

Dynamic Slicing on Java Bytecode Traces

TAO WANG and ABHIK ROYCHOUDHURY

National University of Singapore, Singapore

Dynamic slicing is a well-known technique for program analysis, debugging and understanding. Given a program P and input I , it finds all program statements which directly/indirectly affect the values of some variables' occurrences when P is executed with I . In this paper, we develop a dynamic slicing method for Java programs. Our technique proceeds by backwards traversal of the bytecode trace produced by an input I in a given program P . Since such traces can be huge, we use results from data compression to compactly represent bytecode traces. The major space savings in our method come from the optimized representation of (a) data addresses used as operands by memory reference bytecodes, and (b) instruction addresses used as operands by control transfer bytecodes. We show how dynamic slicing algorithms can directly traverse our compact bytecode traces without resorting to costly decompression. We also extend our dynamic slicing algorithm to perform "relevant slicing". The resultant slices can be used to explain omission errors that is, why some events did not happen during program execution. Detailed experimental results on space/time overheads of tracing and slicing are reported in the paper. The slices computed at the bytecode level are translated back by our tool to the source code level with the help of information available in Java class files. Our *JSlice* dynamic slicing tool has been integrated with the Eclipse platform and is available for usage in research and development.

Categories and Subject Descriptors: D.3.4 [**Programming Languages**]: Processors—*Debuggers*; D.2.5 [**Software Engineering**]: Testing and Debugging—*Debugging aids*; *Testing tools*; *Tracing*

General Terms: Algorithms, Experimentation, Measurement

Additional Key Words and Phrases: Program slicing, Tracing, Debugging

1. INTRODUCTION

Program slicing [Weiser 1984] is a well-known technique for program debugging and understanding. Roughly speaking, program slicing works as follows. Given a program P , the programmer provides a slicing criterion of the form (l, V) , where l is a control location in the program and V is a set of program variables referenced at l . The purpose of slicing is to find out the statements in P which can affect the values of V at l via control and/or data flow. So, if during program execution the values of V at l were "unexpected", the corresponding slice can be inspected to explain the reason for the unexpected values.

Slicing techniques are divided into two categories: static and dynamic. Static (Dynamic) slicing computes the fragment affecting V at l (some occurrence of l) when the input program is executed with any (a specific) input. Thus, for the same slicing criterion, dynamic slice for a given input of a program P is often smaller than the static slice of P . Static slicing techniques typically operate on a program dependence graph (PDG); the nodes of the PDG are simple statements /

A preliminary version of this paper appeared in the International Conference on Software Engineering (ICSE) 2004, see [Wang and Roychoudhury 2004]. *Authors' addresses*: Tao Wang and Abhik Roychoudhury, Department of Computer Science, National University of Singapore, Singapore. E-mail: {wangtao,abhik}@comp.nus.edu.sg

conditions and the edges correspond to data / control dependencies [Horwitz et al. 1990]. Dynamic slicing w.r.t. an input I on the other hand, often proceeds by collecting the execution trace corresponding to I . The dynamic data and control dependencies between the statement occurrences in the execution trace can be pre-computed or computed on demand during slicing [Zhang et al. 2005].

Due to the presence of objects and pointers in programs, static data dependence computations are often conservative, leading to very large static slices. On the other hand, dynamic slices capture the closure of *dynamic* data and control dependencies, hence they are much more precise, and more helpful for narrowing the attention of the programmer. Furthermore, since dynamic slices denote the program fragment affecting the slicing criterion for a *particular* input, they naturally support the task of debugging via running of selected test inputs.

Though dynamic slicing was originally proposed for debugging [Korel and Laski 1988; Agrawal and Horgan 1990], it has subsequently also been used for program comprehension in many other innovative ways. In particular, dynamic slices (or their variants which also involve computing the closure of dependencies by trace traversal) have been used for studying causes of program performance degradation [Zilles and Sohi 2000], identifying isomorphic instructions in terms of their run-time behaviors [Sazeides 2003] and analyzing spurious counter-example traces produced by software model checking [Majumdar and Jhala 2005]. Even in the context of debugging, dynamic slices have been used in unconventional ways e.g. [Akgul et al. 2004] studies reverse execution along a dynamic slice. Thus, dynamic slicing forms the core of many tasks in program development and it is useful to develop efficient methods for computing dynamic slices.

In this paper, we present an infrastructure for dynamic slicing of Java programs. Our method operates on bytecode traces; we work at the bytecode level since slice computation may involve looking inside library methods and the source code of libraries may not always be available. First, the bytecode stream corresponding to an execution trace of a Java program for a given input is collected. The trace collection is done by modifying a virtual machine; we have used the Kaffe Virtual Machine in our experiments. We then perform a backward traversal of the bytecode trace to compute dynamic data and control dependencies on-the-fly. The slice is updated as these dependencies are encountered during trace traversal. Computing the dynamic data dependencies on bytecode traces is complicated due to Java's stack based architecture. The main problem is that partial results of a computation are often stored in the Java Virtual Machine's operand stack. This results in implicit data dependencies between bytecodes involving data transfer via the operand stack. For this reason, our backwards dynamic slicing performs a "reverse" stack simulation while traversing the bytecode trace from the end.

Dynamic slicing methods typically involve traversal of the execution trace. This traversal may be used to pre-compute a dynamic dependence graph or the dynamic dependencies can be computed on demand during trace traversal. Thus, the representation of execution traces is important for dynamic slicing. This is particularly the case for backwards dynamic slicing where the trace is traversed from the end (and hence needs to be stored). In practice, traces tend to be huge. The work of [Zhang et al. 2005] reports experiences in dynamic slicing programs like `gcc` and `perl` where the execution trace runs into *several hundred million instructions*. It

might be inefficient to perform post-mortem analysis over such huge traces. Consequently, it is useful to develop a compact representation for execution traces which capture both control flow and memory reference information. This compact trace should be generated *on-the-fly* during program execution.

Our method proceeds by on-the-fly construction of a compact bytecode trace during program execution. The compactness of our trace representation is owing to several factors. First, bytecodes which do not correspond to memory access (*i.e.* data transfer to and from the heap) or control transfer are not stored in the trace. Operands used by these bytecodes are fixed and can be discovered from Java class files. Secondly, the sequence of addresses used by each memory reference bytecode or control transfer bytecode is stored separately. Since these sequences typically have high repetition of patterns, we exploit such repetition to save space. We modify a well-known lossless data compression algorithm called SEQUITUR [Nevill-Manning and Witten 1997] for this purpose. This algorithm identifies repeated patterns in the sequence on-the-fly and stores them hierarchically.

Generating compact bytecode traces during program execution constitutes the first phase of our dynamic slicer. Furthermore, we want to traverse the compact execution trace to retrieve control and data dependencies for slicing. This traversal should be done without decompressing the trace. In other words, the program trace should be collected, stored and analyzed for slicing – all in its compressed form. This is achieved in our dynamic slicer which traverses the compact bytecode trace and computes the data/control dependencies in compression domain. Since we store the sequence of addresses used by each memory-reference/control-transfer bytecode in compressed format, this involves marking the “visited” part of such an address sequence without decompressing its representation.

Finally, we extend our dynamic slicing method to support “relevant slicing”. Traditional dynamic slicing algorithms explain the values of variables V at some occurrences of a control location l , by highlighting the executed program statements which affect the values of V at l . They do not report the statements which affect the value of V at l , by *not being executed*. To fill this caveat, the concept of relevant slicing was introduced in [Agrawal et al. 1993; Gyimóthy et al. 1999]. We design and implement a relevant slicing algorithm working on our compact bytecode traces.

Summary of Contributions. In summary, the main contribution of this paper is to report methods for dynamic slicing of Java programs. Our slicer proceeds by traversing a compact representation of a bytecode trace and constructs the slice as a set of bytecodes; this slice is then transformed to the source code level with the help of Java class files. Our slicing method is complicated by Java’s stack based architecture which requires us to simulate a stack during trace traversal. We have made our *Jslice* slicing tool available from <http://jslice.sourceforge.net/>. To the best of our knowledge, ours is the first dynamic slicing tool for Java programs.

Since the execution traces are often huge, we develop a space efficient representation of the bytecode stream for a Java program execution. This compressed trace is constructed on-the-fly during program execution. Our dynamic slicer performs backward traversal of this compressed trace *directly* to retrieve data/control dependencies, that is, slicing does not involve costly trace decompression. Our compressed trace representation is interesting in general as a program trace rep-

resentation. The crucial factor leading to compression of program traces is the separation of address sequences used by conditional jump and memory reference instructions. For our subject programs (drawn from standard suites such as the Java Grande benchmark suite or the SPECjvm suite) we obtain compression in varying amounts ranging from 5 – 5000 times. We show that the time overheads for constructing this representation on-the-fly during program execution is tolerable.

We also enhance our dynamic slicing algorithm to capture “omission errors” via “*relevant slicing*” [Agrawal et al. 1993; Gyimóthy et al. 1999]. The relevant slicing algorithm also operates *directly* on the compressed bytecode traces, as our dynamic slicing algorithm. Our experimental results indicate that the additional capability of relevant slices (*i.e.* capturing omission errors) comes at the cost of modest additional overheads in terms of computation time or slice sizes.

Section Organization. The rest of this paper is organized as follows. Section 2 describes our compressed representation of Java bytecode traces. Section 3 and 4 discuss our dynamic slicing and relevant slicing algorithms. Section 5 presents experimental evaluation, while Section 6 describes our JSlice dynamic slicing tool. Planned extensions of the Jslice tool are described in Section 7. Section 8 summarizes related work. We conclude the paper in section 9.

2. COMPRESSED BYTECODE TRACE

In this section, we will discuss how to collect compact bytecode traces of Java programs *on the fly*. This involves a discussion of the compaction scheme as well as the necessary instrumentation. The compaction scheme used by us is exact, lossless and on-the-fly.

2.1 Overall representation

The simplest way to define a program trace is to treat it as a sequence of “instructions”. For Java programs, we view the trace as the sequence of executed bytecodes, instead of program statements. This is because only bytecodes are available for Java libraries, which are used by Java programs. Furthermore, collecting traces at the level of bytecode has the flexibility in tracing/not tracing certain bytecodes. For example, the `getstatic` bytecode loads the value of a static field. This bytecode does not need tracing, because which static field to access is decided at compile-time, and can be discovered from class files during post-mortem analysis.

However, representing a Java program trace as a bytecode sequence has its own share of problems. In particular, it does not allow us to capture many of the repetitions in the trace. Representation of the program trace as a single string loses out structure in several ways.

- The individual methods executed are not separated in the trace representation.
- Sequences of target addresses accessed by individual control transfer bytecodes are not separated out. These sequences capture control flow and exhibit high regularity (*e.g.* a loop branch repeats the same target many times).
- Similarly, sequences of addresses accessed by individual memory load/store bytecodes are not separated out. Again these sequences show fair amount of repetition (*e.g.* a read bytecode sweeping through an array).

In our representation, the compact trace of the whole program consists of trace tables, each of which is used for one method. Method invocations are captured by tracing bytecodes which invoke methods. The last executed bytecode w.r.t. the entire execution is clearly marked. Within the trace table for a method, each row maintains traces of a specific bytecode or of the exit of the method. Monitoring and tracing every bytecode may incur too much time and space overheads. We monitor only the following five kinds of bytecodes to collect the trace, where the first two are necessary to capture data flow of the execution, and the last three are necessary to capture control flow of the execution.

- *Memory allocation bytecodes.* Memory allocation bytecodes record the identities of created objects.

- *Memory access bytecodes.* The bytecodes to access local variables and static fields are not traced since the addresses accessed by these bytecodes can be obtained from the class file. For bytecodes accessing object fields / array elements, we trace the addresses (or identities since an address may be used by different variables in the lifetime of a program execution) corresponding to the bytecode operands.

- *Method invocation bytecodes.* Java programs use four kinds of bytecodes to invoke methods. Two of them, `invokevirtual` and `invokeinterface`, may invoke different methods on different execution instances. These invoked methods have the same method name and parameter descriptor (which can be discovered in class files), but they belong to different classes. So, for every `invokevirtual` and `invokeinterface` bytecode, we record the classes which the invoked methods belong to.

- *Bytecodes with multiple predecessors.* Some bytecodes have multiple predecessors in the control flow graph. For such a bytecode, we record which bytecodes are executed immediately before itself.

- *Method return bytecodes.* If a method has multiple `return` bytecodes, the trace of the method-exit records which `return` bytecodes are executed.

Monitoring the last two kinds of bytecodes (bytecodes with multiple predecessors and method return bytecodes) and marking the last executed bytecode are required due to backward traversal of the trace during post-mortem analysis. On the other hand, if slicing proceeds by forward traversal of the trace, it is not necessary to monitor bytecodes with multiple predecessors and method return bytecodes. Instead, for each conditional branch bytecode we can record which bytecodes are executed immediately after the branch bytecode (*i.e.*, the target addresses).

As mentioned earlier, our trace representation captures each method’s execution in the trace as a trace table. Each row of the trace table for a method m represents the execution of one of the bytecodes of m (in fact it has to be a bytecode which we trace). A row of a trace table thus captures all execution instances of a specific bytecode. The row corresponding to a bytecode b in method m stores the sequence of values taken by each operand of b during execution; if b has multiple predecessors, we also maintain a sequence of the predecessor bytecode of b . Thus, in each row of a trace table we store several sequences in general; *these sequences are stored in a compressed format*. Separating the sequence of values for each bytecode operand allows a compression algorithm to capture and exploit regularity and repetition in

```

1: class Demo{
2:
3:   public int foo(int j){
4:     int ret;
5:     if ( j % 2 == 1 )
6:       ret= 2;
7:     else
8:       ret= 5;
9:     return ret;
10:  }
11:
12:  static public void main (String argsv[]){
13:    int i, k, a, b;
14:    Demo obj= new Demo();
15:    int arr[]= new int[4];
16:
17:    a=2;
18:    b=1;
19:    k=1;
20:    if (a>1){
21:      if (b>1){
22:        k=2;
23:      }
24:    }
25:
26:    for (i=0; i < 4; i++){
27:      arr[i]=k;
28:      k= k + obj.foo(i);
29:    }
30:
31:    System.out.println(k);
32:  }
33: }

public static void main(String[]);
1:   new Class Demo
2:   dup
3:   invokespecial Demo()
4:   astore 5
5:   iconst_4
6:   newarray int
7:   astore 6
8:   iconst_2
9:   istore_3
10:  iconst_1
11:  istore 4
12:  iconst_1
13:  istore_2
14:  iload_3
15:  iconst_1
16:  if_icmple 22
17:  iload 4
18:  iconst_1
19:  if_icmple 22
20:  iconst_2
21:  istore_2
22:  iconst_0
23:  istore_1
24:  iload_1
25:  iconst_4
26:  if_icmpge 39
27:  aload 6
28:  iload_1
29:  iload_2
30:  iastore
31:  iload_2
32:  aload 5
33:  iload_1
34:  invokevirtual foo:(int)
35:  iadd
36:  istore_2
37:  iinc 1, 1
38:  goto 24
39:  getstatic
40:  iload_2
41:  invokevirtual println:(int)
42:  return

Demo();
43:  aload_0
44:  invokespecial Object()
45:  return

public int foo(int);
46:  iload_1
47:  iconst_2
48:  irem
49:  iconst_1
50:  if_icmpne 54
51:  iconst_2
52:  istore_2
53:  goto 56
54:  iconst_5
55:  istore_2
56:  iload_2
57:  ireturn

```

Fig. 1. The left part is a simple Java program, and the right part shows corresponding bytecodes. Method `Demo()` is generated automatically as the class constructor.

the values taken by an operand. This can be due to regularity of control or data flow (*e.g.*, a read bytecode sweeping through an array or a loop iterating many times). Before presenting how to compress trace sequences, let us look at an example to understand the trace table representation. Note that sequences are not compressed in this example for ease of understanding.

Example. The left part of Figure 1 presents a simple Java program, and the right part shows the corresponding bytecode stream. Table I shows the trace tables for methods `main` and `foo`, respectively. The constructor method `Demo` has no trace table, because no bytecode of this method is traced. Each row in the trace table consists of: (a) the id/address for a bytecode (in the *Bytecode* column), and (b) collected traces for that bytecode (in the *Sequences* column).

