

Dynamic Slicing Long Running Programs through Execution Fast Forwarding

Xiangyu Zhang Sriraman Tallam Rajiv Gupta
Department of Computer Science
The University of Arizona
Tucson, Arizona 85721
{xyzhang,tmsriram,gupta}@cs.arizona.edu

ABSTRACT

Fixing runtime bugs in long running programs using tracing based analyses such as dynamic slicing was believed to be prohibitively expensive. In this paper, we present a novel *execution fast forwarding* technique that makes it feasible. While a naive solution is to divide the entire execution by checkpoints, and then apply dynamic slicing enabled by tracing on one checkpoint interval at a time, it is still too costly even with state-of-the-art tracing techniques. Our technique is derived from two key observations. The first one is that long running programs are usually driven by events, which has been taken advantage of by checkpointing/replaying techniques to deterministically replay an execution from the event log. The second observation is that all the events are not relevant to replaying a particular part of the execution, in which the programmer suspects an error happened. We develop a slicing-like technique on the event log such that many irrelevant events are successfully pruned. Driven by the reduced log, the replayed execution is now traced for fault location. This replayed execution has the effect of fast forwarding, i.e the amount of executed instructions is significantly reduced without losing the accuracy of reproducing the failure. We describe how execution fast forwarding is combined with checkpointing and tracing based dynamic slicing, which we believe is the first attempt to integrate these two techniques. The dynamic slices of a set of reported bugs for long running programs are studied to show the effectiveness of dynamic slicing, which is a significant step forward compared to our prior work.

1. INTRODUCTION

During the procedure of debugging, it is often the case that the programmer is interested in a very little part of the entire execution. How to get to this region quickly has been haunting the researchers since debugging long running programs became an issue. The traditional debugging tactics, such as iteratively setting breakpoints and then restarting the program, hardly work because of the re-execution consumes enormous amount of time. More sophisticated methods to tackle this problem include *tracing* and *checkpointing/replaying*.

Tracing is a technique with long history. It was invented for the purpose of replaying an execution. More and more applications have been developed such as performance analysis, software reliability, software understanding, and compiler optimizations. Traces are usually collected once and then they are analyzed multiple times starting from any point. Furthermore many heavy duty analyses can be performed on traces efficiently. As a result, software errors become much more recognizable if appropriate traces are gathered. For example, dynamic slicing, proposed by Korel and Laski [5], is a tracing based technique to help programmers in the pro-

cess of debugging. The dynamic slice of a value computed at a program point in the execution trace includes all those executed statements which were directly or indirectly involved in computation of the value. Our prior work [18, 3, 19, 20] has demonstrated that dynamic slicing is quite effective in automatically isolating the cause effect chain from the root cause to the failed point. Unfortunately, tracing based techniques do not scale for long executions even though state-of-the-art techniques can achieve the space efficiency of 0.1-4 bits per instruction [17, 2]. A simple task as starting Mozilla and browsing a html page may create traces with the size of a few giga bytes.

Checkpointing/replaying is a very attractive technique, the merit of which is the capability of replaying from the intermediate points of the execution once a checkpoint is created. It was invented to facilitate debugging parallel and distributed programs [8, 16]. It quickly gained popularity in general application debugging [11, 12]. A lot of research has been carried out on how to reduce its cost [15, 7] and improving its usability [13]. Most of the existing checkpointing techniques focus on how to faithfully replay an execution. They rarely discuss what to do with the replayed execution or simply suggest that the replayed execution can be debugged with general debuggers such as gdb. However, these debuggers are usually much less powerful than tracing based tools.

Our goal is to apply dynamic slicing, a tracing based technique, to long running programs. A natural question to ask is "*Can we combine tracing and checkpointing?*". It seems tracing and checkpointing are complimentary. Checkpoints divide the whole execution into intervals. Tracing can be applied to one interval at a time, usually the one that interests the programmer. However, this solution is not as simple as it appears for two reasons. First, tracing requires instrumenting the original program. There are two kinds of instrumentation techniques – static and dynamic. Static instrumentation, in which the program is instrumented by compilers, introduces non-trivial execution overhead as tracing cannot be easily turned off. Dynamic instrumentation adaptively instruments the program. It can easily switch from executing the original code to executing the instrumented code or vice versa. Dynamic instrumentation engine usually resides in the process's virtual space and manipulates the virtual memory intensively such that the status of the application process is substantially mixed with the instrumentation engine's status. While checkpoints are often produced by taking snapshots of the virtual memory, it becomes hard to discretely checkpoint the application process. Second, tracing can handle executions up to a few seconds given the speed and storage capacity of today's desktops. Since checkpointing usually produces virtual memory snapshots with the size of a few mega bytes, it is not something that we can afford to perform every second. Checkpoints are

usually created in the interval of, more or less, minutes. The gap between seconds and minutes suggests that it is still too costly to trace a checkpoint interval.

In this paper, we present a novel *execution fast forwarding (EFF)* technique that fills the gap between tracing and checkpointing. It enables dynamic slicing on long executions. Figure 1 illustrates the basic idea. The left part illustrates that an execution, or part of an execution delimited by checkpoints, is usually heavily instrumented for the purpose of dependence tracing. The heavy instrumentation introduces very high runtime overhead and constructs a huge dependence graph, which makes it impractical if the execution gets long. In the right part the fast forwarding technique takes advantage of the characteristics of many long running programs – being driven by events. More precisely, it first collects a full event log from the original execution; given a specific part of the execution that the programmer wants to replay, a meta slicing technique, which is analogous to dynamic slicing but performed on logged events instead of executed instructions, is applied to prune the events irrelevant to replaying the desired execution region. The reduced event log is used to drive the replay, which is also called *the fast forwarded execution*. Compared to the original run, the fast forwarded execution is much smaller as the volume of events passed to the program is significantly lower. As a result, a smaller dependence graph is generated that can be collected through tracing. The contributions of our paper are summarized as follows.

