A Study of Effectiveness of Dynamic Slicing in Locating Real Faults

Xiangyu Zhang Neelam Gupta Rajiv Gupta

Department of Computer Science

The University of Arizona

Tucson, AZ 85737

Contact email: gupta@cs.arizona.edu

Abstract

Dynamic slicing algorithms have been considered to aid in debugging for many years. However, as far as we know, no detailed studies on evaluating the benefits of using dynamic slicing for locating real faults present in programs have been carried out. In this paper we study the effectiveness of fault location using dynamic slicing for a set of real bugs reported in some widely used software programs. Our results show that of the 19 faults studied, 12 faults were captured by data slices, 7 required the use of full slices, and none of them required the use of relevant slices. Moreover, it was observed that dynamic slicing considerably reduced the subset of program statements that needed to be examined to locate faulty statements. Interestingly, we observed that all of the memory bugs in the faulty versions were captured by data slices. The dynamic slices that captured faulty code included 0.45% to 63.18% of statements that were executed at least once.

1 Introduction

The concept of program slicing was first introduced by Mark Weiser [21]. He introduced program slicing as a debugging aid and gave the first *static slicing* algorithm. The static slice of a reference to a variable at a program point is the set of program statements that *can influence* the value of the variable at the given program point under some execution. Therefore every statement in the program from which there is a chain of static data and/or control dependences leading to the variable reference belongs to the static slice of the variable reference. Let us consider a program containing faulty statements. Given a program point at which the value of a variable is output, if the static slice of this variable reference contains one or more faulty statements then under some executions incorrect results may be produced at the output statement. By studying the static slice of the output, a programmer may be able to detect the faulty statements. However, since static slices can be large due to the use of conservative dependence analysis, effort required to locate the faulty statement is expected to be large.

To improve the effectiveness of slicing in program debugging, Korel and Laski proposed the idea of *dynamic slicing* [14]. The dynamic slice of a variable at an execution point includes all those executed statements which *actually effected* the value of the variable at that point during a specific execution. In other words, a statement belongs to the dynamic slice of a variable reference at an execution point if there is a chain of *dynamic* data and/or control dependences from the statement to the variable reference. By studying the dynamic slice of a variable we are in a better position to determine the actual cause which led to the variable having an erroneous value under the specific execution being debugged. Since dynamic slices contain a significantly smaller subset of statements belonging to corresponding static slices, they are more suitable for debugging. Results of our study of dynamic slices reported in [26] show that the number of distinct statements executed at least once during a program execution were 2.46 to 56.08 times more than the number of statements in a dynamic slice. However, the above results were based on computing large number of slices for correct versions of programs. In this paper we will study the effectiveness of dynamic slices for locating faults in program versions with real bugs.

While in the above discussion we have considered dynamic slices that are computed by considering both data and control dependences, previous works have considered three variants of dynamic slices. Different dynamic slicing algorithms use different notions of what they consider as *influencing* the value of a variable for a given program execution. These three variants include:

- *Data slicing.* Statements that directly or indirectly influence the computation of faulty output value through chains of *dynamic data dependences* are included in data slices [25].

- *Full slicing.* Statements that directly or indirectly influence the computation of faulty output value through chains of *dynamic data and/or control dependences* are included in full slices [14].

- *Relevant slicing.* While relevant slices also consider data and control dependences, in addition, they include predicates that actually did not affect the output but could have affected it had they been evaluated differently, direct data dependences of these predicates, and chains of dynamic data and control dependences of these direct data dependences [3, 9].

While dynamic slicing has long been considered useful for debugging [1, 14, 2], experimental studies evaluating the effectiveness of slicing have not been carried out. The main goal of this paper is to experimentally evaluate the three dynamic slicing algorithms. The effectiveness of a given slicing algorithm in fault location is determined by two factors: *How often is the faulty statement present in the slice?* and *How big is the slice, i.e. how many statements are included in the slice?* We present a comparative evaluation of data and full dynamic slicing algorithms with respect to these two criteria. The following relationship holds among the various slices: Static Slice $\supseteq$ Relevant Slice $\supseteq$ Full Slice $\supseteq$ Data Slice. For the class of errors being considered, although, the faulty statement that causes an erroneous output to be produced is guaranteed to be present in the static slice and the relevant slice of the erroneous output, it may or may not be captured by the data slice and full slice.

We carried out experiments with a set of *real* faulty program versions of some widely used programs. The key results of our experimental study are as follows:

- *Applicability.* Our results show that dynamic slicing was found to be applicable in all faults studied. For 15 faults, the dynamic slice considered was the backward dynamic slice of an erroneous output. For 4 faults the program did not produce any output. In these cases we were able to capture the faults in the forward dynamic slice of the minimal failure inducing input [10]. In our study dynamic slicing we found that 12 faults were captured by data slices and 7 faults required the use of full slices. Interestingly, we observed that all of the *memory bugs* in the faulty versions which cause programs to crash due to segmentation fault were captured by the *data slices*.

- *Effectiveness.* It was observed that dynamic slicing considerably reduced the subset of program statements that needed to be examined to locate faulty statements. The dynamic slices that captured faulty code included 0.45% to 63.18% of statements that were executed at least once. These statements represented only a fraction of the total code (0.04% to 8.52%) in the programs.

- *Relevant Slicing.* Although in general faulty code may not be captured by full slices and use of relevant slicing may be required, we observed that for this set of real bugs we did not require the use of relevant slicing.

- *Exploring Slices.* Having computed the fault candidate set in form of a dynamic slice, this is next presented to the programmer who must examine it to locate faulty statements. We found that if the programmer examines statements starting from the statement that computed the erroneous value going backwards in the order of their appearance, 1.35% to 78.12% of the statements in the dynamic slices was examined before the faulty statements were located.

The rest of the paper is organized as follows. In section 2 we give describe data and full dynamic slicing. In section 3 we give the overview of our slicing tool which can be easily adapted to compute different types of slices. Section 4 presents the results of our experiments. Related work is presented in section 5 and the conclusions are given in section 6.