For our example Java program, there are 57 bytecodes altogether, and only 8 of them are traced, as shown in Table I. Bytecodes 1 and 6 (*i.e.* two `new` statements at lines 14 and 15 of the source program) allocate memory for objects, and their traces include o_1 and o_2 , which represent identities of the objects allocated by these

Bytecode	Sequences
1	$\langle o_1 \rangle$
6	$\langle o_2 \rangle$
22	$\langle 19 \rangle$
24	$\langle 23, 38, 38, 38, 38 \rangle$
30	$\langle o_2, o_2, o_2, o_2 \rangle$ $\langle 0, 1, 2, 3 \rangle$
34	$\langle C_{Demo}, C_{Demo}, C_{Demo}, C_{Demo} \rangle$
41	$\langle C_{out} \rangle$

(a)

Bytecode	Sequences
56	$\langle 55, 53, 55, 53 \rangle$

(b)

Table I. Trace table for (a) method `main(String[])` and (b) method `foo(int)` of Figure 1

bytecodes. Bytecode 30 defines an element of an array (*i.e.* define `arr[i]` at line 27 of the source program). Note that for this `iastore` bytecode, two sequences are stored. These sequences correspond to the two operands of the bytecode, namely: identities of accessed array objects (*i.e.* $\langle o_2, o_2, o_2, o_2 \rangle$) and indices of accessed array element (*i.e.* $\langle 0, 1, 2, 3 \rangle$). Both sequences consist of four elements, because bytecode 30 is executed four times and accesses $o_2[0], o_2[1], o_2[2], o_2[3]$ respectively; each element in a sequence records one operand for one execution of bytecode 30. Bytecodes 34 and 41 invoke virtual methods; the operand sequences record classes which invoked methods belong to, where C_{Demo} represents class `Demo` and C_{out} represents the standard output stream class. Bytecodes 22, 24 and 56 have multiple predecessors in the control flow graph. For example, bytecode 56 (*i.e.* `return ret` at line 9 of the source program) has two predecessors: bytecode 53 (*i.e.* after `ret=2` at line 6 of the source program) and bytecode 55 (*i.e.* `ret=5` at line 8 of the source program). The sequence recorded for bytecode 56 (see Table I(b)) captures bytecodes executed immediately before bytecode 56, which consists of bytecodes 55 and 53 in this example. Note that every method in our example program has only one `return` bytecode, so no `return` bytecode is monitored and no trace of the method-exit is collected.

Clearly, different invocations of a method within a program execution can result in different traces. The difference in two executions of a method results from different operands of bytecodes within the method. These different traces are all stored implicitly via the sequences of operands used by the traced bytecodes. As an example, consider the trace table of method `foo` shown in Table I(b). The different traces of `foo` result from the different outcomes of its only conditional branch, which is captured by the trace sequence for predecessors of bytecode 56 in Figure 1, as shown in Table I(b).

2.2 Overview of SEQUITUR

So far, we have described how the bytecode operand sequences representing control flow, data flow, or dynamic call graph are separated in an execution trace. We now employ a lossless compression scheme to exploit the regularity and repetition of these sequences. Our technique is an extension of the SEQUITUR, a lossless data compression algorithm [Nevill-Manning and Witten 1997] which has been used to represent control flow information in program traces [Larus 1999]. First we briefly describe SEQUITUR.

The SEQUITUR algorithm represents a finite sequence σ as a context free grammar whose language is the singleton set $\{\sigma\}$. It reads symbols one-by-one from the input sequence and restructures the rules of the grammar to maintain the following invariants: (A) no pair of adjacent symbols appear more than once in the grammar, and (B) every rule (except the rule defining the start symbol) is used more than once. To intuitively understand the algorithm, we briefly describe how it works on a sequence 123123. As usual, we use capital letters to denote non-terminal symbols.

After reading the first four symbols of the sequence 123123, the grammar consists of the single production rule

$$S \rightarrow 1, 2, 3, 1$$

where S is the start symbol. On reading the fifth symbol, it becomes

$$S \rightarrow 1, 2, 3, 1, 2$$

Since the adjacent symbols 1, 2 appear twice in this rule (violating the first invariant), SEQUITUR introduces a non-terminal A to get

$$S \rightarrow A, 3, A \quad A \rightarrow 1, 2$$

Note that here the rule defining non-terminal A is used twice. Finally, on reading the last symbol of the sequence 123123 the above grammar becomes

$$S \rightarrow A, 3, A, 3 \quad A \rightarrow 1, 2$$

This grammar needs to be restructured since the symbols $A, 3$ appear twice. SEQUITUR introduces another non-terminal to solve the problem. We get the rules

$$S \rightarrow B, B \quad B \rightarrow A, 3 \quad A \rightarrow 1, 2$$

However, now the rule defining non-terminal A is used only once. So, this rule is eliminated to produce the final result.

$$S \rightarrow B, B \quad B \rightarrow 1, 2, 3$$

Note that the above grammar accepts only the sequence 123123.

2.3 Capturing Contiguous Repeated Symbols in SEQUITUR

One drawback of SEQUITUR is that it cannot efficiently represent contiguous repeated symbols, including both terminal and non-terminal symbols. However, contiguous repeated symbols are not uncommon in program traces. Consider the example in Figure 1. Bytecode 24 (*i.e.* $i < 4$ at line 26 of the source program in Figure 1) has two predecessors: bytecode 23 (*i.e.* $i = 0$ at line 26 of the source program in Figure 1) and bytecode 38 (after $i++$ at line 26 of the source program in Figure 1). The `for` loop is iterated four times, so the predecessor sequence for bytecode 24 is: $\langle 23, 38, 38, 38, 38 \rangle$ as shown in Table I(a). To represent this sequence, SEQUITUR will produce the following rules:

$$S \rightarrow 23, A, A \quad A \rightarrow 38, 38$$

In general, if the `for` loop is iterated k times, SEQUITUR needs $O(\lg k)$ rules in this fashion. To exploit such contiguous occurrences in the sequence representation, we propose the Run-Length Encoded SEQUITUR (RLESe).

RLESe constructs a context free grammar to represent a sequence *on the fly*; this contrasts with the work of [Reiss and Renieris 2001] which modifies the SEQUITUR grammar post-mortem. The right side of each rule is a sequence of “**nodes**”. Each node $\langle sym : n \rangle$ consists of a symbol sym and a counter n (*i.e.* run length), representing n contiguous occurrences of sym . RLESe can exploit contiguous repeated symbols, and represent the above trace sequence $\langle 23, 38, 38, 38, 38 \rangle$ of bytecode 24 using the following one rule:

$$S \rightarrow 23 : 1, 38 : 4$$

The RLESe algorithm constructs a context free grammar by reading from the input sequence symbol by symbol. On reading a symbol sym , a node $\langle sym : 1 \rangle$ is appended to the end of the start rule, and grammar rules are re-structured by preserving following three properties. The first property is unique to RLESe, resulting from its maintenance of contiguous occurrences of grammar nodes. The second and third properties are taken (and modified) from SEQUITUR.

- (1) *No contiguous repeated symbols property.* This property states that each pair of adjacent nodes contains different symbols. Continuous repeated symbols will be encoded within the run-length.
- (2) *Digram uniqueness property.* This property means that no *similar* digrams appear in resulting grammar rules. Here a digram refers to two consecutive nodes on the right side of a grammar rule. Two digrams are *similar* if their nodes contain the same pair of symbols *e.g.* $\langle a : 2, X : 2 \rangle$ is similar to $\langle a : 3, X : 4 \rangle$, but $\langle a : 3, X : 2 \rangle$ is not similar to $\langle X : 2, a : 3 \rangle$.
- (3) *Rule utility property.* This rule states that every rule (except the start rule S) is referenced more than once. When a rule is referenced by only one node and the run length n of that node equals 1, the reference will be replaced with the right hand side of this rule.

To maintain the digram uniqueness property in RLESe, we might need to split nodes during grammar construction. This split operation allows the algorithm to obtain duplicated *identical* digrams, and represent them by one grammar rule for potential space saving. Two digrams are *identical* if they have the same pairs of symbols and counters. For example, digram $\langle a : 2, X : 2 \rangle$ is identical to $\langle a : 2, X : 2 \rangle$, but digram $\langle a : 2, X : 2 \rangle$ is not identical to $\langle a : 3, X : 4 \rangle$.

Given two similar digrams $\langle sym_1 : n_1, sym_2 : n_2 \rangle$, and $\langle sym_1 : n'_1, sym_2 : n'_2 \rangle$, we can split at most two nodes to obtain two occurrences of $\langle sym_1 : \min(n_1, n'_1), sym_2 : \min(n_2, n'_2) \rangle$, where $\min(n_1, n'_1)$ denotes the minimum of n_1 and n'_1 . Consider bytecode 24 of Figure 1, which corresponds to the termination condition $i < 4$ of loop at line 26 of the source program. Assume that in some execution, such a loop is executed twice, one time with 6 iterations, and another time with 8 iteration. Recall that bytecode 24 has two predecessors: bytecode 23 (corresponding to $i = 0$ of the source program in Figure 1) and bytecode 38 (after $i++$ of the source program in Figure 1). The predecessor sequence for bytecode 24 is:

$$S \rightarrow 23 : 1, 38 : 6, 23 : 1, 38 : 8$$

To ensure digram uniqueness property, we will split the node $\langle 38 : 8 \rangle$ to a digram $\langle 38 : 6, 38 : 2 \rangle$. This is to remove duplicate occurrences of similar digrams as:

$$S \rightarrow A : 2, 38 : 2 \quad A \rightarrow 23 : 1, 38 : 6$$

The split operation introduces more nodes (at most two) into the grammar, but may save space when the identical digram appears frequently in the sequence being compressed.

In addition to the run-length encoding performed in RLESe, we also need to modify the terminal symbols fed into RLESe algorithm. In particular, we need to employ difference representations in memory reference sequences. For example, the sequence $\langle 0, 1, 2, 3 \rangle$ in Table I(a), which represents the indices of the array elements defined by bytecode 30 in Figure 1, cannot be compressed. By converting it into its difference representation as $\langle 0, 1, 1, 1 \rangle$, RLESe can compactly represent the trace sequence with one rule, as

$$S \rightarrow 0 : 1, 1 : 3$$

As with SEQUITUR [Nevill-Manning and Witten 1997], the RLESe compression algorithm is linear in both space and time, assuming that it takes constant time to find similar digrams. *Detailed time/space complexity analysis of the RLESe compression scheme is presented in our technical report [Wang and Roychoudhury 2007].* Experiments comparing RLESe with SEQUITUR (refer Section 5) show that RLESe can often achieve competitive compression ratio in less time. This is because RLESe can postpone re-constructing the grammar so the grammar rules are re-constructed less frequently. That is, on reading a symbol sym from input, instead of appending node $\langle sym : 1 \rangle$ and re-constructing the grammar immediately, the RLESe algorithm first compares sym against the last node $\langle sym' : n \rangle$ of the start rule. If sym is the same as sym' , the node $\langle sym' : n \rangle$ is updated to $\langle sym' : n + 1 \rangle$ and the grammar is not further re-constructed. If not, node $\langle sym : 1 \rangle$ is appended to the end of the start rule, and the grammar is re-structured so as to preserve the three properties of RLESe.

3. DYNAMIC SLICING

In this section, we focus on how to perform dynamic slicing of Java programs. Our dynamic slicing algorithm operates on the compact bytecode traces described in the last section. Dynamic slicing is performed w.r.t. a slicing criterion (H, α, V) , where H is an execution trace, α represents some bytecodes the programmer is interested in, and V is a set of variables referenced at these bytecodes. The dynamic slice contains all bytecodes which have affected values of variables in V referenced at last occurrences of α in the execution trace H .

Often, the user understands a Java program at the statement level. Thus, the user-defined criterion is often of the form (I, l, V) , where I is an input, and l is a line number of the source program; the user is interested in statements (instead of bytecodes) which have affected values of variables in V referenced at last occurrences of statements at l during the execution with input I . This form is a little different from our bytecode based slicing criterion (H, α, V) . In this case, program execution with input I produces the trace H , and α represents the bytecodes corresponding to statements at l . The user is interested in statements corresponding to bytecodes included in the dynamic slice. *In order to map bytecodes to a line number of the*

source file and vice versa, we use the `LineNumberTable` attribute in a Java's class file [Lindholm and Yellin 1999] which describes such a map.

The dynamic slice includes the closure of dynamic control and data dependencies from the slicing criterion. Formally, a dynamic slice can be defined over *Dynamic Dependence Graph* (DDG) [Agrawal and Horgan 1990]. The DDG captures dynamic control and data dependencies between bytecode occurrences during program execution. Each node of the DDG represents one particular occurrence of a bytecode; edges represent dynamic data and control dependencies. The dynamic slice is then defined as follows.

Definition 3.0.1. Dynamic slice for a slicing criterion consists of all bytecodes whose occurrence nodes can be reached from the node(s) representing the slicing criterion in the DDG.

We can construct the DDG as well as the dynamic slice during a backwards traversal of the execution trace.

3.1 Core Algorithm

Figure 2 presents an inter-procedural dynamic slicing algorithm, which returns the dynamic slice defined in Definition 3.0.1. Before slicing, we pre-compute the static *control flow graph* for the program. In addition, we pre-compute the *control dependence graph* [Ferrante et al. 1987], where each node in the graph represents one bytecode, and an edge from node v to v' represents that bytecode of v' decides whether bytecode of v will be executed. This static control dependence graph is used at lines 22 and 23 of the algorithm in Figure 2 to detect dynamic control dependencies.

Lines 1-5 of Figure 2 introduce five global variables for the slicing algorithm, including the slicing criterion. During dynamic slicing, we maintain δ , a list of variables whose values need to be explained, φ , the set of bytecode occurrences which have affected the slicing criterion, *op_stack*, a operand stack for simulation (see Section 3.3), and *fram*, a stack of frames for method invocations. The dynamic slice includes all bytecodes whose occurrences appear in φ at the end of the algorithm. For every method invocation during trace collection, we create a frame for this invocation during slicing. Each frame contains the method name and a γ set; the γ set includes bytecode occurrences β , where (a) β belongs to this method invocation, and (2) the dynamic control dependencies w.r.t. β need to be explained.

Initially we will set δ , φ , *op_stack*, and *fram* to empty. However, if the program had been aborted in the middle of an execution, the call stack *fram* is initialized differently. In this case, our tracing will record the call stack at the point of abort, call it *stk_abort*. Our dynamic slicing algorithm then initializes *fram* to *stk_abort* and proceeds by backward traversal of the execution trace.

Our slicing algorithm traverses the program's execution trace backwards, starting from the last executed bytecode recorded in the trace H . For each occurrence β of bytecode b_β (i.e. β represents one execution of the bytecode b_β) encountered during the backward traversal for slicing, a frame is created and pushed to *fram* whenever b_β is a return bytecode (lines 9-12 of Figure 2). The γ set of the new frame is initialized to empty (line 11 of Figure 2), since no bytecode occurrence for this method invocation has been traversed. If b_β is a method invocation byte-

code, a frame is popped from *fram* (lines 13-16 of Figure 2). The dynamic slicing algorithm checks whether the encountered bytecode occurrence β has affected the slicing criterion during trace collection, at lines 19-31 of Figure 2. In particular, line 19 of Figure 2 checks if β is the slicing criterion. Line 22 of Figure 2 checks dynamic control dependencies when b_β is a control transfer bytecode. The method `computeControlDependence(b_β , curr_fram, last_fram)` returns true iff. any bytecode occurrence included in the dynamic slice is dynamically control dependent on β , where *curr_fram* is the top of the stack *fram*, and *last_fram* captures the frame popped from *fram* (whenever b_β is a method invocation bytecode). More specifically, the `computeControlDependence` method returns true iff.

- b_β is a conditional branch bytecode, and some bytecode whose occurrence appears in *curr_fram*. γ is statically control dependent on b_β (intra-procedural control dependence check), or
- b_β is method invocation bytecode, and the *last_fram*. γ set is not empty, that is some of the bytecode occurrences in the method body are included in the slice (inter-procedural control dependence check).

Bytecode occurrences which are dynamically control dependent on β are then removed from *curr_fram*. γ at line 24 of Figure 2, because their dynamic control dependencies have just been explained. Line 27 of Figure 2 checks dynamic data dependencies. When the algorithm finds that the bytecode occurrence β has affected the slicing criterion (*i.e.* any of the three checks in lines 19, 22 and 27 of the algorithm in Figure 2 succeeds), β is included into *curr_fram*. γ and used variables are included into δ (at lines 32-34 of Figure 2), in order to find bytecode occurrences which have affected β and hence the slicing criterion. The simulation operand stack *op_stack* is also properly updated at line 35 for further check of data dependencies (this is explained in Section 3.3). The dynamic control and data dependencies checks (lines 22 and 27 of Figure 2) can be reordered, since they are independent of each other.

In the rest of this section, we elaborate on the underlying subtle issues in using the slicing framework of Figure 2. Section 3.2 presents how to traverse the execution backwards without decompressing the compact bytecode trace. Section 3.3 explains the intricacies of our dynamic data dependence computation in presence of Java’s stack based execution. Section 3.4 illustrates the dynamic slicing algorithm with an example and Section 3.5 shows the correctness and cost of our dynamic slicing algorithm.