- We propose a novel EFF technique that performs meta slicing on the event log to eliminate the events that are not relevant to replaying a specific part of the execution. The reduced event log is used to drive the replay to achieve the effect of fast forwarding.
- To implement the EFF technique, we solve the problem of combining tracing and logging/checkpointing. As far as we know, this is the first attempt to put an application process under the supervision of both dynamic instrumentation and logging/checkpointing. Given the strength of these techniques, we believe integrating them has very high potential to impact the existing debugging procedures.
- As the ultimate goal of EFF, dynamic slicing is applied on a set of long running programs, which was not possible previously due to the extremely high expense. The results strongly support our claim in the prior work – *dynamic slicing is very effective in isolating the cause effect chain from the root cause to the failure*.

The remainder of the paper is organized as follows. In section 2 we describe the EFF technique in detail. The system, which is an integration of EFF, tracing and checkpointing, is introduced in section 3. The results of our experiments are presented in section 4. In section 5 we studied the effectiveness of dynamic slicing on long running programs. Conclusions are given in section 6.

2. EXECUTION FAST FORWARDING

Often when a program runs for a long time it is not because it performs a very long and complicated task. Instead, it is often because the program processes a long sequence of simple tasks. For example, programs processing streaming data such as audio, video, and packet data usually carry out the same computation e.g. FFT transformation on a sequence of data; the computation on each data piece tends to be relatively lightweight and independent from each other. Programs that require user interactions display similar properties: the programs spend most of their execution time in handling

user actions and the computation dedicated for each user action is usually simple. Server programs deal with thousands of requests, most of which set off simple computations such as reading a file or retrieving a piece of data from a database. A common feature of these programs is that *they are driven by events*. The events divide the whole execution into small tasks, each one of which corresponds to handling some event. An event is defined as one interaction between the application and the OS. The interaction could be in the forms of: system calls such as *open*, *read*, and *mmap2*; asynchronous or synchronous signals such as *kill* and *segfault*. These events are used to provide OS services, for instance reading/writing a file/socket, to the application program, or to notify something has happened.

The EFF technique is derived from the following observation – *all the events do not need to be replayed in order to replay a particular part of the execution*. Given the fact that the execution is driven by events, we may be able to shrink the replayed execution, and yet reproduce the desired part, if we can prune the irrelevant events.

Figure 2 presents a motivation example. In the original run, the key 'c' was first pressed in order to change the folder name after Mutt, a text based mail user agent, was started; string *"imaps://xyzhang@email.cs.arizona.edu/inbox"* was typed in as the email account, which was followed by the password; after logging in the account, a couple of email messages were accessed; then 'c' was typed again and string *"Hello"* was provided as the new folder name. Since *"Hello"* was not a valid folder name, a warning message was printed on the screen. The events were logged in a file as shown on the left hand side of the figure. The first a few thousands of events present the startup phase of the execution, which is mainly about loading dynamic libraries, allocating virtual memory, and initializing the program state. The shaded events starting from byte position 4898 to position 594803 correspond to the execution related to accessing the email account. Events starting from 594804 contribute to entering the invalid folder name and the warning message was printed by the event at 595007. Let us assume the programmer is interested in reproducing the warning message. Apparently, replaying the entire execution with the full log is an option but not the optimal one. For the event at 595007 to be correctly replayed, we need to replay events at 594804, 594825, ..., 594890, etc. Events from 4898 to 594803 are actually *irrelevant* to replaying the event at 595007. We constructed a new log by removing all the irrelevant events and then drove the replay with the reduced log. The same warning message was successfully reproduced. The execution was actually *fast forwarded* to the desired point by skipping the irrelevant part.

The EFF technique poses two challenges. The first one is how to identify and remove the irrelevant events; the second one is how to replay with the reduced event log. The following subsections describe how we handle these issues.

2.1 Event Dependence Graph

In dynamic slicing, given a value that is observed to be incorrect by the programmer (incorrect value may correspond to an incorrect output or a value that causes the program to crash). A set of executed statements that contributed to the value of the specified variable are computed as its dynamic slice. The executed statements not in the dynamic slice are not relevant to the investigated value. An analogous solution can be applied on the executed events to identify the set of irrelevant events for replaying a given execution region.

Computation of dynamic slices normally consists of two steps: building the dynamic dependence graph (DDG) for a program ex-

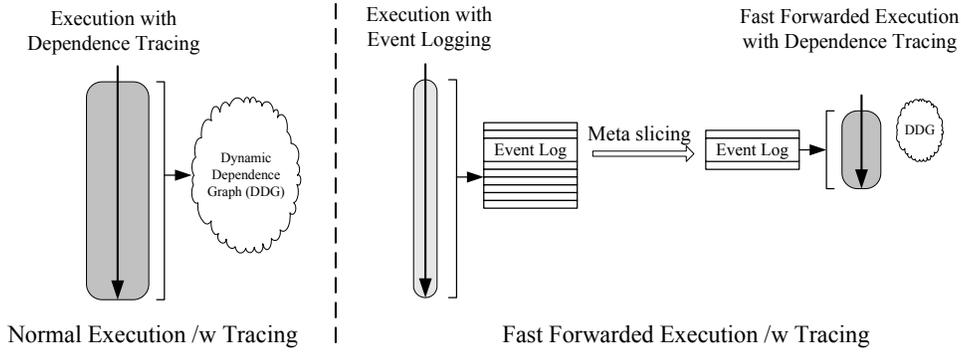


Figure 1: Execution Fast Forwarding.

ecution (where dependences include both data and control dependences); and then traversing the dynamic dependence graph to compute the dynamic slice of the wrong value. To simplify the description, we assume the execution starts from the beginning. We will discuss how to deal with executions starting from checkpoints in later sections.