2 Dynamic Slicing Algorithms

In this section we illustrate the strengths and weaknesses of two dynamic slicing algorithms considered in this paper using examples. We also briefly describe the dynamic information that must be captured in order to compute the dynamic slices.

2.1 Data Slicing

Let us consider the execution of the program on an input that reveals the fault by producing an erroneous output value. Further let us assume that the presence of the faulty statement does not alter the execution control flow, i.e. the set of statements executed for this input are the same whether or not the fault is present. Under these conditions, the erroneous output must have been produced by a fault in the form of a computational error in one of statements whose computed value is related to the output value through a chain of dynamic data dependences. The data slice of the erroneous output value includes all statements that are visited by starting from the output value and then taking the transitive closure over dynamic data dependences. Thus, in the above situation, the faulty statement will be present in the data slice of the erroneous output. Because a dynamic data slice can be small and easy to understand, the faulty statement is easier to locate by examining the data slice.

The example in Figure 1 illustrates data slicing. The program on the left hand side of the figure is a faulty version of the program in which statement 13 contains an error (as indicated in the figure, $z = x - y + 1$ should be replaced by $z = x - y$). For a test input the correct and erroneous output values are shown in the figure. As we can see, this error does not alter the control flow up to the point the program generates an erroneous output value. The computation of the data slice of the erroneous output value of z yields the set of statements $\{5, 6, 13, 14\}$. Apart from the read and output statements, we have statement 13 in this data slice which is the faulty statement.

The computation of a data slice requires the identification of *dynamic data dependences* at runtime. In presence of arrays and pointers we must maintain relevant information to detect dynamic data dependences. An execution of a statement at runtime is uniquely identified by the identity of statement and the execution instance of that statement (because a statement may be executed multiple times at runtime). A dynamic data dependence exists from an

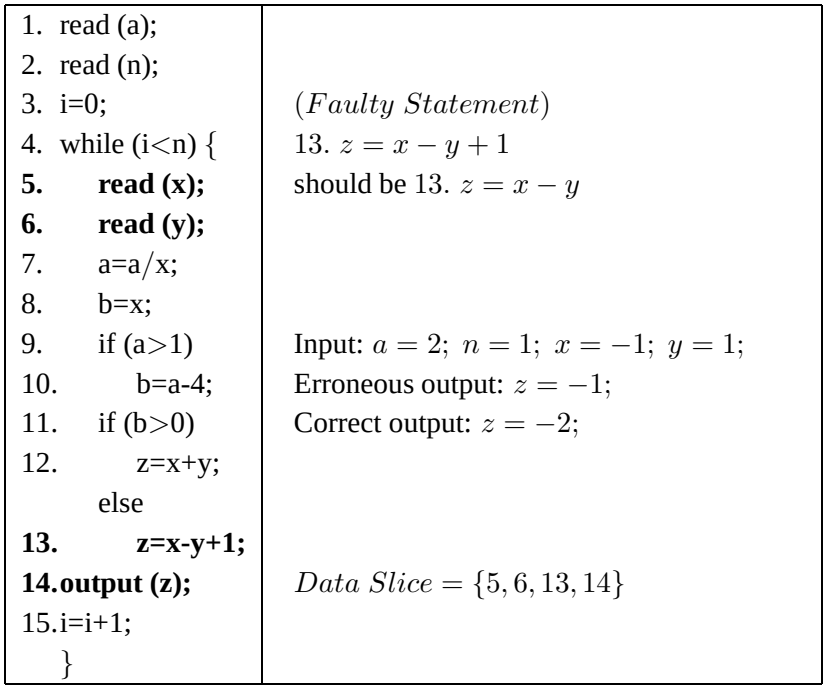

Figure 1. Example of Data Slice.

execution instance of a statement that defines a value and an execution instance of a statement that later uses that value. For each address, we must remember the execution instance of the statement that last wrote to that address. Later when the value is used by an execution instance of a statement, we can establish the dynamic data dependence between the relevant execution instances of two statements.

2.2 Full Slicing

Let us consider another example in which the faulty statement is not captured by the data slice but it is captured by the full slice. Figure 2 shows a faulty program where there is a mistake in statement 10 as shown. When this faulty program is executed on the given input, incorrect output value is produced at statement 14. The program outputs the value 4 at 14 while the correct output value is 0. The faulty statement 10 is not in set $\{5, 6, 12, 14\}$ which is the data slice of z at 14. This is because the fault does not affect value of z at 14 through a chain of dynamic data dependences. Instead fault in statement 10 affects the outcome of predicate at 11 changing the direction of the branch and thus causing statement 12 to be executed instead of statement 13. The value of z thus computed is

altered. The data slice of z at 14 contains statement 12 which is executed by mistake but it does not contain the faulty statement 10.

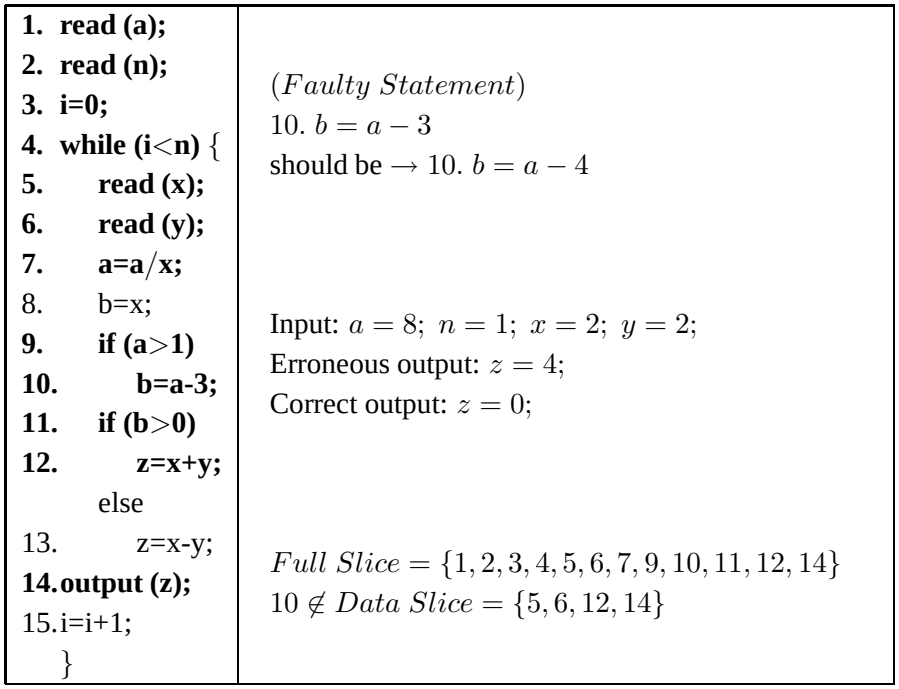

Figure 2. Example of Full Slice vs. Data Slice.