3.2 Backward Traversal of Trace without decompression

The dynamic slicing algorithm in Figure 2 traverses the program execution backwards, starting from the last executed bytecode recorded in the trace *H* (line 7 of Figure 2). The algorithm proceeds by iteratively invoking the `getPrevBytecode` method to obtain the bytecode executed prior to current occurrence β of bytecode b_β during trace collection. Figure 3 presents the `getPrevBytecode` method. The algorithm first retrieves the last executed bytecode within the same method invocation of b_β into b_{last} . It returns b_{last} if b_{last} does not cross method boundaries (lines 2-7 of Figure 3). If b_{last} invokes a method *meth*, the last executed return bytecode of method *meth* is returned (lines 8-15 of Figure 3). If b_{last} represents the start of

```

1  ( $H, \alpha, V$ )= the slicing criterion
2   $\delta = \emptyset$ , a set of variables whose values need to be explained
3   $\varphi = \emptyset$ , the set of bytecode occurrences which have affected the slicing criterion
4  op_stack= empty, the operand stack for simulation
5  fram= empty, the frames of the program execution
6  dynamicSlicing()
7     $b_\beta$  = get last executed bytecode from  $H$ ;
8    while ( $b_\beta$  is defined)
9      if ( $b_\beta$  is a return bytecode)
10         new_fram= createFrame();
11         new_fram. $\gamma = \emptyset$ ;
12         push(fram, new_fram);
13      if ( $b_\beta$  is a method invocation bytecode)
14         last_fram = pop(fram);
15      else
16         last_fram = null;
17       $\beta$  = current occurrence of bytecode  $b_\beta$ ;
18      curr_fram = the top of fram;
19      if ( $\beta$  is the last occurrence of  $b_\beta$  in  $H$ , and  $b_\beta \in \alpha$ )
20         use_vars =  $V \cap$  variables used at  $\beta$ ;
21          $\varphi = \varphi \cup \{\beta\}$ ;
22      if (computeControlDependence( $b_\beta$ , curr_fram, last_fram) )
23          $BC = \{\beta' \mid \beta' \in \text{curr\_fram}.\gamma \text{ and } \beta' \text{ is dynamically control dependent on } \beta\}$ ;
24         curr_fram. $\gamma = \text{curr\_fram}.\gamma - BC$  ;
25         use_vars = variables used at  $\beta$ ;
26          $\varphi = \varphi \cup \{\beta\}$ ;
27      if (computeDataDependence( $\beta$ ,  $b_\beta$ ))
28         def_vars = variables defined at  $\beta$ ;
29          $\delta = \delta - \text{def\_vars}$ ;
30         use_vars = variables used at  $\beta$ ;
31          $\varphi = \varphi \cup \{\beta\}$ ;
32      if ( $\beta \in \varphi$ )
33         curr_fram. $\gamma = \text{curr\_fram}.\gamma \cup \{\beta\}$ ;
34          $\delta = \delta \cup \text{use\_vars}$ ;
35         updateOpStack( $\beta$ ,  $b_\beta$ );
36          $b_\beta = \text{getPrevBytecode}(\beta, b_\beta)$ ;
37      return bytecodes whose occurrences appear in  $\varphi$ ;

```

Fig. 2. The dynamic slicing algorithm

a method, the bytecode which invokes current method is returned (lines 16-17 of Figure 3).

The `getPrevBytecode` method has to retrieve last executed bytecode from the compact bytecode trace H . For this purpose, it needs to traverse the predecessor sequence of a bytecode with multiple predecessors. Since such sequences are compactly stored as RLESe grammars, we need to efficiently traverse RLESe grammars; this is accomplished by the method `getLast`. The `getLast` method gets the last executed predecessor from a RLESe grammar G without decompression, using a root-to-leaf path π in G . The slicing algorithm maintains such a path π for each compressed RLESe sequence G to clearly mark which portion of G has been already visited. We now explain our efficient traversal over the RLESe representation.

```

1  getPrevBytecode ( $\beta$ : bytecode occurrence,  $b_\beta$ : bytecode)
2  if ( $b_\beta$  has exactly one predecessor in the control flow graph)
3       $b_{last}$  = the predecessor bytecode;
4  else
5       $G$  = compressed control flow operand sequence for  $b_\beta$  in the bytecode trace  $H$ ;
6       $\pi$  = a root-to-leaf path for  $G$ ;
7       $b_{last}$  = getLast( $G, \pi$ );
8  if ( $b_{last}$  is a method invocation bytecode)
9       $meth$  = the method invoked by  $\beta$ ;
10     if ( $meth$  has exactly one return bytecode)
11         return the return bytecode;
12     else
13          $G'$  = compressed operand sequence for exit of  $meth$  in trace  $H$ ;
14          $\pi'$  = a root-to-leaf path for  $G'$ ;
15         return getLast( $G', \pi'$ );
16 if ( $b_{last}$  represents the start of a method)
17     return the bytecode which invokes current method;
18 return  $b_{last}$ ;

```

Fig. 3. Get the previous executed bytecode during backward traversal of the execution trace.

In our trace compression scheme, all operand sequences are compressed using RLESe. The dynamic slicing algorithm traverses these sequences from the end to extract predecessors for computation of control flow, and to extract identifies of accessed variables for computation of data flow. For example, consider the operand sequence of array indices for bytecode 56 in Table I(b), which is $\langle 55, 53, 55, 53 \rangle$. During dynamic slicing on the program of Figure 1, we traverse this sequence backwards, that is, from the end. At any point during the traversal, we mark the last visited operand (say $\langle 55, 53, 55, \overline{53} \rangle$) during slicing. The sequence beginning with the marked operand (*i.e.* $\langle \overline{53} \rangle$) has been visited. When the slicing algorithm tries to extract next operand from the operand sequence, we use this mark to find last unvisited element (*i.e.* value 55 in this example). We now describe how such markers can be maintained and updated in the RLESe grammar representation.

The RLESe grammar of a sequence σ can simply be represented as a directed acyclic graph (DAG). The RLESe grammar consists of run-length annotated symbols of the form $\langle sym : n \rangle$ where sym is a terminal or non-terminal symbol and n denotes a run-length, representing n contiguous occurrences of symbol sym . Let us consider an example where the operand sequence is $\langle abbabbabb \rangle$. The RLESe compression algorithm will produce the following rules to represent this operand sequence:

$$S \rightarrow A : 2, c : 1, A : 1 \quad A \rightarrow a : 1, b : 2$$

Small letters denote terminal symbols in the operand sequence, and capital letters denote non-terminal symbols. Figure 4(a) shows corresponding DAG representation, where dashed circles represent non-terminal symbols, circles represent terminal symbols, and edges are annotated by run-lengths. For example, the run-length annotated symbol $\langle A : 1 \rangle$ at the end of the start rule is captured by the node A and the incoming edge $S \xrightarrow{1} A$ for node A . We then use a root-to-leaf path π over the

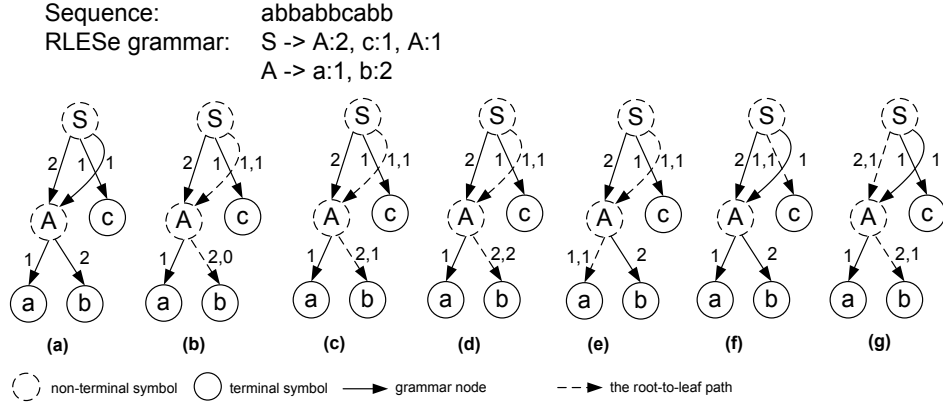


Fig. 4. An example showing extraction of operand sequence (without decompression) from RLESe representation.

DAG representation to mark the symbol in σ that was last visited during backward traversal. For every edge of the path π , we maintain both the run length n of corresponding grammar node $X = \langle sym : n \rangle$, and a visitation counter $k \leq n$, where k denotes the number of times that node X has been visited so far. For example, in the edge $A \xrightarrow{2,1} b$ in Figure 4(c), 2 represents the run length of grammar node $\langle b : 2 \rangle$, and 1 represents that this node has been visited once.

The dynamic slicing algorithm maintains one root-to-leaf path π for every compressed operand sequence. The path π is initialized from the root to the rightmost leaf node in the DAG, and the visitation counter annotated for the last edge in π is set to 0, since no symbol has been visited. The slicing algorithm then uses the path π to find the last unvisited terminal symbol of the RLESe grammar by invoking the `getLast` method of Figure 5. In the `getLast` method, G is the DAG representation of the RLESe grammar for a sequence σ and π is the root-to-leaf path for the symbol in σ that was last visited. The `getLast` method returns the last unvisited symbol and updates the path π . Note that the “immediate left sibling” of an edge $e = x \rightarrow y$ is the edge $e' = x \rightarrow z$ where node z is the immediate left sibling of node y in the graph; this notion is used in lines 10 and 11 of Figure 5.

For the example operand sequence $\langle abbabbcabb \rangle$, Figure 4(b) shows the initialized root-to-leaf path π for the RLESe grammar. Figure 4(c-g) present the resultant path π by calling the `getLast` method of Figure 5 each time, and the symbol of the leaf node pointed by the path π is returned as the last unvisited symbol. For example, Figure 4(c) shows the path π after the first calling the `getLast` method. The path π includes two edges (*i.e.* $S \xrightarrow{1,1} A$ and $A \xrightarrow{2,1} b$), representing both edges have been visited once. The leaf node b is referenced by π . Thus, b is returned by the `getLast` method, which represents the last symbol b in the original sequence $\langle abbabbcabb \rangle$.

With the `getLast` algorithm in Figure 5, we can extract an operand sequence efficiently. Given the grammar for a sequence with length N , the `getLast` method will be invoked N times to extract the entire sequence. The overall space overhead

```

1  getLast( $G$ : Grammar,  $\pi$ : path in  $G$ )
2       $e$  = the last edge of  $\pi$ ;
3      while ( $e$  is defined)
4          let  $e = sym_1 \xrightarrow{n_1, k_1} sym'_1$ ;
5          if ( $k_1 < n_1$ )
6              break;
7          else
8              remove edge  $e$  from  $\pi$ ;
9              change the annotation of edge  $e$  from  $(n_1, k_1)$  to  $(n_1)$ ;
10              $Sib_e$  = immediate left sibling of  $e$  in  $G$ ;
11             if (such a sibling  $Sib_e$  exists)
12                  $e = Sib_e$ ;
13                 break;
14             else
15                  $e$  = last edge of  $\pi$ ;
16         let  $e = sym_2 \xrightarrow{n_2, k_2} sym'_2$ ;
17         change the annotation of edge  $e$  from  $(n_2, k_2)$  to  $(n_2, k_2 + 1)$ ;
18          $G_X$  = DAG rooted at node  $sym'_2$  within  $G$ ;
19         for (each edge  $e'$  from node  $sym'_2$  to rightmost leaf node of  $G_X$ )
20             insert edge  $e'$  into  $\pi$ ;
21             let  $e' = sym_3 \xrightarrow{n_3} sym'_3$ ;
22             change the annotation of edge  $e'$  from  $(n_3)$  to  $(n_3, 1)$ ;
23         return symbol in rightmost leaf node of  $G_X$ ;

```

Fig. 5. One step in the backward traversal of a RLESe sequence (represented as DAG) without decompressing the sequence.

is $O(N)$, and the overall time overhead is $O(N)$. The space overhead is caused by maintaining the root-to-leaf path π , which is used by all invocations of the `getLast` method to extract a sequence. The length of path π is linear in the number of grammar rules (the grammar has no recursive rules). There are fewer grammar rules than grammar nodes, and the number of grammar nodes is bounded by the length of the original sequence. Thus, the space overhead to extract the entire sequence is $O(N)$.

The time overhead to extract the entire sequence comes from the time to access edges and nodes in the DAG. Whenever a node with terminal symbol is accessed, the `getLast` method immediately returns the terminal symbol. So, the total number of times to access node with terminal symbol is $O(N)$ in order to extract a sequence with length N . The total number of accesses to non-terminal nodes is $O(\sum_i \frac{N}{2^i}) = O(N)$. Consequently, the time overhead to extract the entire sequence from a RLESe grammar (by repeatedly invoking `getLast` to get the last symbol which has not been visited) is $O(N)$.

3.3 Computing Data Dependencies

The typical way to detect dynamic data dependencies is to compare addresses of variables defined/used by bytecode occurrences (line 27 of Figure 2). However, this is complicated by Java's stack based architecture. During execution, the Java virtual machine uses an operand stack to hold partial results of execution. Thus, dynamic data dependence exists between bytecode occurrences β and β' , when


```

1  updateOpStack ( $\beta$ : bytecode occurrence,  $b_\beta$ : bytecode)
2      for ( $i = 0$ ;  $i < def\_op(b_\beta)$ ;  $i = i + 1$ )
3          pop(op_stack);
4      for ( $i = 0$ ;  $i < use\_op(b_\beta)$ ;  $i = i + 1$ )
5          push(op_stack,  $\beta$ );

```

Fig. 6. Maintain the simulation stack *op_stack*.

a value is pushed into the operand stack by β and is popped by β' . Consider the program in Figure 1 as an example. Assume that statement 27 of the source program is executed, and the corresponding trace at the level of bytecode is $\langle 27^1, 28^2, 29^3, 30^4 \rangle$ where 27^1 means that the first element of the trace is bytecode 27 and so on. Bytecode occurrence 30^4 (which defines the array element `arr[i]` at line 27 of the source program) is dynamically data dependent on bytecode occurrences 27^1 , 28^2 and 29^3 (which load local variables `k` and `i`, array object reference `arr` at line 27 of the source program, respectively). The three bytecode occurrences push three values into the operand stack, all of which are popped and used by bytecode occurrence 30^4 .

Clearly, dynamic data dependencies w.r.t. local variables and fields can be easily detected by comparing the addresses (or identities) of accessed variables. However, detecting data dependencies w.r.t. the operand stack requires *reverse* simulating the operand stack (since we traverse the trace backwards). Figure 6 presents how to maintain the stack *op_stack* for simulation, which is used by the dynamic slicing algorithm at line 35 of Figure 2. We pop the simulation stack for defined operands, and push used operands into the simulation stack. The function *def_op*(b_β) (*use_op*(b_β)) at line 2 (4) of Figure 6 returns the number of operands defined (used) by bytecode b_β . Note that the stack simulated during slicing does not contain actual values of computation. Instead, each entry of the stack stores the bytecode occurrence which pushed the entry into the stack.

Figure 7 shows the method to determine whether a bytecode occurrence has affected the slicing criterion via dynamic data dependencies. This method is used by the dynamic slicing algorithm at line 27 of Figure 2. If a bytecode occurrence β defines a variable which needs explanation (lines 2-10 of Figure 7), or β defines a partial result which needs explanation (lines 11-13 of Figure 7), the method returns *true* to indicate that β has affected the slicing criterion. A partial result needs explanation if the bytecode occurrence which pushes corresponding entry into the stack has already been included in φ . The function *def_op*(b_β) at line 11 of Figure 7 returns the number of operands defined by bytecode b_β . Our dynamic slicing algorithm needs the addresses of variables accessed by each bytecode occurrence for detecting data dependencies (lines 20, 25, 28 and 30 of Figure 2 and lines 3-8 of Figure 7). If a local variable or a static field is accessed, the address can be found from the class files. If an object field or an array element is accessed, the address can be found from operand sequences of corresponding bytecode in the compact bytecode trace.


When the algorithm detects data dependencies via reverse stack simulation, it requires analyzing some bytecodes whose operands are not traced. Consider the

```

1  computeDataDependence ( $\beta$ : bytecode occurrence,  $b_\beta$ : bytecode)
2      if ( $\beta$  defines a variable)
3          if ( $\beta$  defines a static field or local variable)
4               $def\_loc =$  get address of the defined static field or local variable from class files;
5          if ( $\beta$  defines an object field or an array element)
6               $G =$  compressed operand sequence for  $b_\beta$  in the compact bytecode trace  $H$ 
7               $\pi =$  a root-to-leaf path for  $G$ ;
8               $def\_loc =$   $getLast(G, \pi)$ ;
9          if ( $def\_loc \in \delta$ )
10             return true;
11      $\omega =$  the set of bytecode occurrences in top  $def\_op(b_\beta)$  entries of  $op\_stack$ ;
12     if ( $\omega \cap \varphi \neq \emptyset$ )
13         return true;
14     return false;

```

Fig. 7. Detect dynamic data dependencies for dynamic slicing

Fig. 8. An example shows the *op_stack* after each bytecode occurrence encountered during backward traversal

program in Figure 1 as an example. The statement `if (j%2==1)` at line 5 of the source program corresponds to bytecode sequence $\langle 46^1, 47^2, 48^3, 49^4, 50^5 \rangle$. Figure 8 shows the *op_stack* after processing each bytecode occurrence during backward traversal. Note that the operands of bytecode 48 (*i.e.* bytecode `irem` which stands for the mathematical computation “%”) are not traced. When bytecode occurrence 48^3 is encountered, `computeDataDependence` method in Figure 7 can detect that 50^5 is dynamically data dependent on 48^3 . In addition, the bytecode 48 will also update the *op_stack*, as shown in Figure 8. As we can see from the example, in order to detect implicit data dependencies involving data transfer via the operand stack, it is important to know which bytecode occurrence pushes/pops an entry from the *op_stack*. The actual values computed by the bytecode execution are not important. This highlights difference between our method and the work on Abstract Execution [Larus 1990]. In Abstract Execution, a small set of events are recorded and these are used as guide to execute a modified program. In our work, we record some of the executed bytecodes in our compressed trace representation. However, the untraced bytecodes are not re-executed during the analysis of the compressed trace.