DEFINITION 1. The **Dynamic Dependence Graph** of a program run, $DDG(N, E)$, consists of a set of nodes N and a set of directed edges E where: each node $n_i \in N$ corresponds to i^{th} execution instance of statement n in the program; and each edge $m_j \rightarrow n_i \in E$ corresponds to a dynamic data dependence, dynamic control dependence, or potential dependence of i^{th} execution instance of statement n on the j^{th} execution instance of statement m .

In a DDG, an executed statement is abstracted as $S_j(U, D)$, which means the j instance of statement S . U denotes the set of values used by S_j and D denotes the set of values defined. For example, the execution of statement "store $r_1, [r_2]$ " can be abstracted as "...($U = \{r_1, r_2\}, D = \{[r_2]\}$ ", in which $[r_2]$ represents the memory location addressed by r_2 . A data dependence exists between two executed statements if the U set of one statement overlaps the D set of the other. A control dependence is introduced if the execution of one statement depends on the values in D of the other statement, usually a predicate statement. One executed statement S_j potentially depends on another executed statement, usually a predicate, if and only if the value of the executed statement could have changed if the predicate had taken a different branch. More details about potential dependence can be found in [4, 18].

We already discussed how an executed statement is abstracted. As an event usually corresponds to multiple executed statements, it is important to understand how we deal with events during the DDG construction. Since an event is usually handled inside the OS kernel, a tracing engine which runs in the application space is not able to trace into the kernel. Hence the dependences within the event handler are not captured. Our solution is to summarize the execution of an event into the same abstraction, $E_j(U, D)$, according to the specifications of events. For instance, event " $n=read(fd, Buf, size)$ " can be abstracted as "...($U = \{fd, seek_pointer(fd), size, Buf\}, D = \{seek_pointer(fd), Buf[0], Buf[1], \dots, Buf[n-1]\}$). Note that only the first n elements of Buf are defined according to the specification of event *read*. This event both defines and uses the seek pointer of file fd .

An analogous dependence graph, *Event Dependence Graph (EDG)*, can be constructed to reveal the dependences within events, which can be later on used to prune the irrelevant events.

DEFINITION 2. The **Event Dependence Graph** of a program

run, $EDG(N, E)$, consists of a set of nodes N and a set of directed edges E where: each node $n_i \in N$ corresponds to the i^{th} execution instance of event n in the program; and each edge $m_j \rightarrow n_i \in E$ denotes that there exists a dependence path from m_j to n_i , and there are no other executed events than m_j and n_i on the path.

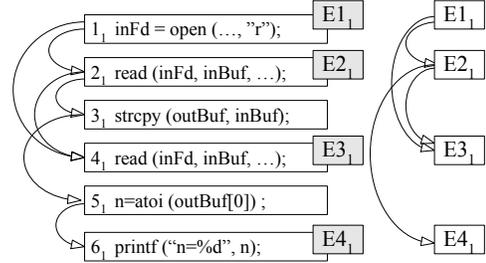


Figure 3: An example of Dynamic Dependence Graph (DDG) and Event Dependence Graph (EDG).

Figure 3 presents an example to illustrate DDG and EDG. The left hand side presents the DDG for the execution of a small piece of code. Statement executions 2_1 and 4_1 data depend on 1_1 because they use the file descriptor defined at 1_1 . 4_1 data depends on 2_1 because 2_1 changes the file seek pointer. The graph on the right hand side shows the EDG. Event execution $E3_1$ depends on $E2_1$ because of the dependence path $2_1 \rightarrow 4_1$. Event execution $E4_1$ depends on $E2_1$ due to the dependence path $2_1 \rightarrow 3_1 \rightarrow 5_1 \rightarrow 6_1$. Note that the *read* events $E2$ and $E3$ are considered as different events because they occur at different program locations.

Control dependence between statements can also lead to dependence between events as demonstrated by another example in Figure 4, where event $E3_1$ depends on event $E2_1$ as the result of 30_1 control depending on 21_1 and 21_1 data depending on 20_1 . The dependence between $E2_1$ and $E3_1$ belongs to control dependence as the execution of $E3_1$ is due to the result of $E2_1$. However, in EDGs we do not distinguish data dependence and control dependence edges.

Precisely constructing the EDG requires accurately tracing each data/control/potential dependence. According to our experience, exactly tracing each data/control dependence on the fly triggers the slow down of up to two orders of magnitude. Potential dependence is even more expensive to trace hence it is usually implemented as a post-mortem analysis. Thus, building precise EDG is a luxury that becomes worthy only when the cost can be amortized by a large number of replays. Otherwise, programmers would rather