Full slices correctly handle the above situation by considering control dependences. A statement s is true (false) control dependent upon a predicate p if and only if p's true (false) outcome determines whether s will be executed. Full slices are computed by taking the transitive closure over both dynamic data and control dependence edges starting from the output value. In the above example, when both types of dependences are considered, statement 10 is included in the full slice. This is because statement 12 is control dependent upon predicate 11 which is data dependent upon statement 10.

To compute full slices, in addition to detecting dynamic data dependences, we must also detect *dynamic control dependences*. While a statement can be statically control dependent upon multiple predicates, at runtime, each execution instance of a statement is dynamically control dependent upon a single predicate. The predicate on which the execution of a statement is control dependent is found as follows. First let us assume that there are no recursive procedures. Given an execution of a statement s, prior to its execution, the most recently executed predicate p

on which s is statically control dependent is found. The execution of s is dynamically control dependent upon this execution of p. Timestamps can be associated with execution instances of statements in order to evaluate the above condition. A second condition is needed in presence of recursion. For a given execution of statement s to be dynamically control dependent upon an execution of a predicate p, the execution instances of both must correspond to the same function invocation.

3 Our Slicing Tool

We have developed a dynamic slicing tool which was used to conduct the experiments described in the next section. Our tool executes `gcc` compiler generated binaries for Intel x86 and computes dynamic slices based upon forward computation algorithms described in the preceding section. Even though our tool works on binary level, the slices can be mapped back to source code level using the debugging information generated by `gcc`.

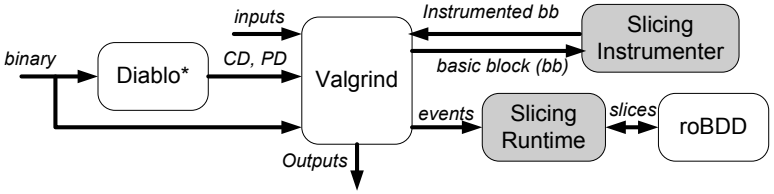

Figure 3. Slicing Infrastructure.

Figure 3 shows the main components of the tool. The *static analysis* component of our tool computes static control dependence (CD) and potential dependence (PD) information required during full and relevant slice computations from the binary. The static analysis was implemented using the *Diablo* [29] retargetable link-time binary rewriting framework as this framework already has the capability of constructing the control flow graph from x86 binary.

The *dynamic profiling* component of our system which is based upon the *Valgrind* memory debugger and profiler [30] accepts the same `gcc` generated binary, instruments it by calling the *slicing instrumenter*, and executes the instrumented code with the support of the *slicing runtime*. The slicing instrumenter and slicing runtime were developed by us to enable collection of dynamic information and computation of dynamic slices. Valgrind's kernel

is a dynamic instrumenter which takes the binary and before executing any new (never instrumented) basic blocks it calls the instrumentation function, which is provided by the slicing instrumenter. The instrumentation function instruments the provided basic block and returns the new basic block to the Valgrind kernel. The kernel executes the instrumented basic block instead of the original one. The instrumented basic block is copied to a new code space and thus it can be reused without calling the instrumenter again. The instrumentation is dynamic in the sense that the user can enforce the expiration of any instrumented basic block such that the original basic block has to be instrumented again (i.e., instrumentation can be turned on and off as desired). Thus, we can easily turn off/on the slicing instrumentation for sake of time performance or for certain code, e.g. library code. The slicing runtime essentially consists of a set of call back functions for certain events (e.g., entering functions, accessing memory, binary operations, predicates etc.). The CD and PD information computed by the static analysis component is represented based on the virtual addresses which can be understood by Valgrind.

Now let us briefly discuss the algorithms used for computing dynamic slices. Two types of methods for computing backward dynamic slices have been proposed: *backward computation* methods [1, 25]; and *forward computation* methods [7, 24]. In backward computation methods the program dependences that are exercised during a program execution are captured and saved in the form of a dynamic dependence graph. Dynamic slices are constructed upon user's requests by backward traversal of the dynamic dependence graph. Although this approach allows computation of *all* dynamic slices of all variables at all execution points, a problem with this method is its space cost. In *forward computation* methods [7, 24] latest backward dynamic slices of all program variables are computed and maintained as sets of statements as the program executes. Advantage of this approach is that the space cost is no longer proportional to the length of execution but rather proportional to the number of variables times the number of statements in the program. Therefore we decided to use forward computation method in this work.

As mentioned above, the forward computation algorithms maintain the latest dynamic slice for each variable/location. These dynamic slices are stored in *reduced ordered Binary Decision Diagram* (roBDD) [17] component of our system. Earlier work [24] identifies three characteristics of dynamic slices: same dynamic slices tend to *reappear* from

time to time during execution, different slices tend to *share statements*, and *clusters of statements* located near each other in the program often appear in a dynamic slice. These characteristics resulted in the observation that roBDD representation of sparse sets was suitable for storing dynamic slices as it was both space and time efficient. The roBDD benefits us in the following respects. Each unique slice is presented by unique integer number in roBDD, which implies that if and only if two slices are identical, they are represented by the same integer number. The whole set of statements in the slice can be recovered from roBDD using that number. This is critical to our design because now for each variable (memory location) we only need to store one integer. Use of roBDD achieves space efficiency because roBDD is capable of removing duplicate, overlapping, and clustered sets which are exactly the characteristics of slices. Using roBDD also provides time efficiency because roBDD implementations of set operations are very efficient. More details about why and how we use roBDD can be found in [24].