3.4 Example

Consider the example program in Figure 1, and corresponding compressed trace in Table I. Assume that the programmer wants to find which bytecodes have affected

β	δ	$fram$		op_stack	$\in \varphi$
		method	γ		
41 ³¹	{}	main	{}	$\langle 41^{31}, 41^{31} \rangle$	
40 ³⁰	{k}	main	{40 ³⁰ }	$\langle 41^{31} \rangle$	*
39 ²⁹	{k}	main	{40 ³⁰ }	$\langle \rangle$	
26 ²⁸	{k}	main	{40 ³⁰ }	$\langle 26^{28}, 26^{28} \rangle$	
25 ²⁷	{k}	main	{40 ³⁰ }	$\langle 26^{28} \rangle$	
24 ²⁶	{k}	main	{40 ³⁰ }	$\langle \rangle$	
38 ²⁵	{k}	main	{40 ³⁰ }	$\langle \rangle$	
37 ²⁴	{k}	main	{40 ³⁰ }	$\langle \rangle$	
36 ²³	{}	main	{40 ³⁰ , 36 ²³ }	$\langle 36^{23} \rangle$	*
35 ²²	{}	main	{40 ³⁰ , 36 ²³ }	$\langle 35^{22}, 35^{22} \rangle$	*
57 ²¹	{}	foo	{57 ²¹ }	$\langle 35^{22}, 57^{21} \rangle$	*
		main	{40 ³⁰ , 36 ²³ }		
56 ²⁰	{ret}	foo	{57 ²¹ , 56 ²⁰ }	$\langle 35^{22} \rangle$	*
		main	{40 ³⁰ , 36 ²³ }		
53 ¹⁹	{ret}	foo	{57 ²¹ , 56 ²⁰ }	$\langle 35^{22} \rangle$	
		main	{40 ³⁰ , 36 ²³ }		
52 ¹⁸	{}	foo	{57 ²¹ , 56 ²⁰ , 52 ¹⁸ }	$\langle 35^{22}, 52^{18} \rangle$	*
		main	{40 ³⁰ , 36 ²³ }		
51 ¹⁷	{}	foo	{57 ²¹ , 56 ²⁰ , 52 ¹⁸ , 51 ¹⁷ }	$\langle 35^{22} \rangle$	*
		main	{40 ³⁰ , 36 ²³ }		
50 ¹⁶	{}	foo	{57 ²¹ , 56 ²⁰ , 50 ¹⁶ }	$\langle 35^{22}, 50^{16}, 50^{16} \rangle$	*
		main	{40 ³⁰ , 36 ²³ }		
49 ¹⁵	{}	foo	{57 ²¹ , 56 ²⁰ , 50 ¹⁶ , 49 ¹⁵ }	$\langle 35^{22}, 50^{16} \rangle$	*
		main	{40 ³⁰ , 36 ²³ }		
48 ¹⁴	{}	foo	{57 ²¹ , 56 ²⁰ , 50 ¹⁶ , 49 ¹⁵ , 48 ¹⁴ }	$\langle 35^{22}, 48^{14}, 48^{14} \rangle$	*
		main	{40 ³⁰ , 36 ²³ }		
47 ¹³	{}	foo	{57 ²¹ , 56 ²⁰ , 50 ¹⁶ , 49 ¹⁵ , 48 ¹⁴ , 47 ¹³ }	$\langle 35^{22}, 48^{14} \rangle$	*
		main	{40 ³⁰ , 36 ²³ }		
46 ¹²	{j}	foo	{57 ²¹ , 56 ²⁰ , 50 ¹⁶ , 49 ¹⁵ , 48 ¹⁴ , 47 ¹³ , 46 ¹² }	$\langle 35^{22} \rangle$	*
		main	{40 ³⁰ , 36 ²³ }		
34 ¹¹	{}	main	{40 ³⁰ , 36 ²³ , 34 ¹¹ }	$\langle 35^{22}, 34^{11}, 34^{11} \rangle$	*
33 ¹⁰	{i}	main	{40 ³⁰ , 36 ²³ , 34 ¹¹ , 33 ¹⁰ }	$\langle 35^{22}, 34^{11} \rangle$	*
32 ⁹	{i, obj}	main	{40 ³⁰ , 36 ²³ , 34 ¹¹ , 33 ¹⁰ , 32 ⁹ }	$\langle 35^{22} \rangle$	*
31 ⁸	{i, obj, k}	main	{40 ³⁰ , 36 ²³ , 34 ¹¹ , 33 ¹⁰ , 32 ⁹ , 31 ⁸ }	$\langle \rangle$	*
30 ⁷	{i, obj, k}	main	{40 ³⁰ , 36 ²³ , 34 ¹¹ , 33 ¹⁰ , 32 ⁹ , 31 ⁸ }	$\langle 30^7, 30^7, 30^7 \rangle$	
29 ⁶	{i, obj, k}	main	{40 ³⁰ , 36 ²³ , 34 ¹¹ , 33 ¹⁰ , 32 ⁹ , 31 ⁸ }	$\langle 30^7, 30^7 \rangle$	
28 ⁵	{i, obj, k}	main	{40 ³⁰ , 36 ²³ , 34 ¹¹ , 33 ¹⁰ , 32 ⁹ , 31 ⁸ }	$\langle 30^7 \rangle$	
27 ⁴	{i, obj, k}	main	{40 ³⁰ , 36 ²³ , 34 ¹¹ , 33 ¹⁰ , 32 ⁹ , 31 ⁸ }	$\langle \rangle$	
26 ³	{i, obj, k}	main	{40 ³⁰ , 26 ³ }	$\langle 26^3, 26^3 \rangle$	*
25 ²	{i, obj, k}	main	{40 ³⁰ , 26 ³ , 25 ² }	$\langle 26^3 \rangle$	*
24 ¹	{i, obj, k}	main	{40 ³⁰ , 26 ³ , 25 ² , 24 ¹ }	$\langle \rangle$	*

Table II. An example to show each stage of the dynamic slicing algorithm in Figure 2. The column β shows bytecode occurrences in the trace being analyzed.

the value of k at line 31 of the source program. Table II shows each stage using the dynamic slicing algorithm in Figure 2 w.r.t. the k . For simplicity, we do not illustrate slicing over the entire execution, but over last executed eight statements – $\langle 26, 27, 28, 5, 6, 9, 26, 31 \rangle$. The corresponding bytecode sequence is a sequence of thirty-one bytecode occurrences shown in the first column of Table II. For each bytecode occurrence, the position number $1, \dots, 31$ of the bytecode occurrence in the sequence is marked as superscript for the sake of clarity.

For each bytecode occurrence β encountered during backward traversal, one row of Table II shows resultant δ , $fram$, op_stack and φ after analyzing β by our dynamic slicing algorithm. The \star in the last column indicates that the corresponding bytecode occurrence has affected the slicing criterion and is included into φ .

When bytecode occurrence 40^{30} is encountered, it is found to be the slicing criterion. The used variable \mathbf{k} is inserted to δ to find which bytecode occurrence defines \mathbf{k} ; the bytecode occurrence 40^{30} is inserted to γ for control dependency check; we pop 41^{31} from the operand stack *op_stack* because 40^{30} loads one value to the operand stack during trace collection. It should be noted that in this simple example, we refer to a variable with its name (*e.g.* \mathbf{k}), since both methods are invoked once, and every variable has a distinct name in this example. This is for simplicity of illustration. In the implementation, we use identifiers to distinguish between variables from same/different method invocation.

After 40^{30} , bytecode occurrence 39^{29} is encountered during backward traversal and so on. We omit the details of the entire traversal but highlight some representative bytecode occurrences.

- After analyzing bytecode occurrence 56^{20} , the slicing algorithm finds that bytecode 56 has two predecessors, and retrieves last unvisited value from operand sequence of bytecode 56 in Table I(b). Therefore, bytecode occurrence 53^{19} is next analyzed.
- When bytecode occurrence 30^7 is encountered, o_2 and 3 are retrieved from operand sequences of bytecode 30 in Table I(a), representing an assignment to array element $o_2[3]$. However, $o_2[3]$ is irrelevant to the slicing criterion, so neither δ nor γ is updated.

3.5 Proof of Correctness and Complexity Analysis

In this section we discuss the correctness proof and complexity analysis of our dynamic slicing algorithm.

THEOREM 3.5.1. *Given a slicing criterion, the dynamic slicing algorithm in Figure 2 returns dynamic slice defined in Definition 3.0.1.*

Proof Sketch: We only present the proof sketch here. *The full proof appears in our technical report [Wang and Roychoudhury 2007], which is available online.*

Let φ_i be the φ set after i loop iterations of the dynamic slicing algorithm in Figure 2, φ_* be the resultant φ set when the algorithm finishes, and β be the bytecode occurrence encountered at the i th loop iteration.

We prove the soundness of the algorithm by induction on loop iterations of the slicing algorithm, *i.e.* for any $\beta' \in \varphi_*$ we show that β' is reachable from the slicing criterion in the dynamic dependence graph (DDG).

We prove the completeness of the slicing algorithm, *i.e.* $\forall \beta'$ reachable from slicing criterion in the Dynamic Dependence Graph (DDG) $\Rightarrow \beta' \in \varphi_*$. Note that there is no cycle in the DDG, so we prove the completeness by induction on structure of the DDG. \square

Now we analyze the cost of the dynamic slicing algorithm in Figure 2. Given the compressed trace for an execution which executes N bytecodes, the space overhead of the slicing algorithm is $O(N)$, and the time overhead is $O(N^2)$.

The space overhead of the algorithm is caused by the maintenance of δ , φ , *op_stack*, *fram*, compressed operand sequences and root-to-leaf paths for every compressed sequence. The sizes of δ , φ , *op_stack* and *fram* are all $O(N)$. For δ , this is because one execution of a bytecode can use a constant number of variables;

for *op_stack*, this is because the number of operands popped from and pushed to the operand stack by one bytecode is bound by a constant; for *fram*, this is because the γ set of each method invocation in *fram* only contains bytecode occurrences for this invocation, and does not overlap with each other. Assume that each bytecode b_i has executed $\eta(b_i)$ times, so $\sum_{b_i} \eta(b_i) = N$. The size of all compressed operand sequences is $\sum_{b_i} O(\eta(b_i)) = O(N)$, because every bytecode has a fixed number of operand sequences, and the size of the compact representation is linear in the length of original operand sequence; *proof of this claim appears in our technical report [Wang and Roychoudhury 2007]*. The size of each root-to-leaf path is bound by the size of corresponding compressed operand sequence. Consequently, the overall space cost of the slicing algorithm is $O(N)$.

During dynamic slicing, the algorithm performs the following four actions for each occurrence β of bytecode b_β encountered during backward traversal of the execution trace.

- (1) extract operand sequences of bytecode b_β from the compressed trace for backward traversal,
- (2) perform slicing criterion, dynamic control/data dependency checks,
- (3) update δ , φ , *op_stack*, and *fram*,
- (4) get previous executed bytecode.

According to the complexity analysis of the `getLast` method which is presented in Section 3.2, it needs $O(\eta(b_i))$ time to extract an operand sequence of bytecode b_i which is executed $\eta(b_i)$ times during trace collection. Note that $\sum_i \eta(b_i) = N$. Consequently, the overall time overhead to perform action (1) to extract all operand sequences is $\sum_i O(\eta(b_i)) = O(N)$, that is, linear in the length of the original execution. After extracting operand sequences, the overall time to perform action (2) and (3) is clearly $O(N^2)$, since they may update/enquire sets δ , φ , *op_stack* and *fram* whose sizes are bound by N . This complexity can be improved further by using efficient data structures for set representation.

The `getPrevBytecode` method of Figure 3 may perform two actions in order to get the previous executed bytecode (action (4) in the preceding) — (a) get the predecessor bytecode from class files, or (b) extract last executed bytecode from compressed operand sequences. The overall time to perform get predecessor bytecode from class files is $O(N)$, since it needs constant time to get the predecessor bytecode from Java class files every time. The overall time to extract operand sequences is $O(N)$ as discussed earlier.

Consequently, the overall time cost of the dynamic slicing algorithm is $O(N^2)$.

4. RELEVANT SLICING

In the previous section, we presented dynamic slicing which can be used for focusing the programmer’s attention to a part of the program. However, there are certain difficulties in using dynamic slices for program debugging/understanding. In particular, dynamic slicing does not consider that the execution of certain statements is wrongly omitted. Consequently, all statements responsible for the error may not be included into the dynamic slice, and the slice may mislead the programmer.

As an example, consider the program fragment presented in Figure 9, which is a simplified version of the program in Figure 1. Let us assume that the statement

```

1  b = 1;
2  k = 1;
3  if (a > 1) {
4      if (b > 1){
5          k = 2
6      }
7  }
8  ... = k

```

Fig. 9. A “buggy” program fragment

$b=1$ at line 1 should be $b=2$ in a correct program. With input $a=2$, the variable k at line 6 is 1 unexpectedly. The execution trace is $\langle 1^1, 2^2, 3^3, 4^4, 6^5 \rangle$, where 1^1 means statement 1 is executed as the *first* statement and so on. If we use dynamic slicing to correct the bug w.r.t. the incorrect value of k at 6^5 , the dynamic slice computed by the slicing algorithm of Section 3 contains only lines 2 and 6. Changing line 2 may fix the bug, since such a change will change the value of k at line 6. Unfortunately, line 1, the actual bug, is excluded from the dynamic slice. In this example the error arises from the execution of line 5 being wrongly omitted, which is caused by the incorrect assignment at line 1. In fact, changing line 1 may cause the predicate at line 4 to be evaluated differently; then line 5 will be executed and the value of k at line 6 will be different. In other words, dynamic slicing does not consider the effect of the unexecuted statement at line 5.

The notion of *relevant slicing*, an extension of dynamic slicing, fills this caveat. Relevant slicing was introduced in [Agrawal et al. 1993; Gyimóthy et al. 1999]. Besides dynamic control and data dependencies, relevant slicing considers *potential dependencies* which capture the potential effects of unexecuted paths of branch and method invocation statements. The relevant slice include more statements which, if changed, may change the “wrong” behaviors w.r.t. the slicing criterion. In the example of Figure 9 with input $a=2$, 6^5 is *potentially dependent* on execution of the branch at line 4 (4^4), because if the predicate at 4^4 is evaluated differently, the variable k is re-defined and then used by 6^5 . Considering that 4^4 is dynamically data dependent on 1^1 , 1^1 has potential effects over 6^5 . Thus, line 1 is included into the resultant relevant slice.

In general, $Dynamic\ Slice \subseteq Relevant\ Slice \subseteq Static\ Slice$. In our experiments (refer Section 5), we show that the sizes of the relevant slices are close to the sizes of the corresponding dynamic slices. Like dynamic slices, the relevant slices are also computed w.r.t. a particular program execution (*i.e.* it only includes executed statements).

The rest of this section is organized as follows. In Section 4.1, we recall *past work* on potential dependencies, a notion which is crucial for computing the relevant slice. We then present our definition of relevant slice. Since the existing works on relevant slicing present a slice as a set of program statements, we first define a relevant slice as a set of statements. However, like our dynamic slicing algorithm, our relevant slicing algorithm (refer Figure 11 in Section 4.2) operates at the level of bytecodes. This is discussed in Section 4.2.

4.1 Background

In this subsection, we recapitulate the definition of potential dependence. The material in this subsection was mostly studied in [Agrawal et al. 1993; Gyimóthy et al. 1999].

Relevant slicing extends dynamic slicing by considering statements which may affect the slicing criterion. These statements are executed and do not affect the criterion. However, if these statements are changed, branches may be evaluated differently, or alternative methods may be invoked. We consider the following two situations for the execution of alternative code fragment due to program changes.

- (1) *Branch*: If the value of a variable used by the predicate of a branch statement is changed, the predicate may be evaluated differently.
- (2) *Method invocation*: If the programmer changes the assignment w.r.t. the object which invokes a method, or the declaration of parameters, an alternative method may be called.

In relevant slicing, we use the notion of *potential dependence* to capture the effects of these branch and method invocation statements. If these statements are evaluated differently, variables may be re-defined and affect the slicing criterion. We define potential dependence as follows; for any statement occurrence β , the corresponding statement is written as S_β .

Definition 4.1.1. Potential Dependence Given an execution trace H of a program P , for any statement instances β and β' appearing in H , β is *potentially dependent* on an earlier branch or method invocation statement instance β' if and only if all of the following conditions hold.

- (1) β is not (transitively) dynamically control/data dependent on β' in H .
- (2) there exists a variable v used by β s.t. v is not defined between β' and β in H and v may be defined along an outgoing edge l of statement $S_{\beta'}$ where $S_{\beta'}$ is the statement at statement occurrence β' . That is, there exists a statement X satisfying the following.
 - (a) X is (transitively) control dependent on $S_{\beta'}$ along the edge l , and
 - (b) X is an assignment statement which assigns to variable v or a variable u which may be aliased to v .
- (3) the edge l is not taken at β' in H .