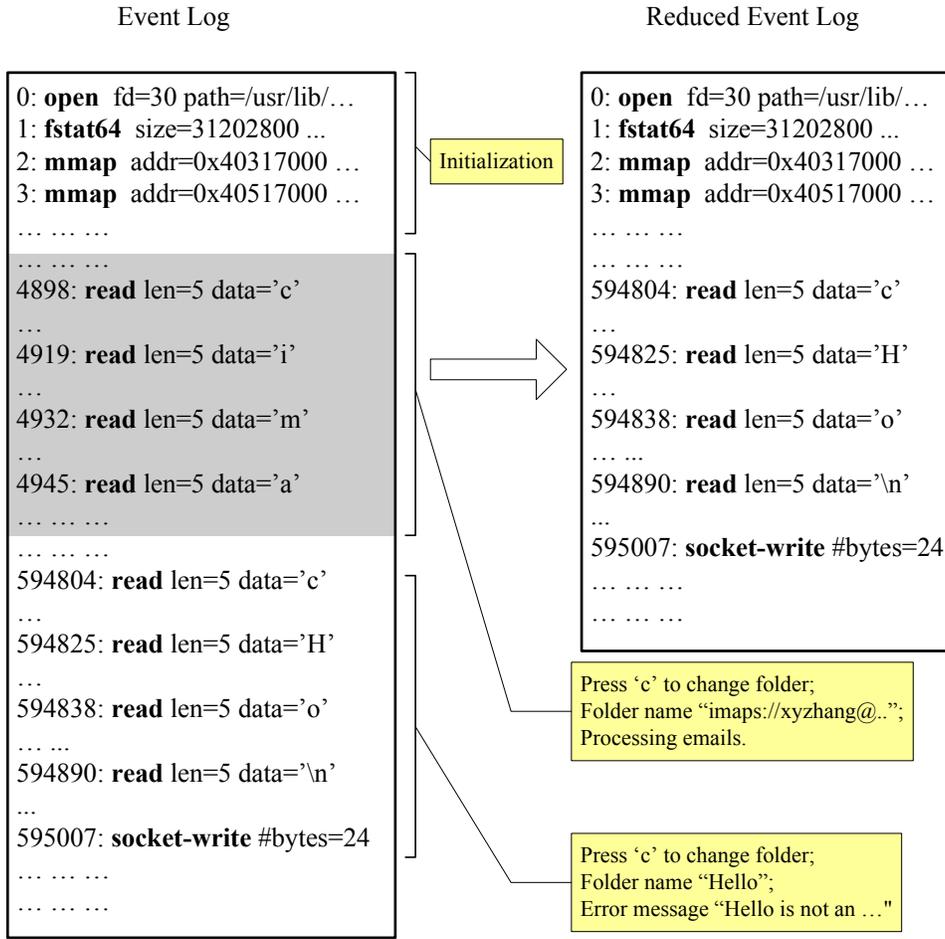


Figure 2: Getting the same warning message by replaying the reduced log for Mutt 1.4.2.1i. The numbers mean the byte positions of the corresponding events in the log.

replay the entire log, which is equivalent to doubling the execution time, than endure the two orders of magnitude slow down in the first place and attain speed up in just a few replays later on. To address this issue, we have to be conservative by constructing an approximate EDG, in which one event depends on the other if and only if they are related by a *static* dependence path. In other words, we only demand a static dependence graph, instead of a dynamic one, together with the event log to build the approximate EDG. The only runtime overhead we pay is for event logging, which is significantly cheaper than tracing each dependence. Because the dependences between events are usually much simpler than the dependences between normal statements, which can be highly complicated due to pointer aliasing, being conservative in EDG construction introduces much less imprecision compared to being conservative in building DDG.

2.2 Meta Slicing on Event Log

Similar to dynamic slicing, given an EDG and an event, which the programmer wants to reproduce, meta slicing on the EDG computes the set of events that are needed in order to replay the given event.

DEFINITION 3. Given $EDG(N, E)$, an event dependence graph, the **Meta Slice** of $e_i \in N$ denoted by $MS(e_i)$ is the subgraph of $EDG(N, E)$ which includes e_i as well as all other nodes and edges

from which e_i is reachable, i.e.

$$MS(e_i) = (\{e_i\}, \{e | e = m_j \rightarrow e_i \in E\}) \cup \bigcup_{\forall m_j \rightarrow e_i} MS(m_j)$$

For example in Figure 3, $MS(E_{41}) = \{E_{11}, E_{21}, E_{41}\}$. Note that we ignore the edges in MS for simplicity. We need to replay E_{11} , which opens the file, and E_{21} , which reads some data from the file, in order to correctly replay E_{41} , which prints some value resulted from the computation over the inputted data. In Figure 4, $MS(E_{31}) = \{E_{11}, E_{21}, E_{31}\}$. E_{21} has to be replayed otherwise the control would not flow to E_{31} .

We have discussed how to find the set of relevant events in order to replay a given event. However, in reality it could be a specific executed statement n_j that the programmer wants to replay. In this case, we need to find out the set of closest events reachable from n_j in the DDG, denoted as $ECut(n_j)$, and then compute the meta slices on these events. For example in Figure 4, $ECut(40_1) = \{20_1\}$, the corresponding meta slice $MS(20_1) = \{10_1, 20_1\}$. Intuitively, both E_{11} and E_{21} need to be replayed in order to replay statement S_1 .

THEOREM 1. The events in $MS(ECut(n_j))$ are sufficient to replay n_j .

Proof Let us assume there is an event e_x not in $MS(ECut(n_j))$,

```

101   ...
      inFd = open (path1, "r"); E11
      ...
201   n = read (inFd, buf, size); E21
211   if (n!=size) {
      ...
301   inFd = open (path2, "r"); E31
      ...
401   S1;
      ...
      }
      ...

```

Figure 4: Another example of Event Dependence Graph.

and e_x needs to be replayed in order to replay n_j . We infer there must exist an executed statement, event or non-event, m_i s.t. n_j is reachable from m_i and m_i is reachable from e_x . In other words, n_j directly/indirectly depends on m_i and m_i directly/indirectly depends on e_x . Otherwise, executing n_j would not require executing e_x . If there is no executed events along the path $e_x \rightarrow m_i \rightarrow n_j$ other than e_x , $e_x \in \text{ECut}(n_j)$, which is contradictory to the assumption; if there exists some executed event other than e_x along the path, let us assume e'_y is the executed event closest to n_j on the path s.t. $e'_y \in \text{eCut}(n_j)$, $e_x \in \text{MS}(e'_y)$ according to the definitions of EDG and meta slicing. It is a contradiction to the assumption. This completes the proof.

Note that in practice ECut has to be conservatively computed as we do not have dynamic dependence information. Our experience shows that it is not a problem because the events in ECut tend to be very close to the desired statement instance in the dependence graph such that very limited number of spurious dependences are brought in during the computation of ECut .

2.3 Replaying with A Reduced Event Log

We have described how meta slicing can be applied to identify the set of events in the log that are relevant to replaying given part of the execution. However, meta slicing is not yet the ultimate solution. It is often the case that a meta slice can not be used directly to drive the replayed execution. For example, in Figure 3, $\text{MS}(E4_1) = \{E1_1, E2_1, E4_1\}$. Replaying with the meta slice fails because $E3_1$ was expected when the control flows to statement 4_1 . This suggests that some events, even though irrelevant to replaying the desired part of the execution, cannot be pruned due to the control flow structure. In this subsection, we are going to describe how an event log is reduced with regard to the meta slice and the intrinsic control flow structure of the application.

