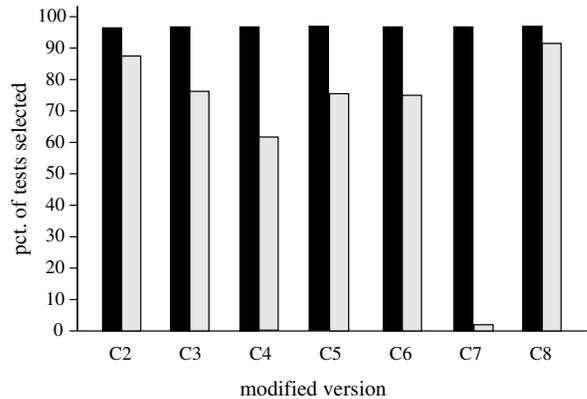

Figure 4. Test selection results for the NOMETA (black) and META (grey) techniques.

As the figure also shows, the META technique always selected a smaller subset of the test suite than the NOMETA technique. In the case of version C7, the difference was extreme: the META technique selected only 1.5% of the test cases in the test suite, whereas the NOMETA technique selected 97% of the test cases. This large difference arose because the changes within C7 are minor, involving few methods, and methods encountered by only a few test cases. On the other versions, differences in selection were more modest, ranging from 6% to 37% of the test suite.

On versions C3, C5 and C6, the META technique selected identical test cases, even though the code changes in those versions differed. This occurred because the code changes involved the same sets of methods. Theoretically, if we had used statement-level regression test selection, the test cases selected for these versions could have differed.

The fact that a regression test selection technique reduces the number of test cases that must be run does not guarantee that the technique will be cost-effective. That is, even though we reduce the number of test cases that need to be rerun, if this does not produce savings in testing time, the reduction in number of test cases will be meaningless. Moreover, savings in testing time might not be proportional to savings in number of test cases (if, for example, the test cases excluded are all inexpensive, while those not excluded are expensive). (See [13] for an applicable cost model.)

In the absence of implementations of the testing techniques and measurements of analysis costs, we cannot determine such savings precisely for this case study; however, we can still gain some insights by considering test execution times. Thus, we recorded execution times for the test cases selected by each technique on each version. Table 4 shows these times.

For each version, the table shows the minutes and seconds required to test that version. The columns show the

Version	NOMETA Test Execution Time	META Test Execution Time
C2	19:44	18:45
C3	19:51	16:57
C4	19:51	13:07
C5	19:52	17:44
C6	19:52	16:40
C7	19:51	00:15
C8	19:49	19:26
average	19:50	14:07
total	138:50	102:54

Table 4. Execution times (minutes:seconds) for test cases selected by the NOMETA and META techniques.

version number, the time required to run the test cases selected by the NOMETA technique, and the time required to run the test cases selected by the META technique. The last two rows show average and total times. On average over the seven modified versions, the META technique reduced testing time from 19 minutes and 50 seconds to 14 minutes and 7 seconds. The total time savings over the sequence of seven versions was 35 minutes and 59 seconds (26% of total time.) In the worst case, for version C8, the META technique saved only 23 seconds (2%) of testing time. In the best case, for version C7, it saved 19 minutes and 36 seconds (99%) of testing time.

Note that these times do not factor in the cost of the analysis required to perform test selection, but in other studies of test selection those costs have been shown to be quite low [20]. Furthermore, these times include only the times required to execute, and not validate test cases; validation would further inflate the times, and increase the savings.

Of course, savings of a few minutes and seconds, such as those exhibited in the differences in testing time seen in this study, may be unimportant. In practice, however, regression testing can require hours, days, or even weeks of effort, and much of this effort may be human-intensive. If results such as those demonstrated by this study scale up, a savings of 26% of the overall testing effort for a sequence of seven releases may be substantial, and a savings of 99% of the testing effort for a version may be a huge win. These results thus provide good evidence that testing with metadata could save significant costs and that metadata could be useful for regression test selection in component-based software.

3.5 Limitations of this Study

Like any empirical study, this study has limitations. We have considered the application of only two regression test selection techniques to a single program and test suite and seven subsequent modified versions of the components that make up that program. Furthermore, we have considered only one measure of test selection effectiveness: percentage reduction in test suite size (although we have buttressed this measure by also considering test execution cost data). Other costs, such as the cost of providing metadata and performing test selection, may be important in practice.

On the other hand, the program and modified versions we used are derived from an actual implementation, and our specification-based test suite represents a test suite that could be used in practice. Furthermore, previous work [3, 20] has illustrated the applicability of our cost measure. Our results thus support an "existence argument": Cases exist in which metadata can produce benefits in regression testing. Thus, these results motivate further research, and the implementation of tools to support the techniques, followed by carefully controlled experimentation, to investigate whether such results will generalize.

4 Conclusion

We have introduced two new techniques for regression testing of component-based applications. The first technique is code-based, and the second technique is specification-based. Both techniques are based on the use of metadata and metamethods to package additional information together with a component. The presence of metadata allows component developers to provide information useful for regression test selection without disclosing the source code of the components they distribute. In particular, only version information, coverage measurement facilities, and information about changes between versions of components need be provided for the techniques to be applicable.

To assess the applicability and effectiveness of the proposed techniques in practice, we have presented a case study performed on a real system. Although there are some limitations to the results of our study, the study does show that cases exist in which the use of metadata can reduce the costs of regression testing component-based applications. In particular, our code-based technique resulted in an average savings of 26% of the testing effort over seven subsequent releases of the considered set of components, with a maximum saving of 99% of the testing effort for one of the versions.