The purpose of potential dependence is to find those statements which will be missed by dynamic slicing. If a branch or method invocation statement has affected the slicing criterion, we do not consider that it has potential influence, since the dynamic slice will always include such a statement. We use condition (1) to exclude dynamic control and data dependencies from potential dependencies. Conditions (2) and (3) ensure that β' has potential influence on β . The statement X in condition 2(b) appears in program P , but its execution is prohibited by the evaluation of branch/method invocation in β' . However, if it is executed (possibly due to a change in the statement $S_{\beta'}$), the value of variable v used at β will be affected. We cannot guarantee that v must be re-defined if β' is evaluated to take edge l . This is because even if β' is evaluated differently, execution of the assignment to v may

be guarded by other (nested) conditional control transfer statements. Furthermore, condition (2) requires computation of static data dependence. In the presence of arrays and dynamic memory allocations, we can only obtain conservative static data dependencies in general. We now proceed to define a relevant slice, and introduce our relevant slicing algorithm to compute such a slice.

4.2 Relevant Slice and the Relevant Slicing Algorithm

To define relevant slices, first we define the notion of an *Extended Dynamic Dependence Graph* (EDDG). The EDDG captures dynamic control dependencies, dynamic data dependencies and potential dependencies w.r.t. a program execution. It is an extension of the *Dynamic Dependence Graph* (DDG) described in [Agrawal and Horgan 1990]. Each node of the DDG represents one particular occurrence of a statement in the program execution; edges represent dynamic data and control dependencies. The EDDG extends the DDG with potential dependencies, by introducing a dummy node for each branch statement occurrence or a method invocation statement occurrence. For each statement occurrence β in the execution trace, a non-dummy node $nn(\beta)$ appears in the EDDG to represent this occurrence. In addition, if β is a branch statement or a method invocation statement, a dummy node $dn(\beta)$ also appears in the EDDG. As far as the edges are concerned, the following are the incoming edges of any arbitrary node $nn(\beta)$ or $dn(\beta)$ appearing in the EDDG.

- dynamic control dependence edges from non-dummy node $nn(\beta')$ to $nn(\beta)$, iff. β' is dynamically control dependent on β .
- dynamic data dependence edges from both non-dummy node $nn(\beta')$ and dummy node $dn(\beta')$ (if there is a dummy node for occurrence β') to $nn(\beta)$, iff. β' is dynamically data dependent on β .
- potential dependence edges from non-dummy node $nn(\beta')$ and dummy node $dn(\beta')$ (if there is a dummy node for β') to $dn(\beta)$, iff. β' is potentially dependent on β .

These dependencies can be detected during backwards traversal of the execution trace.

What is a Relevant Slice. The *relevant slice* is then defined based on the extended dynamic dependence graph (EDDG) as follows.

Definition 4.2.1. Relevant slice for a slicing criterion consists of all statements whose occurrence nodes can be reached from the node(s) for the slicing criterion in the Extended Dynamic Dependence Graph (EDDG).

In order to accurately capture the effects w.r.t. potential dependencies, we have introduced a dummy node into the EDDG for each occurrence of a branch or method invocation statement. The EDDG can then represent dynamic control dependence edges and potential dependence edges separately. This representation allows the reachability analysis for relevant slicing not to consider dynamic control dependence edges which are immediately after potential dependence edges, because such dynamic control dependencies cannot affect behaviors w.r.t. the slicing criterion. Figure 10 shows the EDDG for the example program in Figure 9 for input




Fig. 10. The EDDG for the program in Figure 9 with input $a=2$.

$a = 2$. The execution trace is $\langle 1^1, 2^2, 3^3, 4^4, 6^5 \rangle$. The resultant relevant slice w.r.t. variable k at line 6 consists of lines $\{1, 2, 4, 6\}$, because nodes $1^1, 2^2, 4^4$ and 6^5 can be reached from the criterion 6^5 . Note that statement occurrence 6^5 is potentially dependent on 4^4 , and 4^4 is dynamically control dependent on statement occurrence 3^3 . However, changes related to line 3 will not affect the value of k at 6^5 , nor will it decide whether line 6 will be executed. Therefore, line 3 should not be included into the slice and reported. By using the dummy node to separate dynamic control dependencies and potential dependencies w.r.t. statement occurrence 4^4 , we can easily exclude line 3 of Figure 9 from our relevant slice.

We now discuss how our dynamic slicing algorithm described in Section 3 (operating on compact Java bytecode traces) can be augmented to compute relevant slices. Thus, like the dynamic slicing algorithm, our relevant slicing algorithm operates on the compact bytecode traces described in Section 2.

As in dynamic slicing, relevant slicing is performed w.r.t. a slicing criterion (H, α, V) , where H is the execution history for a certain program input, α represents some bytecodes the programmer is interested in, and V is a set of variables referenced at these bytecodes. Again, the user-defined criterion is often of the form (I, l, V) where I is the input, and l is a line number of the source program. In this case, H represents execution history of the program P with input I , and α represents the bytecodes corresponding to statements at l .

Figure 11 presents a relevant slicing algorithm, which returns the relevant slice as defined in Definition 4.2.1. This algorithm is based on backward traversal of the compressed execution trace H (described in Section 2). The trace H contains execution flow and identities of accessed variables, so that we can detect various dependencies during the traversal. Although the slice can also be computed after constructing the whole EDDG, this approach is impractical because the entire EDDG may be too huge in practice. Before slicing, we compute the *control flow graph*, which is used to detect potential dependencies and to get last executed bytecode. We also pre-compute the static *control dependence graph* [Ferrante et al. 1987] which will be used at lines 22, 23 and 47 of Figure 11, and at line 9 of the `computePotentialDependence` method in Figure 12. Also, prior to running our rel-

evant slicing algorithm we run static points-to analysis [Andersen 1994; Steensgaard 1996]; the results of this analysis are used for determining potential dependencies.

In the relevant slicing algorithm, we introduce a global variable θ , to keep track of variables used by bytecode occurrences β , iff. β is included into the slice because of potential dependencies. In particular, each element in θ is of the form $\langle \beta, prop \rangle$, where β is an bytecode occurrence, and $prop$ is a set of variables. Every variable in $prop$ is used by a bytecode occurrence β' , where

- β' is included into φ (the relevant slice) because of potential dependencies, and
- $\beta' = \beta$, or β' is (transitively) dynamically control dependent on β .

The purpose of the δ set (set of unexplained variables) is the same as in the dynamic slicing algorithm. That is, the δ set includes variables used by bytecode occurrence β for explanation, where β is included into the slice φ because β belongs to the slicing criterion, or there is any bytecode occurrence in φ which is dynamically control/data dependent on β .

For each bytecode occurrence β of bytecode b_β encountered during the backward traversal for slicing, we first check if β has affected the slicing criterion via dynamic control/data dependencies as the in dynamic slicing algorithm of Figure 2. In particular, line 19 checks whether β belongs to the slicing criterion. Line 22 checks dynamic control dependencies if b_β is a conditional control transfer bytecode. Line 27 checks dynamic data dependencies. With the introduction of θ , variables in both δ and θ need explanation. Consequently, the `computeDataDependence` method has been slightly changed, as re-defined in Figure 13. If all the three checks fail, the algorithm proceeds to check the potential dependencies at line 38, by invoking the `computePotentialDependence` method in Figure 12.

For the computation of potential dependencies, we need to pre-compute the effects of various outcomes of a control transfer bytecode. Each such outcome triggers a different code fragment whose effect can be summarized by all possible assignment bytecodes executed. This summarization is used by `computePotentialDependence` method in Figure 12. Note that δ is used to check dynamic data dependencies (line 27 of Figure 11) as well as potential dependencies (line 6 in Figure 12).

The `intersect(MDS, δ)` method used by the `computePotentialDependence` method in Figure 12 checks whether the execution of the alternative path of a bytecode b_β may define some variables which are used by bytecode occurrences in the slice. Here MDS includes variables which may be defined if b_β is evaluated differently, and δ includes variables which have affected the slicing criterion and need explanation. This check is non-trivial because MDS contains static information, while δ contains dynamic information. Let $meth$ be the method that the bytecode b_β belongs to, and $curr_invo$ be the current invocation of $meth$. Note that both MDS and δ include local variables and fields. Every local variable in MDS is also a local variable of method $meth$, and is represented by its identity $\langle var \rangle$. Every local variable in δ for explanation is represented as $\langle invo, var \rangle$, where $invo$ refers to the method invocation which uses the local variable var . Many points-to analysis algorithms [Andersen 1994; Steensgaard 1996] represent abstract memory locations of objects using their possible allocation sites. Thus, we represent an object field in MDS as $\langle site, name \rangle$, where $site$ refers to a possible allocation site of the object,

```

1   $(H, \alpha, V)$  = the slicing criterion
2   $\delta = \emptyset$ , a set of variables whose values need to be explained
3   $\varphi = \emptyset$ , the set of bytecode occurrences which have affected the slicing criterion
4   $op\_stack$  = empty, the operand stack for simulation
5   $fram$  = empty, the frames of the program execution
6   $\theta = \emptyset$ , (bytecode occurrences, used variables) included in slice due to potential dependencies

7  relevantSlicing()
8       $b_\beta$  = get last executed bytecode from  $H$ ;
9      while ( $b_\beta$  is defined)
10         if ( $b_\beta$  is a return bytecode)
11              $new\_fram$  =  $createFrame()$ ;
12              $new\_fram.\gamma$  =  $\emptyset$ ;
13              $push(fram, new\_fram)$ ;
14              $last\_fram$  =  $null$ ;
15         if ( $b_\beta$  is a method invocation bytecode)
16              $last\_fram$  =  $pop(fram)$ ;
17              $\beta$  = current occurrence of bytecode  $b_\beta$ ;
18              $curr\_fram$  = the top of  $fram$ ;
19             if ( $\beta$  is the last occurrence of  $b_\beta$  and  $b_\beta \in \alpha$ )
20                  $use\_vars$  =  $V \cap$  variables used at  $\beta$ ;
21                  $\varphi = \varphi \cup \{\beta\}$ ;
22             if ( $computeControlDependence(b_\beta, curr\_fram, last\_fram)$ )
23                  $BC = \{\beta' \mid \beta' \in curr\_fram.\gamma \text{ and } \beta' \text{ is dynamically control dependent on } b_\beta\}$ ;
24                  $curr\_fram.\gamma = curr\_fram.\gamma - BC$ ;
25                  $use\_vars$  = variables used at  $\beta$ ;
26                  $\varphi = \varphi \cup \{\beta\}$ ;
27             if ( $computeDataDependence(\beta, b_\beta)$ )
28                  $def\_vars$  = variables defined at  $\beta$ ;
29                  $\delta = \delta - def\_vars$ ;
30                  $use\_vars$  = variables used at  $\beta$ ;
31                  $\varphi = \varphi \cup \{\beta\}$ ;
32                 for (each  $\langle \beta', prop' \rangle$  in  $\theta$ )
33                      $prop' = prop' - def\_vars$ ;
34             if ( $\beta \in \varphi$ )
35                  $curr\_fram.\gamma = curr\_fram.\gamma \cup \{\beta\}$ ;
36                  $\delta = \delta \cup use\_vars$ ;
37             else
38                 if ( $computePotentialDependence(\beta, b_\beta)$ )
39                      $\varphi = \varphi \cup \{\beta\}$ ;
40                     if ( $b_\beta$  is a branch bytecode)
41                          $use\_vars$  = variables used at  $\beta$ ;  $\theta = \theta \cup \{\langle \beta, use\_vars \rangle\}$ ;
42                     if ( $b_\beta$  is a method invocation bytecode)
43                          $o$  = the variable to invoke a method;  $\theta = \theta \cup \{\langle \beta, \{o\} \rangle\}$ ;
44                     if ( $b_\beta$  is a branch bytecode or method invocation bytecode)
45                          $prop = \emptyset$ ;
46                         for (each  $\langle \beta', prop' \rangle$  in  $\theta$ )
47                             if ( $\beta' = \beta$  or  $\beta'$  is control dependent on  $\beta$ )
48                                  $prop = prop \cup prop'$ ;  $\theta = \theta - \{\langle \beta', prop' \rangle\}$ ;
49                              $\theta = \theta \cup \{\langle \beta, prop \rangle\}$ ;
50                      $updateOpStack(\beta, b_\beta)$ ;
51                      $\beta = getPrevBytecode(\beta, b_\beta)$ ;
52             return bytecodes whose occurrences appear in  $\varphi$ ;

```

Fig. 11. The relevant slicing algorithm.

```

1  computePotentialDependence ( $\beta$ : bytecode occurrence,  $b_\beta$ : bytecode)
2      if ( $b_\beta$  is a branch or method invocation bytecode )
3          for (each possible outcome  $x$  of  $b_\beta$ )
4              if (outcome  $x$  of  $b_\beta$  did not occur at  $\beta$  )
5                   $MDS$  = the set of variables which may be defined
                        when outcome  $x$  occurs;
6                  if ( $intersect(MDS, \delta)$ )
7                      return true;
8                  for (each  $\langle \beta', prop' \rangle$  in  $\theta$ )
9                      if ( $\beta'$  is not control dependent on  $\beta$  and  $intersect(MDS, prop')$  )
10                         return true;
11 return false;

```

Fig. 12. Detect potential dependencies for relevant slicing.

```

1  computeDataDependence ( $\beta$ : bytecode occurrence,  $b_\beta$ : bytecode)
2      if ( $\beta$  defines a variable)
3          if ( $\beta$  defines a static field or local variable)
4               $def\_loc$  = get address of the defined static field or local variable from class files;
5          if ( $\beta$  defines an object field or an array element)
6               $G$  = compressed operand sequence for  $b_\beta$  in the compact bytecode trace  $H$ 
7               $\pi$  = a root-to-leaf path for  $G$ ;
8               $def\_loc$  =  $getLast(G, \pi)$ ;
9          if ( $def\_loc \in \delta$  )
10             return true;
11         for (each  $\langle \beta', prop' \rangle$  in  $\theta$ )
12             if ( $def\_loc \in prop'$ )
13                 return true;
14      $\omega$  = the set of bytecode occurrences in top  $def\_op(b_\beta)$  entries of  $op\_stack$ ;
15     if ( $\omega \cap \varphi \neq \emptyset$ )
16         return true;
17 return false;

```

Fig. 13. Detect dynamic data dependencies for relevant slicing.

and $name$ refers to the name of the field. We also represent an object field in δ as $\langle site, timestamp, name \rangle$, where $site$ refers to the allocation site of the object, $timestamp$ distinguishes between objects created at the same allocation site, and $name$ refers to field name. Note that $timestamp$ is only important for detecting dynamic data dependencies. The $intersect(MDS, \delta)$ method returns true iff:

- there is a local variable where $\langle var \rangle \in MDS$ and $\langle curr_invo, var \rangle \in \delta$, or
- there is a field where $\langle site, name \rangle \in MDS$ and $\langle site, timestamp, name \rangle \in \delta$.

We do not consider partial results (*i.e.* operands in the operand stack) for potential dependencies because partial results will not be transferred between bytecodes of different statements.

The proof of correctness of the relevant slicing algorithm in Figure 11 is similar to that of the dynamic slicing algorithm in Figure 2. *The detailed proof of correctness can be found in our technical report [Wang and Roychoudhury 2007].*

Now we analyze the cost of the relevant slicing algorithm in Figure 11. The space overheads of the slicing algorithm are $O(N^2 + m^3)$, and the time overheads are $O(m^2 \cdot N^3)$, where N is the length of the execution, and m is the number of bytecodes of the program. Since the relevant slicing algorithm is similar with the dynamic slicing algorithm in Figure 2, the cost analysis is also similar except costs w.r.t. the θ and MDS .

The θ set contains at most N elements of the form $\langle \beta, prop \rangle$, because every bytecode occurrence β has at most one element $\langle \beta, prop \rangle$ in θ . The size of each $prop$ is $O(N)$, because at most N bytecode occurrences are (transitively) dynamically control dependent on β and every bytecode occurrence uses a constant number of variables. Consequently, the space overheads of θ are $O(N^2)$.

Each MDS includes the set of variables which may be defined when a specific outcome of a branch bytecode occurs. If m is the number of bytecodes in the program, clearly there are at most m MDS s. How do we bound the size of each MDS ? Each MDS may have at most m assignments and each of these assignments may affect at most m locations (provided we distinguish locations based on allocation sites as is common in points-to analysis methods). Thus, the size of each MDS is $O(m^2)$. Since there are at most m MDS s, the space overheads of maintaining the MDS s are $O(m^3)$. The other portions of the relevant slicing algorithm are taken from dynamic slicing; the space overheads of these portions of the relevant slicing algorithm are $O(N)$, as explained in Section 3. So, the overall space overheads of our relevant slicing algorithm are $O(N^2 + m^3)$.

We now calculate the time overheads of relevant slicing. First we estimate the time overheads for maintaining the θ set in the relevant slicing algorithm. Note that the θ set contains $O(N)$ elements of the form $\langle \beta, prop \rangle$. The set difference operation at line 33 of Figure 11 is executed $O(N^2)$ times. Since each $prop$ set contains $O(N)$ variables, the overall time overheads to perform the set difference operation at line 33 of Figure 11 are $O(N^3)$. The total time overheads to perform the set union operation at lines 41, 43 and 49 of Figure 11 are $O(N^2)$, because lines 41, 43 and 49 are executed $O(N)$ times and the size of the θ set is $O(N)$. Given a bytecode occurrence β and the θ set, there are a constant number of elements $\langle \beta', prop' \rangle \in \theta$, where β' is (directly) dynamically control dependent on β . This is because there is no explicit `goto` statement in Java programs. Different occurrences of the same bytecode are dynamically control dependent on different bytecode occurrences. Note that there are a constant number of bytecodes which are statically control dependent on the bytecode b_β of occurrence β . Thus, there are a constant number of bytecode occurrences which are *directly* dynamically control dependent on β . In addition, every bytecode occurrence β' has at most one $\langle \beta', prop' \rangle$ in θ . Consequently, line 48 is executed $O(N)$ times, and the overall time overheads to execute line 48 are $O(N^3)$. The total time overheads to perform check at line 12 of Figure 13 are $O(N^3)$, which is similar to perform set difference operation at line 33 of Figure 11.