We also implemented a simple debugging interface which provides limited capabilities including setting breakpoints, continuing execution, stopping after certain steps of execution, slicing on a register, slicing on a memory location, and slicing on the latest instance of a predicate.

4 Experimental Evaluation

In this section we present results of experiments that we conducted. For these experiments we collected faulty versions of commonly used programs. Unlike our previous study [23] of dynamic slicing algorithms that used faulty versions of programs created by injecting faults in them, this study uses real programs with real bugs that were reported by users of these programs. We carried out two main experiments. The first experiment involves a study of the data and full slices of these programs. This experiment enabled the comparison of data and full dynamic slicing in terms of their applicability (i.e., their ability to capture faults) and effectiveness (i.e., their sizes). The second experiment shows how the computed dynamic slices may be explored by the programmer to locate faults.

The faulty versions of the programs along with the descriptions of the faults are given in Table 1. The source from which the faulty version was obtained is also given. As we can see, these programs are widely used. In

addition we would like to note that the first nine faults (i.e., faults in programs `grep 2.5` through `make 3.80`) cause the programs to produce wrong outputs while the last ten faults (i.e., faults in programs `gzip-1.2.4` through `mc-4.5.55`) contain memory bugs lead to a segmentation error. Memory bugs essentially cause memory corruption problems and when a corrupted memory location is accessed, the program crashes with a segmentation fault error message.

Table 1. Faults Used in the Study.

Program	Bug Description	Source
grep 2.5	using -i -o together produces wrong output	http://savannah.gnu.org
grep 2.5.1	(a) using -F -w together produces wrong output	http://savannah.gnu.org
	(b) using -o -n together produces wrong output	http://comments.gmane.org/ gmane.comp.gnu.grep.bugs/
	(c) ”echo doîe — grep doîe” finds no match	http://comments.gmane.org/ gmane.comp.gnu.grep.bugs/
flex 2.5.31	(a) some variable is not defined with option -l, which fails the compilation of xfree86	http://soureforge.net
	(b) string ”]]” is not allowed in user’s code	http://soureforge.net
	(c) the generated code contains extra #endif	http://soureforge.net
make 3.80	(a) backslashes in dependency names are not removed	http://savannah.gnu.org
	(b) fail to recognize the updated file status while there are multiple target in the pattern rule	http://savannah.gnu.org
gzip-1.2.4	1024 byte long filename overflows into global variable	AccMon [27]
ncompress-4.2.4	1024 byte long filename corrupts stack return address	AccMon [27]
polymorph-0.4.0	2048 byte long filename corrupts stack return address	AccMon [27]
tar-1.13.25	wrong loop bounds lead to heap object overflow	AccMon [27]
bc-1.06	misuse of bounds variable corrupts heap objects	AccMon [27]
tidy-34132	memory corruption problem	AccMon [27]
mutt-1.4.2.1i	heap buffer bound miscalculation	http://www.securiteam.com/
pine-4.44	(a) missing end quote corrupts stack	http://www.xatrix.com/
	(b) special characters corrupt heap buffer	http://www.securityfocus.com/
mc-4.5.55	uninitialized string corrupts stack	http://www.securityfocus.com/

4.1 Data Slices vs. Full Slices

Applicability and Effectiveness of Data and Full Dynamic Slicing. Our first experimented involved computing the dynamic data slices and dynamic full slices for the failed runs that exercise the faults. Before we compute dynamic slices we must identify a value in the failed run on which to perform dynamic data/full slicing. We encountered three kinds of situations in these faults which were handled as follows:

- For programs that produced an incorrect output value, backward dynamic slicing was performed starting at the first incorrect output value produced during the failed run.
- For the programs that crashed, the value which when referenced caused the crash served as the basis for computing the backward dynamic slice.
- For the four faults in `grep`, it was not possible to perform backward dynamic slicing. When these four faults were exercised the program did not crash but rather it produced incorrect output. However, this incorrect output essentially was *no output*. Since no output was produced, we did not have a value on which to base backward slicing computation. To handle these situations we found the minimal failure inducing input [10] which is the part of the input that triggered the failure. The faulty code was then captured in the *forward* dynamic slice of the failure inducing input.

The results of dynamic slicing are shown in Table 2. The column *In* indicates whether the faulty code was captured by the data slice (DS), in this case it is also captured by the full slice, or whether it is only captured by the full slice (FS). As we can see, out of the 19 faults considered, 12 faults were captured by dynamic data slices, and the remaining 7 faults were captured only by dynamic full slices. We would like to mention that in case of faults in `pine` and `mc`, where the faults are captured by the dynamic data slices, we were unable to compute the sizes of the full dynamic slices. For `pine`, the version of `diablo` used in our system was not able to handle the compiled binary because it is very large (over thirty megabytes) and thus control dependence analysis could not be performed. For `mc`, we ran out of shadow space used by `valgrind` for computing full slices. However, we were

Table 2. Data Slices and Full Slices.