Because of these promising initial results, we plan to perform further research on the use of metadata for regression testing. Our first goal is to build a set of tools that allow us to automate the application of the presented techniques and to integrate them into the ARISTOTLE analysis system [2] and DejaVu [19]. In this way, we will be able to run extensive experiments to further validate the code-based approach. In parallel, we will study the applicability of the specification-based approach on real examples. Finally, we will study other applications of component metadata and their effectiveness for software engineering tasks.

Acknowledgements

References

[1] W. R. Adrion, M. A. Branstad, and J. C. Cherniavsky. Validation, verification, and testing of computer software. *ACM Computing Surveys*, 14(2):159–192, June 1982.

[2] Aristotle Research Group. ARISTOTLE: Software engineering tools. http://www.cc.gatech.edu/aristotle/, 2000.

[3] J. Bible, G. Rothermel, and D. Rosenblum. Coarse- and fine-grained safe regression test selection. *ACM Trans. Softw. Eng. and Meth.*, to appear.

[4] P. Brereton and D. Budgen. Component-Based Systems: A Classification of Issues. *IEEE Computer*, 33(11):54–52, November 2000.

[5] N. Brown and C. Kindel. *Distributed Component Object Model protocol: DCOM/1.0*. January 1998.

[6] C. Canal, L. Fuentes, J. Troya, and A. Vallecillo. Extending CORBA interfaces with p-calculus for protocol compatibility. In *Technology of Object-Oriented Languages and Systems (TOOLS'00)*, pages 208–225, June 2000.

[7] A. Carzaniga, D. Rosenblum, and A. L. Wolf. Achieving scalability and expressiveness in an internet-scale event notification service. In *Nineteenth ACM Symp. Princ. Distr. Computing*, pages 219–227, July 2000.

[8] Y. Chen, D. Rosenblum, and K. Vo. TestTube: A system for selective regression testing. In *16th Int'l. Conf. Softw. Eng.*, pages 211–222, May 1994.

[9] R. Cherinka, C. M. Overstreet, and J. Ricci. Maintaining a COTS integrated solution — Are traditional static analysis techniques sufficient for this new programming methodology? In *Int'l. Conf. Softw. Maint.*, pages 160–169, November 1998.

[10] Enterprise JavaBeans technology. http://java.sun.com/products/ejb/index.html, October 2000.

[11] T. Graves, M. J. Harrold, J.-M. Kim, A. Porter, and G. Rothermel. An empirical study of regression test selection techniques. In *Int'l. Conf. Softw. Eng.*, pages 188–197, April 1998.

[12] F. Huber, A. Rausch, and B. Rumpe. Component interface diagrams: Putting components to work. Institutsbericht, Technische Universitaet Muenchen, Institut fuer Informatik, December 1998.

[13] H. K. N. Leung and L. J. White. A cost model to compare regression test strategies. In *Conf. Softw. Maint.*, pages 201–208, October 1991.

[14] U. Lindquist and E. Jonsson. A map of security risks associated with using COTS. *IEEE Computer*, 31(6):pages 60–66, June 1998.

[15] P. M. Maurer. Components: What if they gave a revolution and nobody came. *IEEE Computer*, 33(6):28–34, June 2000.

[16] A. Orso, M. J. Harrold, and D. S. Rosenblum. Component metadata for software engineering tasks. In W. Emmerich and S. Tai, editors, *EDO '00*, Lecture Notes in Computer Science. Springer-Verlag / ACM Press, November 2000. (to appear).

[17] T. Ostrand and M. Balcer. The category-partition method for specifying and generating functional tests. *Comm. ACM*, 31(6), June 1988.

[18] G. Piccinelli and S. Lynden. Concept and tools for e-service development. In *7th Workshop HP Openview Univesity Association (OVUA'00)*, June 2000.

[19] G. Rothermel and M. Harrold. A safe, efficient regression test selection technique. *ACM Trans. on Softw. Eng.and Meth.*, 6(2):173–210, April 1997.

[20] G. Rothermel and M. Harrold. Empirical studies of a safe regression test selection technique. *IEEE Trans. Softw. Eng.*, 24(6):401–419, June 1998.

[21] G. Rothermel, M. J. Harrold, and J. Dedhia. Regression test selection for C++ software. *J. Softw. Testing, Verif., and Rel.*, 10(2), June 2000.

[22] C. Szyperski. *Component Oriented Programming*. Addison-Wesley, first edition, 1997.

[23] J. Troya and A. Vallecillo. On the addition of properties to components. In J. Bosch and S. Mitchell, editors, *Object-Oriented Technology: ECOOP'97 Workshop Reader*, volume 1357 of *Lecture Notes in Computer Science*, pages 374–378. Springer, 1997.

[24] J. Voas. The challenges of using COTS software in component-based development. *IEEE Computer*, 31(6):44–45, June 1998.

[25] J. Voas. Maintaining component-based systems. *IEEE Software*, 15(4):22–27, July 1998.

[26] E. Weyuker. Testing component-based software: A cautionary tale. *IEEE Software*, 15(5):54–59, Sept–Oct 1998.

[27] L. White and H. Leung. A firewall concept for both control-flow and data-flow in regression integration testing. In *Conf. Softw. Maint.*, pages 262–270, November 1992.

[28] XOTcl - extended object Tcl. http://nestroy.wi-inf.uni-essen.de/xotcl/, November 2000.