Now, we analyze the overall time overheads to perform the `intersect` operation by the `computePotentialDependence` method in Figure 12. The `intersect` operation at line 6 of Figure 12 can be executed at most N times. In each execution of the `intersect` operation we compare the contents of MDS and δ . Since the size of MDS is $O(m^2)$ and the size of the set δ is $O(N)$, therefore the time overheads

of a single `intersect` operation are $O(N \cdot m^2)$. Thus, the total time overheads to execute all the `intersect` operations at line 6 of Figure 12 are $O(N \cdot N \cdot m^2)$. Similarly, the total time overheads to execute all the `intersect` operations at line 9 of Figure 12 are $O(N^3 \cdot m^2)$, since this `intersect` operation is executed $O(N^2)$ times.

The other portions of the relevant slicing algorithm are taken from dynamic slicing; the time complexity of these portions of the relevant slicing algorithm is $O(N^2)$, as discussed in Section 3. This leads to a time complexity of $O(N^3 \cdot m^2)$ for our relevant slicing algorithm.

5. EXPERIMENTAL EVALUATION

We have implemented a prototype slicing tool and applied it to several subject programs. The experiments report time and space efficiency of our trace collection technique. We compress traces using both SEQUITUR and RLESe, and compare the time overheads and effectiveness of the improved compression algorithm against the previous one, in order to investigate the cost effectiveness of the proposed compression algorithm. In addition, the experiments report time overheads to compute dynamic and relevant slices, and the sizes of these slices.

To collect execution traces, we have modified the Kaffe virtual machine [Kaffe] to monitor interpreted bytecodes. In our traces, we use object identities instead of addresses to represent objects. Creation of structures such as multi-dimensional arrays, constant strings etc. may implicitly create objects. We trace and allocate identities to these objects as well. The virtual machine may invoke some methods automatically when “special” events occur (*e.g.* it may invoke the static initializer of a class automatically when a static field of the class is first accessed). These event are also stored, and used for backward traversal of the execution trace.

Most Java programs use libraries. Dynamic dependencies introduced by the execution of library methods are often necessary to compute correct slices. Our implementation does not distinguish between library methods and non-library methods during tracing and slicing. However, after the slicing finishes, we will filter statements inside the libraries from the slice. This is because libraries are often provided by other vendors, and programmers will not look into them.

Finally, note that static points-to analysis results are required in our relevant slicing algorithm (for computing potential dependencies). To compute points-to sets, we used the *spark* toolkit [Lhoták 2002] which is integrated in the compiler optimization framework *soot* [Vallée-Rai et al. 1999].

5.1 Subject Programs

The subjects used in our experiments include six subjects from the Java Grande Forum benchmark suite [JGF], three subjects from the SPECjvm suite [SPECjvm98 1998], and one medium sized Java utility program [Berk and Ananian 2003]. Descriptions and inputs of these subjects are shown in Table III. We ran and collected execution traces of each program with inputs shown in the third column in Table III. Corresponding execution characteristics are shown in Table IV. The second column in Table IV shows the number of bytecodes of these subjects. The column *Executed bytecodes* in Table IV presents the number of distinct bytecodes executed during one execution, and the fourth column in Table IV shows corresponding total num-

Subject	Description	Input
Crypt	IDEA encryption and decryption	200,000 bytes
SOR	Successive over-relaxation on a grid	100×100 grid
FFT	1-D fast Fourier transform	2^{15} complex numbers
HeapSort	Integer sorting	10000 integers
LUFact	LU factorisation	200×200 matrix
Series	Fourier coefficient analysis	200 Fourier coefficients
._201_compress	Modified Lempel-Ziv method (LZW)	228.tar in the SPECjvm suite
._202_jess	Java Expert Shell System	fullmab.clp in the SPECjvm suite
._209_db	Performs multiple database functions	db2 & scr2 in the SPECjvm suite
JLex	A Lexical Analyzer Generator for Java	sample.lex from the tool's web site

Table III. Descriptions and input sizes of subject programs.

Subject	Total # of Bytecodes	Executed Bytecodes	# Bytecode Instances	# Branch Instances	# Method invocations
Crypt	2,939	1,828	103,708,780	1,700,544	48
SOR	1,656	740	59,283,663	1,990,324	26
FFT	2,327	1,216	72,602,818	2,097,204	37
HeapSort	1,682	679	11,627,522	743,677	15,025
LUFact	2,885	1,520	98,273,627	6,146,024	41,236
Series	1,800	795	16,367,637	1,196,656	399,425
._201_compress	8,797	5,764	166,537,472	7,474,589	1,999,317
._202_jess	34,019	14,845	5,162,548	375,128	171,251
._209_db	8,794	4,948	38,122,955	3,624,673	19,432
Jlex	22,077	14,737	13,083,864	1,343,372	180,317

Table IV. Execution characteristics of subject programs.

Subject	Orig. Trace	Trace Table	RLESe Sequences	All	All/Orig. (%)
Crypt	64.0M	8.9k	8.8k	17.8k	0.03
SOR	73.4M	7.6k	10.8k	18.5k	0.02
FFT	75.2M	8.2k	87.3k	95.5k	0.12
HeapSort	23.6M	7.7k	1.7M	1.7M	7.20
LUFact	113.1M	9.1k	179.7k	188.9k	0.16
Series	24.4M	7.7k	444.4k	452.2k	1.81
._201_compress	288.6M	23.7k	8.8M	8.8M	3.05
._202_jess	25.7M	79.2k	3.4M	3.4M	13.23
._209_db	194.5M	29.0k	39.2M	39.2M	20.15
Jlex	49.9M	62.6k	1.5M	1.5M	3.01

Table V. Compression efficiency of our bytecode traces. All sizes are in bytes.

ber of bytecode occurrences executed. The last two columns present the number of branch bytecode occurrences and method invocations, respectively. Bytecodes from Java libraries are not counted in the table. However, bytecodes of user methods called from library methods are included in the counts shown.




Fig. 14. Time overheads of RLESe and SEQUITUR. The time unit is *second*.

5.2 Time and Space Efficiency of Trace Collection

We first study the efficiency of our compact trace representation. Table V shows the compression efficiency of our compact trace representation. The column *Orig. Trace* represents the space overheads of storing uncompressed execution traces on disk. To accurately measure the compression achieved by our RLESe grammars, we *leave out* the untraced bytecodes from the *Orig. Trace* as well as the compressed trace. Next two columns *Trace Table* and *RLESe Sequences* show the space overheads to maintain the trace tables and to store the compressed operand sequences of bytecodes on disk. The column *All* represents the overall space costs of our compact bytecode traces, *i.e.* sum of space overheads of *Trace Table* and *RLESe Sequences*. The % column indicates *All* as a percentage of *Orig. Trace*. For all our subject programs, we can achieve at least 5 times compression. Indeed for some programs we get more than two orders of magnitude (*i.e.* 100 times) compression.

Figure 14 presents the absolute running time of the subject programs without instrumentation, tracing with RLESe and tracing with SEQUITUR. All experiments were performed on a Pentium 4 3.0 GHz machine with 1 GB of memory. From this figure we can see that the slowdown for collecting traces using RLESe compared to the original program execution is 2 – 10 times for most subject programs, and the slowdown is near 20 times for *_209.db* which randomly accesses a database. This shows that our method is suitable for program executions with many repetitions, but the time overhead may be not scalable to program executions with too many random accesses. Note that these additional time overheads are caused by compression using RLESe/SEQUITUR. The results reflect one limitation of RLESe (as well as SEQUITUR) — the time efficiency of the compression algorithm heavily depends on the efficiency of checking the digram uniqueness property. To speed up this check, a sophisticated index of digrams should be used. However, the choice of the index has to trade-off between time and space efficiency, since the index will introduce additional space overheads during trace collection. In our implementation, we use the B+ tree (instead of sparse hash) to index digrams during compressing, and the index is no longer needed after compression is completed.

Subject	RLESe %	SEQUITUR %
Crypt	0.03	0.07
SOR	0.02	0.04
FFT	0.12	0.34
HeapSort	7.20	7.20
LUFact	0.16	0.40
Series	1.81	2.50
_201_compress	3.05	3.12
_202_jess	13.23	13.62
_209_db	20.15	18.35
Jlex	3.01	4.01

Table VI. Comparing compression ratio of RLESe and SEQUITUR.

We now compare the space efficiency of RLESe against SEQUITUR. Both algorithms were performed on operand sequences of bytecodes. For a fair comparison, we use the same index to search for identical/similar digrams in the implementation of both algorithms. Table VI compares the space costs of both algorithms by presenting their compression ratio (in percentage). From the tables, we can see that the space consumption of compressed traces produced by both algorithms are somewhat comparable. RLESe outperforms SEQUITUR for eight subject programs, and SEQUITUR outperforms for one (where the `_209_db` program randomly accesses a database and does not have many contiguous repeated symbols in operand traces). Since RLESe employs run-length encoding of terminal and non-terminal symbols over and above the SEQUITUR algorithm, nodes in grammars produced by RLESe are usually less than those produced by SEQUITUR. However, each grammar node of RLESe has an extra field to record the run-length. This will lead to higher memory consumption when RLESe cannot reduce enough nodes (*e.g.* compressed traces for `_209_db`).

Figure 14 compares the time overheads of RLESe and SEQUITUR. Clearly RLESe outperforms SEQUITUR in time on studied programs. The time overheads of both algorithms are mainly caused by checking digram uniqueness property. RLESe usually produces less nodes in grammars, so that similar digrams can be found more efficiently. In addition, RLESe checks this property after contiguous repeated symbols have finished, whereas SEQUITUR does this on reading every symbol. Table VII shows this by representing the frequency of checking the digram uniqueness property. When there are many contiguous repeated symbols in the execution traces (*e.g.* the `LUFact`, `SOR` subjects), RLESe checks the property much less frequently than SEQUITUR. Consequently the tracing overheads of RLESe are also much less than those of SEQUITUR.

5.3 Overheads for Computing Dynamic Slice and Relevant Slice

We now measure the sizes of dynamic and relevant slices and the time overheads to compute these slices. Thus, apart from studying the absolute time/space overheads of dynamic slicing, these experiments also serve as comparison between dynamic and relevant slicing. For each subject from the Java Grande Forum benchmark suite, we randomly chose five distinct slicing criteria; for other bigger subjects, we randomly chose fifteen distinct slicing criteria.

Subject	RLESe	SEQUITUR
Crypt	800,605	20,499,488
SOR	393,202	25,563,918
FFT	2,020,976	25,722,054
HeapSort	2,262,916	7,066,784
LUFact	561,233	42,586,136
Series	2,501,001	9,507,982
_201_compress	10,240,652	80,923,387
_202_jess	2,090,153	6,666,460
_209_db	24,386,746	54,033,314
Jlex	4,555,338	15,497,211

Table VII. The number of times to check digram uniqueness property by RLESe and SEQUITUR.

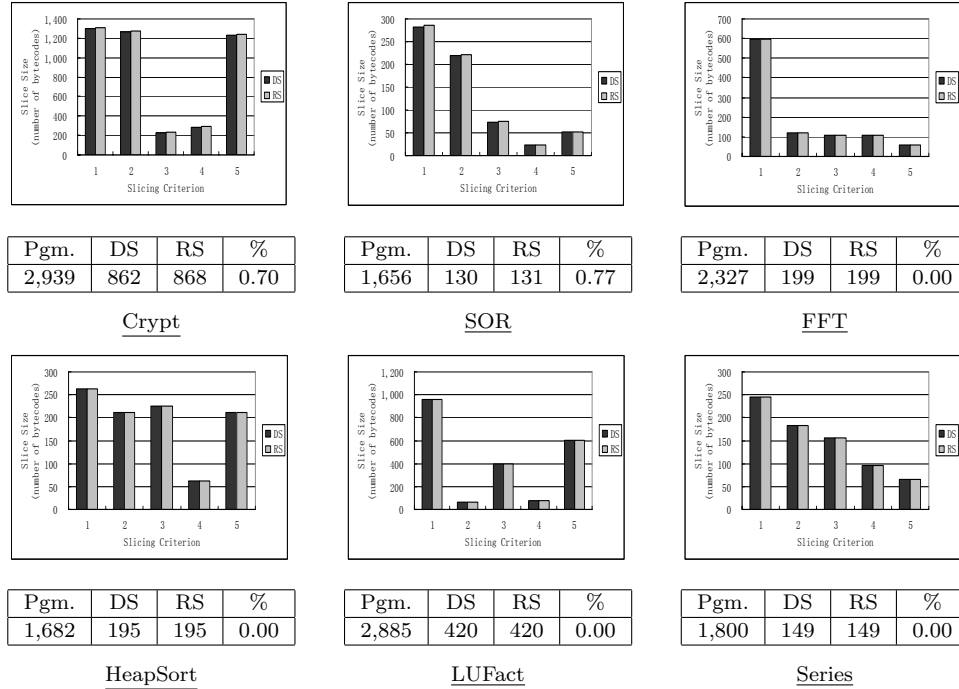


Fig. 15. Compare sizes of relevant slices with dynamic slices. *Pgm.* represents size of the program. *DS* and *RS* represent average sizes of dynamic slices and relevant slices, respectively. All sizes are reported as the number of bytecodes. The last % column represents the increased sizes of relevant slices (*i.e.* $\frac{RS-DS}{DS}$) in percentage.

Dynamic slice and relevant slice. Figure 15 and 16 present the slice sizes. As we can see, most dynamic slices and relevant slices were relatively small, less than 30% of corresponding source programs on average in our experiments. This is because programs often consist of several parts which work (almost) independently. The irrelevant portions are excluded from both dynamic slices and relevant slices. In our experiments, control and data structures of subject programs from Java Grande

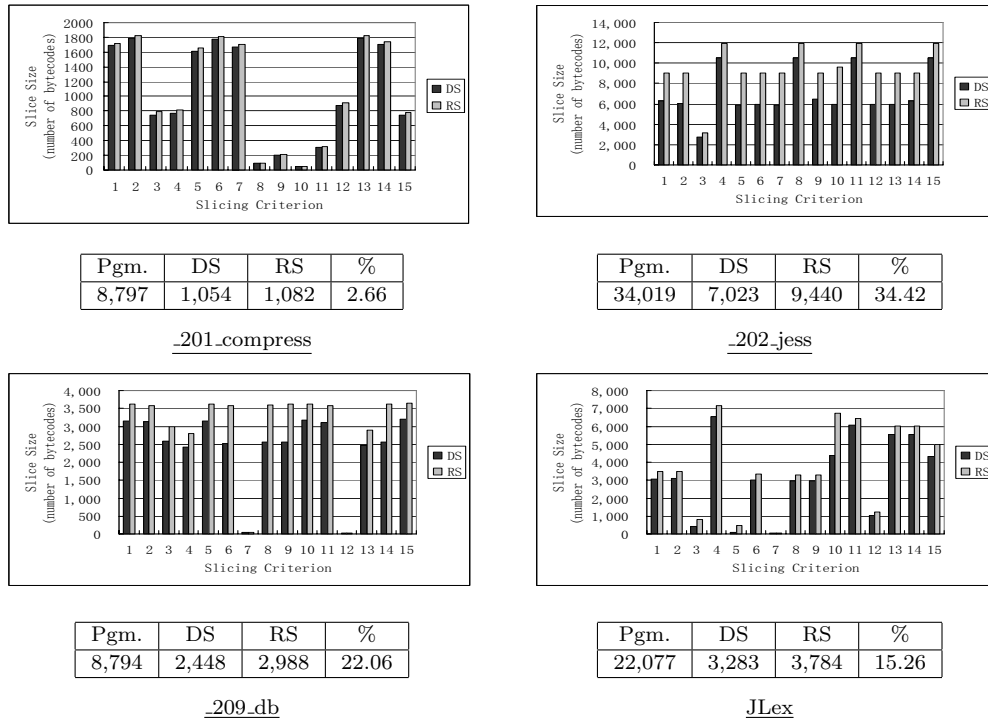


Fig. 16. Compare sizes of relevant slices with dynamic slices. *Pgm.* represents size of the program. *DS* and *RS* represent average sizes of dynamic slices and relevant slices, respectively. All sizes are reported as the number of bytecodes. The last % column represents the increased sizes of relevant slices (*i.e.* $\frac{RS-DS}{DS}$) in percentage.

Forum benchmark suite are quite simple. Furthermore, class hierarchies of these programs are simple, with only limited use of inheritance. The relevant slices are almost the same as corresponding dynamic slices for these programs, as shown in Figure 15.