Before we present the algorithm, let us first study an example that clearly explains how it is made possible to reduce a log without losing the validity. In Figure 5, the program displayed in the left column takes user commands from *stdin*. Different actions are taken based on different commands. For instance, messages are printed on the screen if 'a'/'c' is pressed; a file is opened if 'o' is pressed; the opened file is read if 'r' is read; if the data read does not match the size required, an error message is delivered. The event log for a particular execution is presented in the right column. During the execution, a file is opened and then read for twice; the second read does not satisfy the size wanted such that an error message is printed at 94₁; in between of these events, a number of events happen as the results of 'a'/'c' being pressed. Let us assume 94₁ is the event we want to replay. $\text{MS}(94_1)$ is denoted as the shaded events in the log. Apparently, the meta slice is not legit-

	Event Log
5	gettimeofday()
...	5 ₁ gettimeofday
10	while (1) {
...	20 ₁ getchar
20	switch (c = getchar()) {
...	31 ₁ printf ("..A..")
30	case 'a':
31	printf ("case A\n");
...	20 ₂ getchar
50	case 'c':
51	printf ("case C\n");
...	80 ₁ open
80	case 'o':
81	fd = open (... , "r");
...	20 ₃ getchar
90	case 'r':
91	n = read (fd, buf, size);
92	if (n!=size) {
93	gettimeofday()
94	printf ("Error: ...\n");
...	51 ₁ printf ("..C..")
...	20 ₄ getchar
...	51 ₂ printf ("..C..")
...	20 ₅ getchar
...	91 ₁ read
...	20 ₆ getchar
...	51 ₂ printf ("..C..")
...	20 ₇ getchar
...	51 ₃ printf ("..C..")
...	20 ₈ getchar
...	91 ₂ read
...	93 ₁ gettimeofday
...	94 ₁ printf ("Err...")

Figure 5: An example on reducing the event log. The shaded events are those in $\text{MS}(94_1)$.

imate for replay as event 5₁ (*gettimeofday*), which is not in the meta slice, is expected at the beginning of the replayed execution. While 5₁ is not removable, events 20₁ and 31₁ can be removed without any problem. The important observation here is that 20₂ and 20₁ are *compatible* and thus 20₂ can be moved up to replace 20₁ such that the event in between, 31₁, is pruned.

DEFINITION 4. An event execution e_i is compatible with another event execution e_j iff their calling contexts are identical and they occur at the same program point.

All the events 20_x in Figure 5 are compatible to each other. This example suggests we are able to alter the replayed execution by replacing an event with its compatible peer.

The algorithm to reduce a log given the meta slice is presented as follows. `Get_next_event()` gets the next event from the log file; `get_next_marked_event()` gets the next event belongs to the meta slice, which we assume is precomputed, in the log file. These two methods share the same file seek pointer, which can be set by `set_file_pointer(...)`.

```

Input: the original log Log
Output: the reduced log RLog
Initialize: RLog ← ϕ
while (em=get_next_marked_event(Log))!=EOF do
  e=get_next_event(Log)
  for each et from e to em in Log do
    if et.context ≡ em.context then
      goto L1
    endif
    Rlog ← Rlog · et
  endfor
L1:
  Rlog ← Rlog · em
  set_file_pointer(Log, em)
endwhile

```

The basic idea of the algorithm is that given a marked event e_m , an event in the meta slice, we find the earliest compatible event e_t in between e and e_m s.t. moving e_m up to replace e_t maximizes the savings. All the events between e and e_t including e are copied to

the new log to satisfy the control flow structure confinement. The events between e_t and e_m are discarded.

Table 1 presents the reduction procedure of the example in Figure 5. As shown in the table, during iteration one, 5_1 is the first event retrieved from the log, and 20_2 is the first marked event. 20_2 can be moved up to replace 20_1 such that 5_1 and 20_2 are the two events appended to the new log. During the second iteration, 80_1 is the next event and also the next marked event such that it is simply copied to the new log. In iteration three, moving 20_5 up to replace 20_3 results in cutting the events from 20_3 to 50_1 . The final reduced log is shown in the last row of the table. The reduce log can be used to drive the replayed execution to reproduce the error message at 94_1 .

Table 1: Computation table for figure 5.

Iteration	e	e_m	RLog
1	5_1	20_2	$5_1 20_2$
2	80_1	80_1	$5_1 20_2 80_1$
3	20_3	20_5	$5_1 20_2 80_1 20_5$
4	91_1	91_1	$5_1 20_2 80_1 20_5 91_1$
5	20_6	20_8	$5_1 20_2 80_1 20_5 91_1 20_8$
6	91_2	91_2	$5_1 20_2 80_1 20_5 91_1 20_8 91_2$
7	93_1	94_1	$5_1 20_2 80_1 20_5 91_1 20_8 91_2 93_1 94_1$

2.4 Dynamic Slicing during Replaying

Dynamic slicing was believed to be too expensive to apply for long executions. With sophisticated compression techniques [17] we can achieve the space efficiency of four bits per executed instruction, which is still not powerful enough for executions that run for minutes, hours, or days. The EFF technique can reproduce the failure without going through most of the irrelevant part of the execution. As a result, dynamic slicing becomes feasible for the fast forwarded executions. According to our previous study [18], dynamic data slicing, in which slices are computed by considering only the data dependence, is quite effective for memory type of bugs. Therefore, we only compute data slices in this paper due to the fact that crashes are usually the type of bugs reported for long running programs. In the rest part of the paper, we mean dynamic data slices when we mention dynamic slices. Note that the dynamic slicing in this phase is different from the meta slicing mentioned earlier: meta slicing is performed on the event dependence graph and generates a reduced log; dynamic slicing is performed on the statement level dynamic dependence graph which is constructed during the fast forwarded replay.