Program	LOC	Exec (LOC%)	DS (Exec%)	FS (Exec%)	In	Min (LOC%)
grep 2.5	8581	1157 (13.48%)	67 (5.79%)	731 (63.18%)	FS	731 (8.52%)
grep 2.5.1 (a)	8587	509 (5.93%)	15 (2.95%)	32 (6.29%)	FS	32 (0.37%)
grep 2.5.1 (b)	8587	1123 (13.08%)	90 (8.02%)	599 (53.34%)	FS	599 (6.98%)
grep 2.5.1 (c)	8587	1338 (15.58%)	6 (0.45%)	12 (0.90%)	DS	6 (0.07%)
flex 2.5.31 (a)	26754	1871 (6.99%)	159 (8.59%)	695 (37.15%)	FS	695 (2.60%)
flex 2.5.31 (b)	26754	2198 (8.22%)	89 (4.05%)	272 (12.37%)	FS	272 (1.07%)
flex 2.5.31 (c)	26754	2053 (7.67%)	24 (1.17%)	50 (2.44%)	DS	24 (0.09%)
make 3.80 (a)	29978	2277 (7.60%)	388 (17.04%)	981 (43.08%)	FS	981 (3.27%)
make 3.80 (b)	29978	2740 (9.14%)	588 (21.46%)	1290 (47.08%)	FS	1290 (4.30%)
gzip-1.2.4	8164	118 (1.45%)	14 (11.86%)	34 (28.81%)	DS	14 (0.17%)
ncompress-4.2.4	1923	59 (3.07%)	13 (22.03%)	18 (30.51%)	DS	13 (0.68%)
polymorph-0.4.0	716	45 (6.29%)	17 (37.78%)	21 (46.67%)	DS	17 (2.38%)
tar-1.13.25	25854	445 (1.72%)	44 (9.89%)	105 (23.60%)	DS	44 (0.17%)
bc-1.06	8288	636 (7.67%)	76 (11.95%)	204 (32.07%)	DS	76 (0.92%)
tidy-34132	31132	1519 (4.88%)	148 (9.74%)	554 (36.47%)	DS	148 (0.48%)
mutt-1.4.2.1	71774	2551 (3.55%)	242 (9.49%)	1052 (41.24%)	DS	242 (0.34%)
pine-4.44 (a)	253832	3930 (1.55%)	102 (2.60%)	-	DS	102 (0.04%)
pine-4.44 (b)	253832	8956 (3.53%)	605 (6.76%)	-	DS	605 (0.24%)
mc-4.5.55	66944	3154 (4.71%)	48 (1.52%)	-	DS	48 (0.07%)

able to compute dynamic data slices for these programs. Sinces faults were captured by the dynamic data slices, the relevant results for these programs are being reported.

Now let us see how dynamic slicing reduces the amount of code the programmer has to examine to locate faulty code. In Table 2, *LOC* is the lines of code in each program, *Exec* represents the lines of code that are actually executed during the failed run (i.e., the remaining lines of code are not executed during the failed run) – the number in parenthesis is the value of *Exec* expressed as a percentage of *LOC*. *DS* and *FS* give the lines of code that are not only executed but also belong to the dynamic data slices and full slices respectively – the numbers in parenthesis are the values of *DS* and *FS* expressed as a percentage of *Exec*. Finally, *Min* is the number of lines of code in the smallest of *Exec*, *DS*, and *FS* that actually captures the faulty code – the number in parenthesis is the value of *Min* expressed as a percentage of *LOC*. In other words, *Min* is the fault candidate set that must be examined by the

programmer to locate faulty code. From the data in Table 2 we can make several observations.

By analyzing the above data we observe the following. First we notice that the lines of code in *Exec* is a small percentage ranging from 1.45% to 15.58% of the total lines of code *LOC* in the program. Since *Exec* is a small percentage of *LOC*, even this rudimentary dynamic information is quite effective in reducing the size of the fault candidate set presented to the programmer for examination. Second, we observe that the sizes of dynamic data and full slices are significantly smaller than *Exec*. The sizes of *DS* range from 0.45% to 37.78% of the sizes of *Exec* and the sizes of *FS* range from 0.90% to 63.18% of the sizes of *Exec*. We also observe that sizes of dynamic data slices are significantly smaller than sizes of dynamic full slices in most of the cases. Finally, the *Min* column we present the size of the fault candidate set that is of significance for the programmer. We observe that the lines of code in *Min* is a very small percentage ranging from 0.04% to 8.52% of *LOC* the total lines of code in the program. Thus we conclude that dynamic information offers significant reductions in the size of the fault candidate set.

Memory Bugs. One key issue is when to use dynamic data slices and when to use full dynamic slices. We observe that for all faults that are memory bugs dynamic data slices captured the faulty code. It is easy to identify that the program has been effected by a memory bug when it crashes with a segmentation fault error. In such situations the user can use dynamic data slicing instead of full dynamic slicing. Through further analysis that we next describe, we determined the reason due to which dynamic data slices are so effective for programs with memory bugs that cause program to crash with a segmentation fault. In other words, even though data slicing is not effective in capturing faulty statement in general, it is very effective for memory related bugs. Since the data slices can be significantly smaller than the full slices (e.g., `tar`, `bc`, etc.) and therefore using data slices for memory related bugs it quite advantageous.

The reason why data slices are so effective for memory bugs is that the program crash is caused due to the presence of an *unexpected dynamic data dependence* between the point at which memory is corrupted and the later point at which the corrupted value is used. In fact the memory corruption typically corrupts a pointer and its use causes a crash because it dereferences the pointer. Dynamic data slice captures all appropriate dynamic data

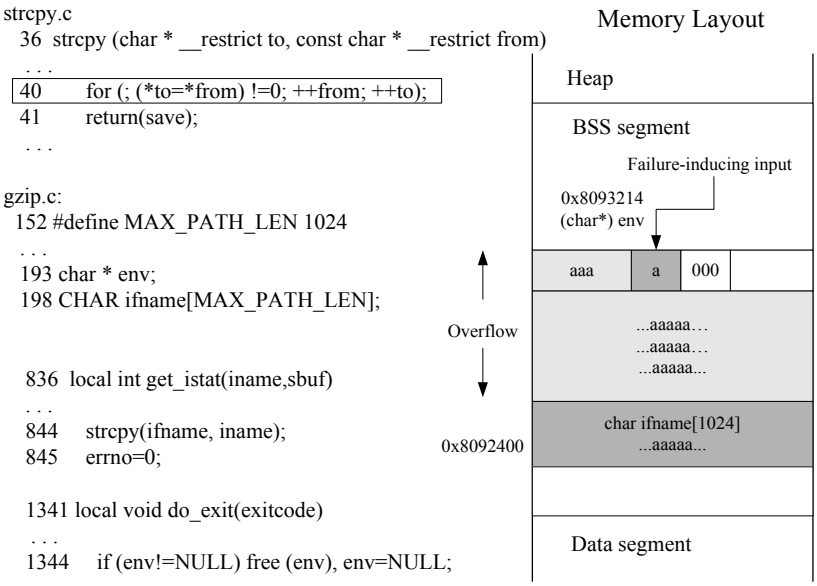

Figure 4. `gzip` **Buffer Overflow Bug.**

dependences including the unexpected dynamic data dependence and therefore it is able to capture faulty code. To illustrate the above, let us consider the case of `gzip` which contains a buffer overflow problem. In Figure 4, on the left hand side we show the relevant code segment for the problem. The problem happens in the `strcpy` statement at line 844. Variable `ifname` is a global array defined at line 198. The size of the array is defined as 1024. Before the `strcpy` statement at line 844, there is no check on the length of the string `iname`. If the length of string `iname` is longer than 1024, then the buffer overflows. If the length of string `iname` is larger than 3604, the value of `env` is changed due to buffer overflow. This is because according to the memory layout shown in Figure 4, the difference between `env` and `ifname` is 3604 bytes. Later when at line 1344 *free(env)* is executed, the program crashes due to presence of an illegal memory address in `env`. When dynamic data slice is computed for this illegal address, the faulty statement in `strcpy` is captured in the dynamic data slice.