On the other hand, other subject programs (*i.e.* _201_compress, _202_jess, _209_db, and JLex) are more complex. These programs have more sophisticated control structures. In addition, these programs use inheritance and method overloading and overriding, so that different methods may be called by the same method invocation bytecode. Both of the factors lead to potential dependencies between bytecode occurrences. More importantly, these subject programs involve substantial use of Java libraries, such as collection classes and I/O classes. These library classes contribute a lot to relevant slices, because of their complex control structures and usage of object-oriented features like method overloading (which lead to potential dependencies). In fact, if we do not consider potential dependencies inside such library classes, the average sizes of relevant slices for _201_compress, _202_jess, _209_db, and JLex were 1072, 8393, 2523, and 3728, which were 1.71%, 19.51%, 3.06%, 13.55% bigger than corresponding dynamic slices, respectively.




Fig. 17. Compare time overheads of relevant slicing with dynamic slicing. *DS* and *RS* represent average time to perform dynamic slicing and relevant slicing, respectively. All time are reported in second. The last % column represents the increased time for relevant slicing (i.e. $\frac{RS-DS}{DS}$) in percentage.

Time overheads. Figure 17 and 18 show the time overheads to compute dynamic slices and relevant slices. Clearly, the time overheads to compute dynamic slices and relevant slices are sensitive to choice of programs and slicing criteria. According to our slicing algorithms in Figure 2 and 11, the time overheads mainly come from: (1) extracting operand sequences of bytecodes and backwards traversal, and (2) updating/comparing various sets to detect dependencies. For a particular execution trace, the first task is common for slicing w.r.t. every slicing criterion, and the second task is sensitive to the sizes of the slices. For bigger slices, their sets to detect dependencies are also bigger during slicing, resulting in larger time overheads.

5.4 Summary and Threats to Validity

In summary, the RLESe compaction scheme achieves comparable or better compression ration than SEQUITUR. The time overheads for the online compaction in RLESe is found to be less than the compaction time in SEQUITUR. Our compact trace representation can be used to perform dynamic slicing and relevant slicing

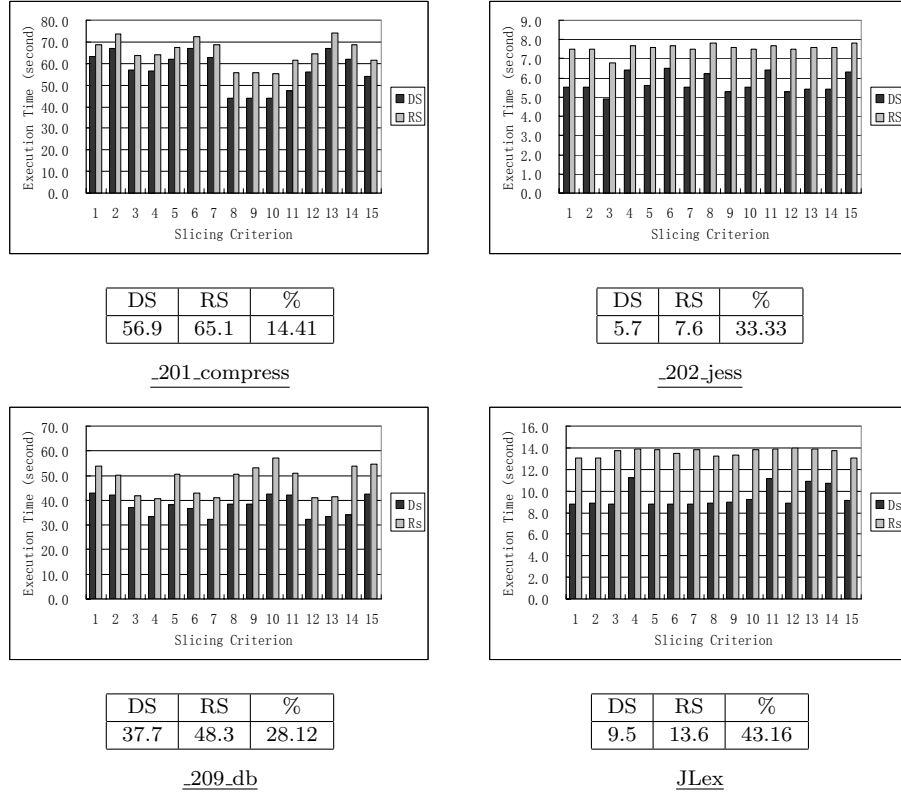


Fig. 18. Compare time overheads of relevant slicing with dynamic slicing. *DS* and *RS* represent average time to perform dynamic slicing and relevant slicing, respectively. All time are reported in second. The last % column represents the increased time for relevant slicing (i.e. $\frac{RS-DS}{DS}$) in percentage.

efficiently, both in time and space. The resultant dynamic slices are relatively small, compared against the entire program. They can guide further analysis (e.g. program debugging) working on a small but important portion of the program. Relevant slices are often bigger than dynamic slices (since they consider *potential dependencies*), but relevant slices are still relatively small compared against the entire program. Of course, relevant slicing for a given criterion takes more time than dynamic slicing for the same criterion. However, the increased time overheads, which depend on the choice of slicing criterion, the control structures and data flow structures of the program, were tolerable in our experiments.

We note that there are various threats to validity of the conclusion from our experiments. Our conclusions on less time overheads of RLESe (as compared to SEQUITUR) can be invalidated if the execution trace has very few contiguous repeated symbols. In this case, RLESe checks the diagram uniqueness property almost as frequently as SEQUITUR. This can happen in a program with random data accesses (e.g., the _209_db subject program used in our experiments).

Our conclusions on the sizes of dynamic and relevant slices can be invalidated for programs with little inherent parallelism, and having long chains of control/data dependencies. For such programs, the dynamic and the relevant slices may not be substantially smaller than the program. Furthermore, as the slice sizes increase so do the time to compute these slices, since additional statements in the slice will result in more unexplained variables in the δ set. It is also possible to encounter a situation where the dynamic slice is quite small compared to the program, but the relevant slice is much bigger than the dynamic slice. Since relevant slice computation involves detecting potential dependencies and potential dependencies involve computing static data dependencies, a lot depends on the accuracy of the points-to analysis used to detect static data dependencies. If the points-to analysis is very conservative, it may lead to a large relevant slice.

6. JSlice: A DYNAMIC SLICING TOOL FOR JAVA PROGRAMS

We have made our Java dynamic slicing tool available for use by researchers and developers. The tool is integrated with the Eclipse platform¹. Programmers can specify the dynamic slicing criterion, perform dynamic slicing, and inspect the dynamic slice in the Integrated Development Environment (IDE) of Eclipse. Note that programmers also develop and debug Java programs in this IDE. They can use this dynamic slicing tool easily during their development. We will briefly introduce the dynamic slicing tool in the following. More details are available in the tool's web site given in the following.

<http://jslice.sourceforge.net/>

Figure 19 presents the architecture of our *JSlice* dynamic slicing tool. The tool consists of two parts: (a) a front end, which is the user interface, and (b) a back end, which collects traces and performs dynamic slicing.

The front end is implemented as a Eclipse plug-in. It helps programmers to specify the dynamic slicing criterion, and read the dynamic slice. For example, programmers can specify the slicing criterion by selecting a line in the source code editor. All variables referenced at the selected line will be included into the slicing criterion. Figure 20 demonstrates how statements included in the dynamic slice are highlighted in the source code editor, so that programmers can easily find and inspect them.

The back end is integrated with the open-source Kaffe Virtual Machine. When Kaffe is executing a Java program, we instrument the interpreter of the virtual machine, and collect the compact trace at the level of bytecode. After the Java program terminates, we perform dynamic slicing on the compact trace without decompression, w.r.t. the slicing criterion specified via the front end. The resultant slice is then transformed to the source code level with the help of information available in Java class files.

The dynamic slicing tool supports various ways to specify the slicing criterion. When a programmer selects a line in the source code editor as the slicing criterion in the tool, he/she can specify whether the last occurrence or all occurrences of

¹Details about Eclipse can be found at <http://www.eclipse.org/>.

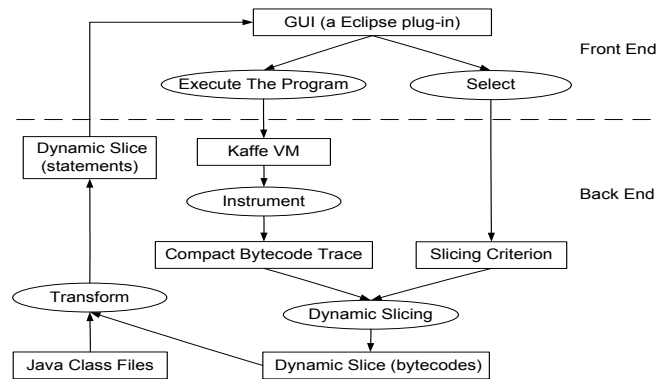


Fig. 19. The architecture of the JSlice tool.

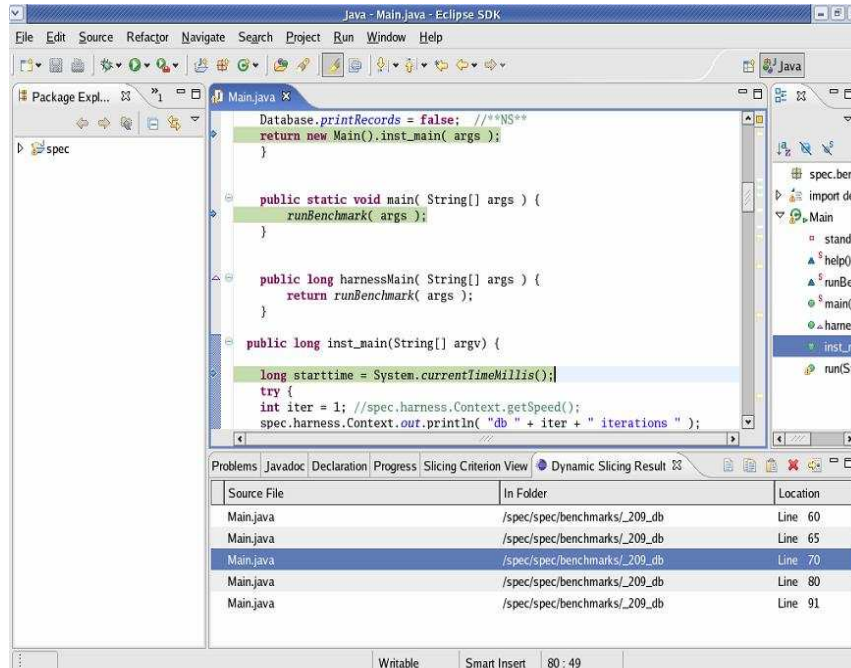


Fig. 20. Highlighted statements capture the dynamic slice computed (by the JSlice tool).

this statement should be considered as the criterion. When a programmer is interested in variables referenced at a specific occurrence of a statement in the middle of the execution, he/she can specify the statement as the slicing criterion, and terminate the program’s execution as soon as the “interesting” statement occurrence appears. The tool will then return the slice w.r.t. variables referenced at the “interesting” statement occurrence. The next example shows how our dynamic slicing tool supports this feature. Suppose we have the following program fragment,

```

1.  for ( $i = 0; i < 10; i++$ ){
2.       $j = k + f(i)$ ;
3.  }
```

and we are interested in how the value of variable j in the 6th loop iteration is computed. We can specify statement 2 as the dynamic slicing criterion, and add two lines (*i.e.* line 3 and 4) into the program as shown in the following.

```

1.  for ( $i = 0; i < 10; i++$ ){
2.       $j = k + f(i)$ ;
3.      if ( $i \geq 5$ )
4.          java.lang.System.exit(1);
5.  }
```

Our tool will then return the dynamic slice w.r.t. variables referenced at the 6th occurrence of statement 2.

7. ONGOING EXTENSIONS OF OUR SLICING TOOL

The dynamic slicing tool JSlice supports most features of the Java programming languages, such as object, field, inheritance, polymorphism and etc. In future, we plan to enhance JSlice to support more features of the Java programming language. In particular, exceptions, reflection and multi-threading are widely used features of the Java programming language. We plan to extend our dynamic slicing tool to handle these features in the following manner.

7.1 Exceptions

When a program violates any semantic constraint of the Java programming language, the Java virtual machine throws an exception to signal this error [Joy et al. 2000]. This exception will cause a non-local transfer of control from the point where the exception occurred to the exception handler which can be specified by the programmer. It is necessary to store this non-local transfer of control during trace collection, so that we can reconstruct such a control transfer during the backward traversal for dynamic slicing.

JSlice maintains traces of each bytecode separately in the trace tables for the program, as discussed in Section 2. Thus, the non-local transfer of control should be stored in the traces of the first bytecode of the exception handler. Note that this control transfer will cause the Java virtual machine to change the call stack, in the process of looking for an appropriate exception handler. In particular, the virtual machine will pop method invocations from the call stack up to the method invocation *invo_except* (which the exception handler belongs to), and then execute the exception handler. For each invocation *invo* which is popped from or revised in the call stack, we need to record the following (assume that *invo* is an invocation of method *meth*).

- the class name of *meth*, and
- the method name of *meth*, and
- the signature of *meth*, and
- the id/address of last executed bytecode of *invo*, and
- the size of the operand stack of *invo* before *invo* is popped or revised.

When the first bytecode b of an exception handler is encountered during the backward traversal for slicing, the dynamic slicing algorithm should retrieve information from the traces of b , and reconstruct the call stack, so that the backward traversal can continue.

Exception handling also introduces extra dependencies into the program: the dynamic control dependence between (a) the bytecode occurrence which throws the exception and (b) the exception handler which catches the exception [Sinha and Harrold 2000]. This means that, when any bytecode instance in the exception handler is included into the dynamic slice, the bytecode occurrence which throws the exception should also be included into the slice.

Finally block. For Java programs, exception handlers often come with a `finally` block, where the Java Virtual Machine ensures that the finally block is always executed even if an unexpected exception occurs. However, the usage of the finally block complicates the construction of the control flow graph of a Java method, as discussed in the following. During the execution of a Java method, a finally block is always entered by executing a `JSR` bytecode. The semantics of the `JSR` bytecode is very similar to that of the `goto` bytecode. However, when a `JSR` bytecode b is executed, the address of the bytecode b' which immediately follows b is stored into the operand stack. When the finally block finishes execution, the saved address of the bytecode b' is retrieved and the execution continues from the bytecode b' . In other words, the bytecode b' , which is executed after the finally block, is not represented as an operand in the last bytecode of the finally block. As a result, it is not clear which bytecodes may be executed after a finally block. In order to discover this kind of information, the algorithm in Figure 21 is used. Given a method $meth$, the algorithm in Figure 21 returns an array $succ$, where for every bytecode $exit$ which is the last bytecode of a finally block, $succ[exit]$ represents the set of bytecodes which may be executed after the bytecode $exit$. The algorithm proceeds by traversing the bytecode sequence of the method $meth$ twice. During the first traversal, we mark the entry bytecode $entry$ of each finally block, and maintain $next[entry]$, the set of bytecodes which may be executed after the finally block. During the second traversal, for every $entry$ bytecode of a finally block, we detect the corresponding $exit$ bytecode which exits the finally block. Additionally, we set $succ[exit]$ to $next[entry]$, so that we can get the control flow information w.r.t. these $exit$ bytecodes and construct the control flow graph as usual. The stack $entryStack$ is required here because of nested finally blocks.

7.2 Reflection

Reflection gives the Java code access to internal information of classes in the Java Virtual Machine, and allows the code to work with classes selected during execution, not in the source code. The main difficulty to support reflection for slicing lies in the fact that many reflection methods are implemented as native methods, and JSlice cannot trace details of native methods. Of all the native reflection methods, the following two kinds of methods are particularly important for dynamic slicing.

—*Native methods which invoke a Java method* e.g. `java.lang.reflect.Method.invoke`
 Clearly, there exists a control transfer from the native method to the callee Java method. This control transfer is important for dynamic slicing, since we need

```

1  findNext (meth: a Java method)
2      initialize each element of the array next and succ to  $\emptyset$ ;
3      initialize entryStack to null;
4      for (each bytecode b of the method meth)
5          if (the bytecode b is a JSR bytecode)
6              entry = the operand bytecode of the JSR bytecode b;
7              mark the bytecode entry as an entry of a finally block;
8              b' = the bytecode which immediately follows the JSR bytecode b;
9              next[entry] = next[entry]  $\cup$  {b'};
10     for (each bytecode b of the method meth)
11         if (the bytecode b is an entry of a finally block)
12             push(entryStack, b);
13         if (the bytecode b is the last bytecode of a finally block)
14             entry = pop(entryStack);
15             exit = b;
16             succ[exit] = next[entry];
17     return succ;

```

Fig. 21. The algorithm to find the bytecodes which may be executed after each **finally** block.

to traverse the callee Java method for dynamic data dependence analysis. Here, we have to explicitly record the control transfer. The class name, method name, and signature of the callee method should be recorded.

—*Native methods which read/write fields or arrays*, where we can deduce which variables are accessed according to the parameters and the invoking objects. For example, field access methods in the `java.lang.reflect.Field` class fall into this category. These native methods are also essential for dynamic data dependence analysis. Note that these methods behave similarly with field/array access bytecodes, we can trace and analyze these methods in a similar way as corresponding bytecodes. That is, we trace the address (or identity) corresponding to the object/array, and the field name or the index of the array element. During dynamic slicing, such information is retrieved to detect dynamic data dependencies.