3. THE EFF SYSTEM

As we mentioned earlier, tracing can handle the execution of up to a few seconds, whereas checkpoints are usually created in the intervals of minutes. The ultimate goal of EFF is to fill the gap between tracing and checkpointing such that dynamic slicing can be applied. We have discussed how EFF fast forwards an execution from the beginning by replaying a reduced log. However, there is nothing fundamental that prevents EFF from being applied to executions resumed from checkpoints. Therefore, in this section we are going to describe how EFF, checkpointing, and tracing are integrated together. The composed system can be used to debug long running programs.

The system is presented in Figure 6. It consists of four components: *dynamic instrumentation* component, whose primary duty is to provide the infrastructure for tracing; *logging/checkpointing* component; *slicing* component; and the *EFF* component. The sys-

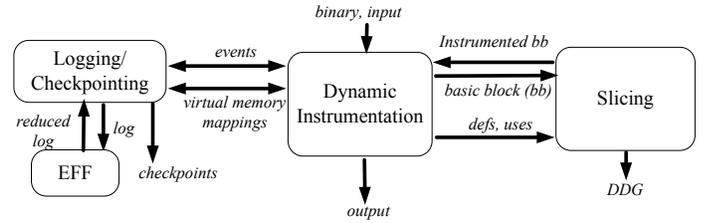


Figure 6: System infrastructure.

tem works as follows. In the original run, the slicing component is deactivated to reduce runtime overhead. The dynamic instrumentation engine traps each system call and forwards it to the corresponding handler in the logging model, which in turn logs the event. Checkpoints are created occasionally until a crash happens. In the replayed run, a smaller log file produced by EFF is supplied to drive the replay; in the mean time, the slicing component is turned on to trace the exercised data dependences till the crash point. The constructed dependence graph is studied to identify the root cause of the crash.

Dynamic Instrumentation The dynamic instrumentation engine is adapted from *valgrind* [14], which takes a x86 binary and before executing any new (never instrumented) basic blocks it calls the instrumentation function provided. The instrumentation function instruments the given basic block and returns the new basic block to *valgrind*. *Valgrind* 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). In our case, we can easily turn on/off the slicing component for the sake of performance or for certain part of the code, e.g. library code.

Logging/Checkpointing Our logging component is modified from *jockey* [13], which is an industry-strength checkpointing/replaying library executed in the application’s space. Compared to the checkpointing techniques executed in the kernel space, *jockey* has superior usability. Setting `LD_LIBRARY_PRELOAD=libjockey.so` is the only command required to activate *jockey*. Once loaded in, *jockey* calls the initialization method before the application gains control. During the initialization, *jockey* scans through all the binaries including the libraries loaded by the application, looking for any system call sites. Those system calls are redirected to *jockey* in order to log the corresponding events or, during the replay, retrieve the events from the log file without actually passing them on to the OS. Checkpoints can be created by setting a timer, such that the application is not even aware of the existence of *jockey*, or by making a library call to *jockey* inside the application. In the latter case, the application has to include *jockey*’s header file and be explicitly linked with the library. On receiving a checkpointing request, *jockey* obtains the layout of the application’s virtual space, which is *jockey*’s space as well, by parsing `/proc/self/maps`. A checkpoint is created by dumping all the virtual memory segments that do not belong to *jockey*.

Slicing The slicing component is inherited from our prior work [18]. The main difference is that we do not trace control dependence in this system because according to our study [18] tracing only data dependence is powerful enough to capture the root causes of memory bugs, which are the ones usually reported for long running programs. Another difference is that we augment the com-

ponent such that it stops at the execution points where an illegal memory access occurs, for example writing to an unallocated address. These points are usually earlier than the actual crash points.

EFF The *EFF* component implements the technique described in previous section. It takes an event log file dumped by the logging model and then computes the meta slice for a given set of events. The meta slice is used to prune the event log. The resulting smaller event log is used to drive the replayed execution. The computation of meta slice requires a static dependence graph, which is created by profiling the executed dependences in a few profiling runs due to the lack of an implementation of points-to analysis for x86 binaries.

One of the greatest challenges is to integrate logging/checkpointing model into the dynamic instrumentation engine. The integration is very meaningful because of the following reasons. Dynamic instrumentation is becoming more and more widely used in recent years. Not only is it attractive for the purpose of adaptive profiling/tracing, but also performance improvement can be achieved by executing a regular application on a dynamic instrumentation engine. Binary translation, a very promising technique that is derived from dynamic instrumentation, can virtually execute an architecture specific binary on a different architecture. Logging/checkpointing, on the other hand, has already been very popular for fault tolerance, debugging, etc. We believe logging/checkpointing should become a standard functionality of a dynamic instrumentation infrastructure in the near future. Therefore, the issues we are addressing here may be general to the integration of tools with similar functionalities. The first issue is the separation of the virtual space. Both *valgrind* and *jockey* are residents in the application’s space. They both assume total control over the entire virtual space such that they reserve certain address space for their own purposes. The reservations conflict each other. For instance, *jockey* reserves 0x7200000-0x7800000 for its heap, the mapping of the log file, etc. The same address space is also reserved by *valgrind* for tracing. Our solution is to make them aware the existence of each other by separating the application’s address space into two parts – the *valgrind*’s space and the *jockey*’s space. The application is actually executed within the *valgrind*’s space. The second issue is about adjusting the system call trapping mechanism in *jockey*. *Jockey* traps system calls by directly overwriting the application’s code. As a result, *valgrind* traces into *jockey* and tries to instrument the *jockey* code, which is undesirable. Our solution is to avoid any direct interactions between the application code and the *jockey* code. *Jockey* can only interface with *valgrind*. More precisely, we use *valgrind* to trap the system calls and then call the *jockey* event handlers inside *valgrind*. The third issue is to discretely checkpoint the execution. A naive solution only checkpoints the application’s program status. The reality is that the application’s program status is so mixed up with the *valgrind*’s status that *valgrind* fails to resume from the checkpoints during replay if only the application’s status has been checkpointed. Our solution is to treat the *valgrind*’s status as part of the application’s status such that it is checkpointed as well. Some of the *valgrind*’s status should be excluded such as the *valgrind*’s log file descriptor, which should be reopened at the beginning of the replay. There are some other minor issues in order to make both *valgrind* and *jockey* run correctly such as some of the *valgrind*’s sanity checks have to be turned off. We are not able to cover all these issues due to the space limit.