Relevant Slicing. It has been observed [3, 9] that in some situations faulty statements are not captured by full slices. Consider the following faulty version of a program. Let us consider the situation in which statement $y = 0;$ is erroneous and it causes the predicate $y > 0$ to evaluate to false instead of being true. False evaluation of the predicate causes the execution of the assignment to x inside the if-statement to be bypassed leading to incorrect

value of x to be output. Since the statement inside the if-statement is not executed there is not dependence between the output statement and the faulty statement $y = 0;$. In other words, the dynamic full slice does not capture the faulty statement.

$x = 1;$

$y = 0;$

`if` $y > 0$ `then`

  $x = 2;$

`endif`

`output`$(x);$

In general, the basic reason is that some statements which should have been executed did not get executed due to the fault. To handle the above situation a new form of dependence needs to be introduced between certain predicate outcomes and uses. Given a use u, let us define a *potentially depends* set $PD(u)$ such that the set contains members of the form that specify predicates and their outcomes (i.e., p^T or p^F). If p^T (p^F) is present in $PD(u)$, it means that if prior to the execution of u predicate p was executed, and its outcome was T (F), then while no definition corresponding to u was encountered, it could have been encountered if p had evaluated to F (T). For the above example this means that $(y > 0)^F \in PD(output(x))$ because when the outcome of predicate $y > 0$ is F, no definition of x is encountered after execution of $y > 0$ while if $y > 0$ had evaluated to T the definition $x = 2$ would have been encountered. The potentially depends property is a static property of u which is precomputed and later used at runtime to compute relevant slices.

In an earlier study [23] we reported that when faults are present in predicate statements, full slices are sometimes inadequate and therefore one must use dynamic relevant slices. In this earlier study faults were artificially injected in predicates and studied. In contrast, the results reported in this paper are based upon some real bugs reported by users. We observed that for these real bugs relevant slicing was not needed even though some of these bugs did influence the outcomes of predicates during the failed run. To understand why relevant slices were needed in the

earlier study but not in this new study we further studied the nature of bugs in the programs. In the earlier study based upon Siemens suite we noticed that many bugs were injected by changing the predicates and even shortening the predicates by eliminating part of the condition. As a result the situation of the type illustrated earlier where code that should have been executed is bypassed arose requiring the need for relevant slicing. On the other hand, when we studied the incorrect evaluations of predicates in real bugs we noticed a very different behavior. In most of the cases incorrect evaluation of predicates was present in programs with buffer overflow bugs. Here incorrect evaluation of a look predicate caused the loop body to be executed too many times leading to buffer overflow and memory corruption which caused the program to crash. In other words, the incorrect evaluation of predicates did not cause execution of code to be bypassed and hence the need for using relevant slicing did not arise.

4.2 Exploring Dynamic Slices

A dynamic slice provides a fault candidate set that the programmer must examine to identify the faulty statement. Therefore smaller the set of statements that the user has to examine the better it is. Even though dynamic slices produce fault candidate sets that are small in comparison to the set of executed statements, it can still be quite a lot of work to examine all of the statements in these slices. Therefore we considered a strategy in which the statements in the dynamic slice are *ordered* and the programmer examines the statements in that order. Once the faulty code is encountered by the programmer, the fault is located and the programmer need not examine rest of the dynamic slice. In other words, the programmer need not always explore the entire dynamic slice. The strategy we used orders the statements according to the *dependence distance* between them and the statement at which error was observed. More precisely, the dependence distance of a statement in the dynamic slice is the length of the minimum length chain of dependences starting from the statement and ending at the statement at which error was observed.

The results of this experiment are discussed next. In Table 3 the column *Slice Type* indicates the kind of slice that was explored in this experiment. As we can see, we explored the dynamic data slices (DS) for programs with memory bugs and dynamic full slices (FS) for other programs. Based upon the observations of the preceding sections this choice is most appropriate. The column *Slice Size* (SS) gives the size of the dynamic slice being

Table 3. Exploring Dynamic Slices.

Program	Slice Type	Slice Size (SS)	Explored SS (ESS)	EDD
grep 2.5	FS	731	86 (11.76%)	9
grep 2.5.1 (a)	FS	32	25 (78.12%)	8
grep 2.5.1 (b)	FS	599	157 (26.21%)	11
grep 2.5.1 (c)	FS	12	6 (50.00%)	3
flex 2.5.31 (a)	FS	695	13 (1.87%)	5
flex 2.5.31 (b)	FS	272	109 (40.07%)	31
flex 2.5.31 (c)	FS	50	3 (6.00%)	2
make 3.80 (a)	FS	981	187 (19.06%)	21
make 3.80 (b)	FS	1290	53 (4.11%)	19
gzip-1.2.4	DS	14	2 (14.28%)	2
ncompress-4.2.4	DS	13	1 (7.69%)	1
polymorph-0.4.0	DS	17	4 (23.53%)	3
tar-1.13.25	DS	44	5 (11.36%)	4
bc-1.06	DS	76	4 (5.26%)	3
tidy-34132	DS	148	2 (1.35%)	2
mutt-1.4.2.1	DS	242	17 (7.02%)	4
pine-4.44 (a)	DS	102	3 (2.94%)	3
pine-4.44 (b)	DS	605	38 (6.28%)	18
mc-4.5.55	DS	48	2 (4.17%)	2

explored and *Explored Slice Size* gives the size of portion of the slice that was explored before the faulty code was encountered. The size of ESS as a percentage of SS is also given in parenthesis. As we can see, the lines of code in the dynamic slice that were explored as a percentage of the total lines of code in the dynamic slice ranges from 1.35% to 78.12%. In seven out of eleven cases this number is in single digits. Thus, using our proposed strategy, in practice, a programmer has to examine far fewer statements. Finally the maximum dependence distance up to which the dynamic slice was explored (EDD) is given. As we can see this dependence distance was found to be small for programs where dynamic full slices were used and for programs with memory bugs this distance was mostly one.