7.3 Multi-threading

We plan to extend JSlice to support multi-threaded Java programs. The trace representation for a multi-threaded Java program could be similar to that for a single-threaded Java program. That is, each method has one trace table, and each row of the table maintains the control and data flow traces of a specific bytecode (see Section 2 for the trace representation). However, Java threads often communicate with each other through inter-thread events, such as shared variable access events, wait/notify events. The order of these events is required for dynamic slicing, because such an order is essential to reason/detect inter-thread dynamic dependencies. Levrouw et al. have proposed an efficient mechanism which can be used to trace the order of these inter-thread events [Levrouw et al. 1994]. We now briefly describe this approach in the following.

Levrouw’s approach is based on the Lamport Clocks [Lamport 1997]. During the execution, each thread t_i has a scalar clock c_t^i , and each object o also maintains a clock c_o . These clocks are initialized to 0. Whenever there is an inter-thread

event e where the thread t_i accesses the object o , this event is recorded with a time stamp $c_e = \max(c_t^i, c_o) + 1$. The function \max returns the maximum value of the two inputs. Additionally, c_t^i and c_o are updated to c_e . These recorded time stamps actually impose a partial order on all inter-thread events. Levrouw et al. show that we can replay the original execution and re-construct the dynamic dependencies, by enforcing inter-thread events following the partial order.

There is one practical problem in employing the above scheme for tracing multi-threaded Java programs — all objects can be accessed by different threads, and it is often non-trivial to know which objects are shared before the execution. As a result, every access to an object should be considered as an inter-thread event. However, if we trace the time stamp for every object access, the trace size may explode. Fortunately, Levrouw et al. show that it is not necessary to trace all time stamps to record the partial order. In particular, for an inter-thread event e where the thread t_i accesses the object o , let c_t^i be the time stamp of t_i , and c_o of be the time stamp of o . We only need to trace the increment of c_t^i before and after the event e , if $c_t^i < c_o$. The reader is referred to [Ronsse and Bosschere 1999; Levrouw et al. 1994] for details. Note that the tracing scheme given here will work for multi-threaded Java programs running on multi-processor platforms as well.

The dynamic slicing algorithm for multi-threaded programs is similar to that for single-threaded programs (see Section 3). However, the algorithm should now maintain several operand stacks and call stacks, each of which corresponds to one thread. At any specific time, only one operand stack and one call stack are active. When we encounter an inter-thread event during the backward traversal, we pause the traversal along this thread until we have traversed all inter-thread events with a bigger time stamp. In addition, besides dynamic control and data dependencies, the slicing algorithm should also consider inter-thread dependencies, such as the dependencies introduced by wait-notify operations.

8. RELATED WORK

In this section, we survey related literature on compressed program trace representations, dynamic slicing and extensions of dynamic slicing for detecting omission errors.

8.1 Work on Dynamic Slicing

Weiser originally introduced the concept of program slicing [Weiser 1984]. In the last two decades, program slicing, in particular dynamic slicing, has been widely used in many software engineering activities, such as program understanding, debugging and testing [Agrawal et al. 1993; Korel and Rilling 1997; Lucia 2001; Tip 1995]. The first dynamic slicing algorithm was introduced by Korel and Laski [Korel and Laski 1988]. In particular, they exploited *dynamic flow concepts* to capture the dependencies between occurrences of statements in the execution trace, and generated executable dynamic slices. Later, Agrawal and Horgan used the *dynamic dependence graph* to compute non-executable but precise dynamic slices, by representing each occurrence of a statement as a distinct node in the graph [Agrawal and Horgan 1990; Agrawal 1991]. A survey of program slicing techniques developed in the eighties and early nineties appears in [Tip 1995].

Static and dynamic slicing of object-oriented programs based on dependence graphs have been studied in [Larsen and Harrold 1996] and [Xu et al. 2002] respectively. Computation of control and data dependencies between Java bytecodes has been discussed in [Zhao 2000]. The work of [Ohata et al. 2001] combined dynamic and static slicing to avoid the space overheads of processing traces, at the cost of precision of computed slices. Recently, Zhang et al. [Zhang et al. 2005] studied the performance issues in computing the control/data dependencies from the execution traces for slicing purposes.

In [Dhamdhere et al. 2003], Dhamdhere et al. present an approach for dynamic slicing on compact execution traces. However, they do not employ any data compression algorithm on the execution trace. They classify execution instances of statements as critical or non-critical, and store only the latest execution instances for non-critical statements. The classification of statements as critical/non-critical is sensitive to the slicing criterion. In contrast, our compression scheme is not related to the slicing criterion and exploits regularity/repetition of control/data flow in the trace. Our slicing algorithm operates directly on this compressed trace achieving substantial space savings at tolerable time overheads.

Extensions of Dynamic Slicing for Detecting Omission Errors. In the past, *relevant slicing* has been studied as an extension of dynamic slicing for the purpose of detecting omission errors in a program [Agrawal et al. 1993; Gyimóthy et al. 1999]. The approach in [Agrawal et al. 1993] relies on the huge dynamic dependence graph. Furthermore, if b is a branch statement with which statements in the slice have potential dependencies, [Agrawal et al. 1993] only computes the closure of data and potential dependencies of b . In other words, control dependencies are ignored w.r.t. statements on which b is data dependent. Gyimóthy proposed an forward relevant slicing algorithm in [Gyimóthy et al. 1999], and it avoided using the huge dynamic dependence graph. However, such an algorithm will compute many redundant dependencies since it is not goal directed. In addition, while computing the dependencies of a later occurrence of certain branch statements (those which appear in the slice due to potential dependencies), the algorithm also includes statements which affect an early occurrence of the same branch statement. In this paper, we define a relevant slice over the *Extended Dynamic Dependence Graph* (EDDG), and it is more accurate than previous ones. We have also presented a space-efficient relevant slicing algorithm which operates directly on the compressed bytecode traces.

8.2 Work on Compact Trace Representations

Various compact trace representation schemes have been developed in [Goel et al. 2003; Larus 1999; Pleszkun 1994; Zhang and Gupta 2001; 2004] to reduce the high space overheads of storing and analyzing traces. Pleszkun presented a two-pass trace scheme, which recorded basic block's successors and data reference patterns [Pleszkun 1994]. The organization of his trace is similar to our trace table. However, Pleszkun's technique does not allow traces to be collected on the fly, and the trace is still large because the techniques for exploiting repetitions in the trace are limited. Recently, Larus proposed a compact and analyzable representation of a program's dynamic control flow via the on-line compression algorithm SEQUITUR [Larus

1999]. The entire trace is treated as a single string during compression, but it becomes costly to access the trace of a specific method. Zhang and Gupta suggested breaking the traces into per-method traces [Zhang and Gupta 2001]. However, it is not clear how to efficiently represent data flow in their traces. In a later work [Zhang and Gupta 2004], Zhang and Gupta presented a unified representation of different types of program traces, including control flow, value, address, and dependence.

The idea of separating out the data accesses of load/store instructions into a separate sequence (which is then compressed) is explored in [Goel et al. 2003] in the context of parallel program executions. However, this work uses the SEQUITUR algorithm which is not suitable for representing contiguous repeated patterns. In our work, we have developed RLESe to improve SEQUITUR’s space and time efficiency, by capturing contiguous repeated symbols and encoding them with their run-length. RLESe is different from the algorithm proposed by Reiss and Renieris [Reiss and Renieris 2001], since it is an on-line compression algorithm, whereas Reiss and Renieris suggested modifying SEQUITUR grammar rules in a post processing step.

Another approach for efficient tracing is not to trace all bytecodes/instructions during trace collection. This approach is explored by the *abstract execution* technique [Larus 1990], which differs from ours in how to handle untraced instructions. In particular, abstract execution executes a program P to record a small number of significant events, thereby deriving a modified program P' . The program P' then executes with the significant events as the guide; this amounts to re-executing parts of P for discovering information about instructions in P which were not traced. On the other hand, our method records certain bytecodes in an execution as a compressed representation. Post-mortem analysis of this compressed representation does not involve re-execution of the untraced bytecodes. To retrieve information about untraced bytecodes we detect dynamic dependencies via a lightweight flow analysis.

9. DISCUSSION

In this paper, we have developed a space efficient scheme for compactly representing bytecode traces of Java programs. The time overheads and compression efficiency of our representation are studied empirically. We use our compact traces for efficient dynamic slicing. We also extend our dynamic slicing algorithm to explain omission errors (errors arising from omission of a statement’s execution due to incorrect evaluation of branches/method invocations). Our method has been implemented on top of the open source Kaffe Virtual Machine. The traces are collected by monitoring Java bytecodes. The bytecode stream is compressed online and slicing is performed post-mortem by traversing the compressed trace without decompression. We have released our *JSlice* dynamic slicing tool as open-source software for usage in research and development.

ACKNOWLEDGMENTS

This work was partially supported by a Public Sector research grant from the Agency of Science Technology and Research (A*STAR), Singapore.

REFERENCES

- AGRAWAL, H. 1991. Towards automatic debugging of computer programs. Ph.D. thesis, Purdue University.
- AGRAWAL, H. AND HORGAN, J. 1990. Dynamic program slicing. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. ACM, New York, USA, 246–256.
- AGRAWAL, H., HORGAN, J., KRAUSER, E., AND LONDON, S. 1993. Incremental regression testing. In *International Conference on Software Maintenance (ICSM)*. IEEE Computer Society, Washington, DC, USA, 348–357.
- AKGUL, T., MOONEY, V., AND PANDE, S. 2004. A fast assembly level reverse execution method via dynamic slicing. In *International Conference on Software Engineering (ICSE)*. IEEE Computer Society, Edinburgh, Scotland, UK, 522–531.
- ANDERSEN, L. O. 1994. Program analysis and specialization for the c programming language. Ph.D. thesis, University of Copenhagen.
- BERK, E. J. AND ANANIAN, C. S. 2003. A lexical analyzer generator for Java. website: <http://www.cs.princeton.edu/~appel/modern/java/JLex/>.
- DHAMDHARE, D., GURURAJA, K., AND GANU, P. 2003. A compact execution history for dynamic slicing. *Information Processing Letters* 85, 145–152.
- FERRANTE, J., OTTENSTEIN, K., AND WARREN, J. 1987. The program dependence graph and its use in optimization. *ACM Transactions on Programming Languages and Systems* 9, 3, 319–349.
- GOEL, A., ROYCHOUDHURY, A., AND MITRA, T. 2003. Compactly representing parallel program executions. In *ACM Symposium on Principles and Practice of Parallel Programming (PPoPP)*. ACM, San Diego, CA, USA, 191–202.
- GYIMÓTHY, T., BESZÉDES, A., AND FORGÁCS, I. 1999. An efficient relevant slicing method for debugging. In *7th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. Springer-Verlag, Toulouse, France, 303–321.
- HORWITZ, S., REPS, T., AND BINKLEY, D. 1990. Interprocedural slicing using dependence graphs. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 12, 1, 26–60.
- JGF. The Java Grande Forum Benchmark Suite. website: <http://www.epcc.ed.ac.uk/javagrande/seq/contents.html>.
- JOY, B., STEELE, G., GOSLING, J., AND BRACHA, G. 2000. *Java(TM) Language Specification (2nd Edition)*. Prentice Hall PTR.
- KAFFE. The kaffe virtual machine. website: <http://www.kaffe.org>.
- KOREL, B. AND LASKI, J. W. 1988. Dynamic program slicing. *Information Processing Letters* 29, 3, 155–163.
- KOREL, B. AND RILLING, J. 1997. Application of dynamic slicing in program debugging. In *International Workshop on Automatic Debugging*. Springer, Linkoping, Sweden.
- LAMPORT, L. 1997. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM* 21, 558–565.
- LARSEN, L. AND HARROLD, M. 1996. Slicing object-oriented software. In *Proceedings of ACM/IEEE International Conference on Software Engineering (ICSE)*. IEEE Computer Society, Berlin, Germany, 495–505.
- LARUS, J. 1990. Abstract execution: A technique for efficiently tracing programs. *Software - Practice and Experience (SPE)* 20, 1241–1258.
- LARUS, J. R. 1999. Whole program paths. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. ACM, Atlanta, Georgia, USA, 259–269.
- LEVROUW, L. J., AUDENAERT, K. M. R., AND CAMPENHOUT, J. M. 1994. A new trace and replay system for shared memory programs based on lamport clocks. In *The second Euromicro Workshop on Parallel and Distributed Processing*. IEEE Computer Society, ELIS, Universiteit Gent, Belgium, 471–478.
- LHOTÁK, O. 2002. Spark: A flexible points-to analysis framework for Java. M.S. thesis, McGill University.

- LINDHOLM, T. AND YELLIN, F. 1999. *The Java(TM) Virtual Machine Specification (2nd Edition)*. Prentice Hall PTR.
- LUCIA, A. D. 2001. Program slicing: Methods and applications. In *IEEE International Workshop on Source Code Analysis and Manipulation*. IEEE Computer Society, Florence, Italy, 142–149.
- MAJUMDAR, R. AND JHALA, R. 2005. Path slicing. In *Intl. Conf. on Programming Language Design and Implementation (PLDI)*. ACM, Chicago, IL, USA.
- NEVILL-MANNING, C. G. AND WITTEN, I. H. 1997. Linear-time, incremental hierarchy inference for compression. In *Data Compression Conference (DCC)*. IEEE Computer Society, Snowbird, Utah, USA, 3–11.
- OHATA, F., HIROSE, K., FUJII, M., AND INOUE, K. 2001. A slicing method for object-oriented programs using lightweight dynamic information. In *Asia-Pacific Software Engineering Conference*. IEEE Computer Society, Macau, China.
- PLESZKUN, A. R. 1994. Techniques for compressing program address traces. In *IEEE/ACM International Symposium on Microarchitecture (MICRO)*. ACM, San Jose, CA, USA, 32–39.
- REISS, S. P. AND RENIERIS, M. 2001. Encoding program executions. In *ACM/IEEE International Conference on Software Engineering (ICSE)*. IEEE Computer Society, Toronto, Ontario, Canada, 221–230.
- RONSSSE, M. AND BOSSCHERE, K. D. 1999. Replay: a fully integrated practical record/replay system. *ACM Transactions on Computer Systems (TOCS)* 17, 133–152.
- SAZEIDES, Y. 2003. Instruction isomorphism in program execution. *Journal of Instruction-Level Parallelism* 5.
- SINHA, S. AND HARROLD, M. J. 2000. Analysis and testing of programs with exceptionhandling constructs. *IEEE Transactions on Software Engineering* 26, 9, 849–871.
- SPECJVM98. 1998. Spec JVM98 benchmarks. website: <http://www.specbench.org/osg/jvm98/>.
- STEENSGAARD, B. 1996. Points-to analysis in almost linear time. In *ACM SIGPLAN-SIGACT symposium on Principles of programming languages*. ACM, St. Petersburg Beach, Florida, USA, 32–41.
- TIP, F. 1995. A survey of program slicing techniques. *Journal of Programming Languages* 3, 3, 121–189.
- VALLÉE-RAI, R., CO, P., GAGNON, E., HENDREN, L., LAM, P., AND SUNDARESAN, V. 1999. Soot - a Java bytecode optimization framework. In *Conference of the Centre for Advanced Studies on Collaborative research (CASCON)*. IBM Press, Mississauga, Ontario, Canada, 13.
- WANG, T. AND ROYCHOUDHURY, A. 2004. Using compressed bytecode traces for slicing Java programs. In *ACM/IEEE International Conference on Software Engineering (ICSE)*. IEEE Computer Society, Edinburgh, Scotland, UK, 512–521. Available from <http://www.comp.nus.edu.sg/~abhik/pdf/icse04.pdf>.
- WANG, T. AND ROYCHOUDHURY, A. 2007. Dynamic slicing on Java bytecode traces. Tech. Rep. TRB3/07, National University of Singapore. March. <http://www.comp.nus.edu.sg/~abhik/pdf/JSlice-TR.pdf>.
- WEISER, M. 1984. Program slicing. *IEEE Transactions on Software Engineering* 10, 4, 352–357.
- XU, B., CHEN, Z., AND YANG, H. 2002. Dynamic slicing object-oriented programs for debugging. In *IEEE International Workshop on Source Code Analysis and Manipulation*. IEEE Computer Society, Montreal, Canada.
- ZHANG, X. AND GUPTA, R. 2004. Whole execution traces. In *IEEE/ACM International Symposium on Microarchitecture (Micro)*. IEEE Computer Society, Portland, OR, USA, 105–116.
- ZHANG, X., GUPTA, R., AND ZHANG, Y. 2005. Cost and precision tradeoffs of dynamic data slicing algorithms. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 27, 631–661.
- ZHANG, Y. AND GUPTA, R. 2001. Timestamped whole program path representation and its applications. In *ACM SIGPLAN Conference on Programming Language Design and Implementation*. ACM, Snowbird, Utah, USA, 180–190.
- ZHAO, J. 2000. Dependence analysis of Java bytecode. In *IEEE Annual International Computer Software and Applications Conference*. IEEE Computer Society, Taipei, Taiwan, 486–491.

ZILLES, C. B. AND SOHI, G. 2000. Understanding the backward slices of performance degrading instructions. In *International Symposium on Computer Architecture (ISCA)*. IEEE Computer Society, Vancouver, BC, Canada, 172–181.