4. EXPERIMENTATION

The evaluation was hard to carry out. The first issue is that what benchmarks we should use. The programs we select should be able to run for a long time. We looked at the set of bugs studied in [6,

10, 9] and picked the programs that can execute for a long time. Table 2 presents the set of programs we selected. Most of them are user interactive programs. We ignored *apache* since *apache* creates multiple processes while our logging model can handle only one process at the current stage. The second issue is that we need the input that can drive the execution for a long time and then crash the execution. On the other hand, the execution should not be so long that it becomes too heavy a task for us to collect the data. Unfortunately, the input coming with the selected bugs usually leads to very short executions. Given the fact that most benchmarks are interactive, we constructed a long input by first performing a lot of user actions and then apply the failure inducing input – the input comes with the benchmarks. For example in *mutt*, we took the following actions: (i) opening an email account; (ii) going through all the emails one by one, the total is about six hundreds; (iii) trying to switch to an invalid folder; repeating steps (ii) and (iii) two more times; providing the failure inducing input and crashing the program. We collected the user time as the performance indicator since the real time may significantly differ each time depending on the user’s behavior.

Table 2: Description of the benchmarks

Benchmark	Description	LOC	Bug Type
bc-1.06	interactive calculator	14.4K	heap overfbw
mc-4.5.55	file manager	86.2K	stack overfbw
mutt-1.4.2.li	email client	453.6K	heap overfbw
pine-4.44	email client	211.9K	stack overfbw
pine-4.44	email client	211.9K	heap overfbw
squid-2.3	web proxy cache server	93.5K	heap overfbw

We investigated four execution scenarios: *orig.* denotes the original execution; *traced* denotes the original execution plus the dependence tracing; *logged* represents the original execution plus logging; *EFF* represents the fast forwarded execution plus the dependence tracing. In the logged run, an event log is created. The *EFF* technique is applied to reduce the log. The statement instance we want to replay is where the crash happened. The *EFF* technique is able to reproduce the crash in a much shorter execution. Due to the complexity of our system, our implementation is not sound at the current stage. Some times we have to hard code a few event dependences, otherwise the reduced log is not valid to drive the replay which is manifested as an event missing when it is expected or the presence of an extra event. We expect to have it fixed in the camera ready version if the paper gets accepted.

Table 3: Performance comparison of different execution scenarios.

Benchmark	Orig. (sec.)	Traced (sec.)	Logged (sec.)	EFF (sec.)
bc-1.06	13.6	2040.4	16.2	0.05
mc-4.5.55	11.3	499.3	-(1)	-
mutt-1.4.2.li	19.7	3237.7	26.1	0.06
pine-4.44(stack)	14.4	2088.4	36.8	0.12
pine-4.44(heap)	13.9	2102.2	34.4	0.20
squid-2.3	14.6	1131.6	25.6	0.17
Benchmark	Traced/Orig.	Logged/Orig.	Traced/EFF (sec.)	
bc-1.06	150.6	1.19	40808.8	
mc-4.5.55	44.5	-	-	
mutt-1.4.2.li	164.5	1.32	53960.8	
pine-4.44 (stack)	145.1	2.55	17403.6	
pine-4.44 (heap)	151.5	2.47	10510.9	
squid-2.3	77.3	1.75	6656.4	

(1) *mc* crashed our system for some reason that we are still investigating.

Table 3 compares the performance under the four scenarios. We

can see the original runs, which were terminated by crashes, consume user time ranging from 11.3 to 19.7 seconds, which corresponds to the real time of a few minutes. They are not long by simply looking at the absolute numbers, but they well exceed the capability of our dependence tracing technique. We can easily extend the executions by repeating the user actions. The side effect is the increased difficulty of collecting the time for the executions in the *traced* scenario. Note that even though checkpointing is supported in our system, the original execution does not last long enough to trigger it. Fortunately, it does not affect the evaluations of the EFF technique and the effectiveness of dynamic slicing on long running programs. From table 3, we have the following observations.

- Dependence tracing introduces 44.5 to 164.5 times slow down. A programmer may bear it for a short run but highly unlikely for a long run.
- The slow down factors for logging range from 1.19 to 2.55, which are significantly smaller than the tracing slow down factors. For user interactive programs, the overhead is not noticeable.
- EFF can greatly shorten the executions such that dependence tracing becomes bearable.

Table 4: Comparison of the event logs.

Benchmark	# of events in Orig.	# of events in EFF	Orig./EFF
bc-1.06	340509	7	48644.0
mc-4.5.55	-	-	-
mutt-1.4.2.1i	262559	489	536.9
pine-4.44	7365830	3028	2432.6
pine-4.44	8707316	27279	319.2
squid-2.3	1620988	795	2038.9

Table 5: Comparison of the dependence graphs.