4.3 Cost of Dynamic Slicing

The cost of dynamic slicing consists of two main components: the space cost which is the memory needed to store the dynamic dependence graph (DDG) required for computing the dynamic slices; and the execution time cost which includes the time to collect the runtime information and build the dynamic dynamic dependence graph and the time to perform dynamic slicing. The above costs for the faults studied are given in Table 4. The size of the dynamic dependence graph is given in column *DDG Size*. The size of the graph depends upon the length of the failing program run. As we can see the size varies from 173 KB to nearly 209 MB. The time spent on building the dynamic dependence graph, given by column *DDG Time*, ranges from 0.4 seconds to 284.1 seconds. As we can see, the time is typically proportional to the length of the run, i.e. the size of the DDG. The slicing times are given in column *Slicing Time* and they range from 0.01 to 6.71 seconds.

Table 4. Dynamic Dependence Graph Size and Execution Times.

Program	DDG Size (KB)	Slicing Time (seconds)	DDG Time (seconds)
grep 2.5	760	0.04	35.5
grep 2.5.1(a)	794	0.04	29.2
grep 2.5.1(b)	333	0.02	4.4
grep 2.5.1(c)	968	0.06	20.1
flex-3.51(a)	196131	4.39	135.5
flex-3.5.31(b)	202441	3.14	138.9
flex-3.5.31(c)	199170	6.71	130.2
make 3.80(a)	17409	0.24	28.0
make 3.80(b)	15801	1.74	34.6
gzip-1.2.4	164	0.01	1.2
ncompress	211	0.03	1.1
polymorph	173	0.03	0.4
tar	420	0.01	10.9
bc	1404	0.15	6.7
tidy	92872	0.53	17.5
mutt-1.4.2.1	74358	4.34	284.1
pine-4.44 (a)	44108	6.16	63.5
pine-4.44 (b)	70266	4.33	68.4
mc-4.5.55	208849	0.6	120.7

5 Related Work

Dynamic slicing was introduced as an aid to debugging by Korel and Laski in 1988 [14]. Although the idea seems very promising, it has not been used in practice. There is a practical reason for this. The problem of the high cost of computing dynamic slices had not been addressed till recently. In recent work [26, 24], we developed practical implementations of dynamic slicing for both backward computation [26] and forward computation [24] algorithms have been developed. We demonstrated that dynamic slices of program runs that were 67 million to 140 million instructions in length, on an average, took 1.92 seconds to 36.25 seconds to compute [26].

Dynamic slicing has been studied as an aid to debugging by many researchers [2, 13, 15, 4, 18]. Agrawal et al. [4] proposed subtracting a single correct execution trace from a single failed execution trace. In [18], Pan and Spafford presented a family of heuristics for fault localization using dynamic slicing. Compared to these previous works, we are the first one to compare the effectiveness of dynamic slicing algorithms in fault location.

General studies of dynamic slice sizes have been conducted. For example in our work in [26] showed that the number of distinct statements executed at least once during a program execution were 2.46 to 56.08 times more than the number of statements in the dynamic slice. However, these results are based upon computing dynamic slices of randomly selected values computed by correct versions of programs. In another study [23] we computed dynamic slices based upon failed runs of faulty versions of programs. These faults had been injected into the programs. In contrast the study presented in this paper consider a set of real faults reported by users of widely used programs. Some of the observations of this study based upon real faults are different from those of the previous study of injected faults. The differences and the explanation for these differences are as follows:

- First, in our current study, for the faults in `grep`, no output was produced and hence instead of backward dynamic slices we had to make use of forward dynamic slices. Similar situation did not arise for the Siemens suite programs used in the earlier study [23].

- Second, in the earlier study the need for using relevant slicing arose while in our current study data and full dynamic slices were able to capture all faults. As explained earlier, in our current study that contains many

memory bugs, most of the situations where predicates evaluated incorrectly, the incorrect evaluation did not cause bypassing of the code but rather execution of the code that should have been bypassed. Therefore, relevant slicing was not needed.

We would also like to point out that, while examining the statements in a dynamic slice, the relevance of dependence distance from the erroneous output has long been considered useful [5, 6, 16]. Therefore in tools for visualizing dynamic slices, ways have been explored to communicate to the programmer the dependence distance information. For example, in [16], Krinke uses different shades of gray to highlight the statements in the dynamic slice. In particular, the darker the shade, the smaller is the dependence distance. In this paper, through experiments, we have validated the merit of using dependence distance information while exploring dynamic slices.

A lot of interesting research other than dynamic slicing have been carried on in fault location. Zeller has presented a series techniques [10, 22, 8] from isolating the critical input to isolating cost-effect chains in both space and time. The basic idea is to find the specific part of the *input/program state* which is critical to the program failure by minimizing the difference between the *input/program state* leading to a passing run and that leading to a failing run. We believe our technique can be combined with Zeller's technique in many aspects, for instance, the isolated *causes* are perfect slicing criteria starting from which dynamic slicing may produce a much smaller fault candidate set than from the failure point. Renieris and Reiss [20] presented a technique that selects the single passing run that most resembles to the failing run and reports the difference between these two runs. Jones [12] presented a technique that uses software visualization to assist fault location. Their technique provides a ranking of each statement according to its ratio of failing tests to correct tests.

6 Conclusions

The development of dynamic slicing was motivated by the problem of locating the faulty code when an execution of a program fails. There has been a significant amount of research on developing algorithms for computing different types of dynamic slices. The contribution of this paper is to present an experimental evaluation of effectiveness of

dynamic slices for the benefit of using them to locate a faulty statement in a program. In particular, this is the first study based upon real faults reported by users of widely used programs. From our experiments we found that data slices were found to be very effective for memory related faults and for remaining faults full slicing was adequate. None of the faults required the use of relevant slicing. Finally, we found that even if the slice size is large, the user may have to examine only a subset of statements in the slice before encountering the faulty statement.

Acknowledgements:

We would like to thank the reviewers for their suggestions that encouraged us to do even more through job of revising the original submission.

References

[1] H. Agrawal and J. Horgan, "Dynamic Program Slicing," *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 246-256, 1990.

[2] H. Agrawal, R. DeMillo, and E. Spafford, "Debugging with Dynamic Slicing and Backtracking," *Software Practice and Experience*, Vol. 23, No. 6, pages 589-616, 1993.

[3] H. Agrawal, J.R. Horgan, E.W. , and S.A. London, "Incremental Regression Testing", *IEEE Conference on Software Maintenance*, pages 348-357, Montreal, Canada, 1993.

[4] H. Agrawal, J. Horgan, S. London, and W. Wong, "Fault Localization Using Execution Slices and Dataflow Tests," *6th IEEE International Symposium on Software Reliability Engineering*, pages 143-151, 1995.

[5] G. Antoniol, R. Fiutem, G. Lutteri, P. Tonella, S. Zanfei, and E. Merlo, "Program Understanding and Maintenance with the CANTO Environment," *International Conference on Software Maintenance*, pages 72-, Bari, Italy, October 1997.

[6] T. Ball and S.G. Eick, "Visualizing Program Slices," *IEEE Symposium on Visual Languages*, pages 288-295, St. Louis, Missouri, October 1994.

[7] A. Beszedes, T. Gergely, Z.M. Szabo, J. Csirik, and T. Gyimothy, “Dynamic Slicing Method for Maintenance of Large C Programs,” *5th European Conference on Software Maintenance and Reengineering*, pages 105-113, March 2001.

[8] H. Cleve and Andreas Zeller, “Locating Causes of Program Failures”, *27th International Conference on Software Engineering*, pages 342-351, 2005.

[9] T. Gyimothy, A. Beszedes, I. Forgacs, “An Efficient Relevant Slicing Method for Debugging”, *7th European Software Engineering Conference and 7th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, pages 303-321, Toulouse, France, 1999.

[10] R. Hildebrandt and A. Zeller, “Simplifying Failure-inducing Input”, *International Symposium on Software Testing and Analysis*, pages 135-145,2000.

[11] M. Hutchins, H. Foster, T. Goradia, and T. Ostrand, “Experiments on the Effectiveness of Dataflow- and Controlflow-based Test Adequacy Criteria,” *16th International Conference on Software Engineering*, pages 191-200, 1994.

[12] J.A. Jones, ”Fault Localization Using Visualization of Test Information”, *26th International Conference on Software Engineering*, page 54-56,2004.

[13] M. Kamkar, “Interprocedural Dynamic Slicing with Applications to Debugging and Testing,” *PhD Thesis*, Linkoping University, 1993.

[14] B. Korel and J. Laski, “Dynamic Program Slicing,” *Information Processing Letters*, Vol. 29, No. 3, pages 155-163, 1988.

[15] B. Korel and J. Rilling, “Application of Dynamic Slicing in Program Debugging,” *3rd International Workshop on Automatic Debugging*, pages 43-58, Linkoping, Sweden, 1997.

[16] J. Krinke, “Visualization of Program Dependence and Slices,” *International Conference on Software Maintenance*, pages 168-177, 2004.

[17] J. Lin-Nielsen. “BuDDy, A Binary Decision Diagram Package,” Department of Information Technology, Technical University of Denmark, *http://www.itu.dk/research/buddy/*.

[18] H. Pan and E. H. Spafford, “Heuristics for Automatic Localization of Software Faults”, *Technical Report SERC-TR-116-P*, Purdue University, 1992.

[19] S. Narayanaswamy, G. Pokam, and B. Calder, “BugNet: continuously recording program execution for deterministic replay debugging,” *32nd International Symposium on Computer Architecture*, pages 284-295, 2005.

[20] M. Renieris and S. Reiss, “Fault Localization with Nearest Neighbor Queries,” *IEEE International Conference on Automated Software Engineering*, pages 30-39, 2003.

[21] M. Weiser, “Program Slicing,” *IEEE Transactions on Software Engineering*, Vol. SE-10, No. 4, pages 352-357, 1982.

[22] A. Zeller, “Isolating Cause-effect Chains from Computer Programs”, *10th ACM SIGSOFT Symposium on Foundations of Software Engineering*, pages 1-10, Charleston, South Carolina, 2002.

[23] X. Zhang, H. He, N. Gupta, and R. Gupta, “Experimental Evaluation of using Dynamic Slices for Fault Location,” *SIGSOFT-SIGPLAN Sixth International Symposium on Automated and Analysis-Driven Debugging*, pages 33-42, Moterey, California, September 2005.

[24] X. Zhang, R. Gupta, and Y. Zhang, “Effective Forward Computation of Dynamic Slices Using Reduced Ordered Binary Decision Diagrams,” *IEEE International Conference on Software Engineering*, pages 502-511, Edinburgh, UK, 2004.

[25] X. Zhang, R. Gupta, and Y. Zhang, “Precise Dynamic Slicing Algorithms,” *IEEE/ACM International Conference on Software Engineering*, pages 319-329, Portland, Oregon, May 2003.

[26] X. Zhang and R. Gupta, “Cost Effective Dynamic Program Slicing,” *ACM SIGPLAN Conference on Programming Language Design and Implementation*, pages 94-106, June 2004.

[27] P. Zhou, W. Liu, L. Fei, S. Lu, F. Qin, Y. Zhou, S.P. Midkiff, and J. Torrellas, “AccMon: automatically detecting memory-related bugs via program counter-based invariants,” *37th Annual International Symposium on Microarchitecture*, pages 269-280, 2004.

[28] *http://www.cse.unl.edu/~galileo/sir*

[29] *http://www.elis.ugent.be/diablo/*

[30] *http://valgrind.org/*