Benchmark	# of dep. in Orig.	# of dep. in EFF	Orig./EFF
bc-1.06	2.18×10^{10}	4.9×10^5	44489.8
mc-4.5.55	0.67×10^{10}	-	-
mutt-1.4.2.1i	4.86×10^{10}	4.21×10^7	1154.4
pine-4.44	1.95×10^{10}	2.68×10^7	727.61
pine-4.44	2.78×10^{10}	1.55×10^8	179.4
squid-2.3	1.1×10^{10}	1.93×10^6	5699.5

Table 4 compares the numbers of events before and after event reduction. We can see the reduction factors range from 319.2 to 48644.0, which well explain why the fast forwarded executions become so short. Table 5 presents the numbers of the exercised data dependences in the original and the fast forwarded executions. We want to point out that these numbers are collected after the intra-basic-block optimization [17] which eliminates considerable redundant dependences. We can tell that the numbers for the fast forwarded executions are much smaller. The constructed dependence graphs can be stored even without further compression [17].

5. DYNAMIC SLICING ON A SET OF LONG RUNNING BUGS

The performance of a set of long running bugs has been studied in the last section. As the original motivation, dynamic slicing is applied and evaluated to show the effectiveness.

5.1 Mutt

Mutt [21] is a text based mail user agent (MUA) for Unix based Operating Systems. It has many features including customizability,

POP3 and IMAP support, and ability to handle multiple mailbox formats. According to [22], mutt version 1.4 has a known memory bug which is as follows. The Mutt Mail User Agent (MUA) has support for accessing remote mailboxes through the IMAP protocol. When mutt has to convert the name of the folder from its internal UTF-8 representation to UTF-7 it calls the function *utf8_to_utf7* in module *imap/utf7.c*. When this function does the conversion, it miscalculates the length of the output string. To conduct our experiment, after Mutt is executed for a long time, we supply a UTF-8 folder name that contains some special characters. The heap buffer is overflowed and a segmentation fault is flagged. We reduce the event log using EFF and then replay the execution with the new log. Dynamic slicing is activated in the replayed execution. Figure 7 shows the computed dynamic slice.

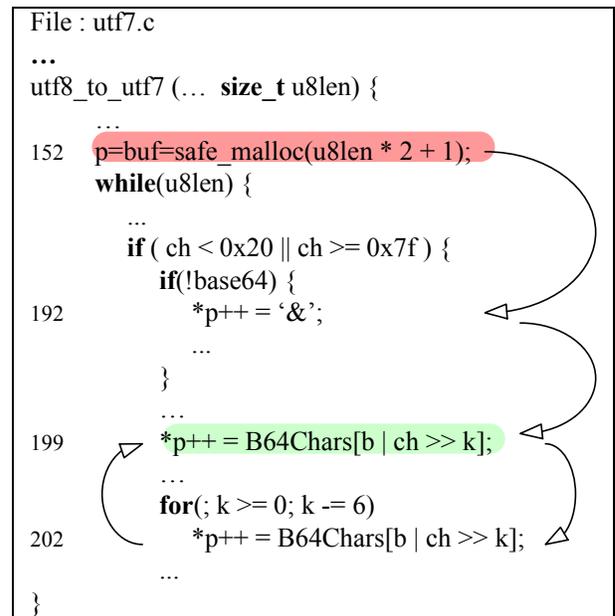


Figure 7: Mut 1.4.2.1i

As we mentioned earlier, our slicing component also monitors for any attempt for illegal memory access. After detecting a write to a memory region not allocated at line number 199, we now inspect the data slice to find the root cause. We find that the last instance of line 199 is data dependent on line 202 and vice-versa through variable 'p'. The arrows indicate the data dependence. The data dependence chain in the slice leads us to the first instance of line 199 which is data dependent on line 192 and this in turn is data dependent on line 152, because of variable 'p', which is the root cause of the bug as there is an error in calculating the buffer length at this point. We needed to inspect just 8 static statements before getting to the root cause, and the dependence chain provides a very clear explanation on the cause effect relations.

5.2 Pine

Pine [25] is a popularly used application for reading, sending and managing email messages and is distributed with the Linux operating system. Pine version 4.44 has two buffer overflow errors. One is a stack overflow and the other is a heap overflow. We look at both errors in the following subsections.

5.2.1 Pine Stack Overflow

According to [26] pine has a stack overflow error. Pine calls an error prone API when it accesses mailboxes. By asking pine to handle a mailbox that has some special characters this bug can be triggered causing pine to crash.

We are able to capture the root cause of the bug again using dynamic slicing. Our tracking infrastructure reports an illegal memory access at line number 589 of file *mail.c*, where the statement is "for (...;(c=*t++)!='');" ". We look at the slice at this point and find that there is a loop carry self dependence. This line is also the root cause of the BUG as variable 't', which is the pointer to a string, is incremented beyond its allocated region (on stack) if the provided string does not have the end quote. We needed to inspect 3 static statements to nail the root cause.

5.2.2 Pine Heap Overflow

According to [27] pine has a bug that when triggered can overflow the heap memory causing a potential crash. This can occur when pine processes the "From" field of email headers. Certain special characters in the header can cause the bug. Figure 8 shows the code where the bug is present.

```

File : bldaddr.c
int est_size(a) {
...
7269 cnt += ...
...
return(max(cnt,50));
}

File : bldaddr.c
char *addr_list_string(...) {
...
7126 list = (char *) fs_get(...est_size(adrlist));
...
7128 rfc822_write_address_decode(list, ...);
}

File : rfc822.c
void rfc822_cat (char *dest, ...) {
...
dest += strlen(dest);
*dest++ = "";
...
for(;s = strpbrk (src,"\\"); ...) {
strncpy (dest, ...);
dest += i;
260 *dest++ = '\\';
*dest++ = *s;
}
}

```

Figure 8: Pine 4.44 heap overflow.

There is an illegal heap access detected by our infrastructure at line number 260 in file *rfc822.c*. However, the root cause of the bug is at line number 7269 of file *bldaddr.c*. The buffer *dest* in *rfc822_cat* is allocated in *addr_list_string*. The size of the alloca-

tion is miscalculated in *est_size* because it does not consider special characters. The figure shows the dependences that we tracked to get to the root cause from the error point. This is an example where the root cause and the symptom are in different functions. We had to examine 10 static code statements to get to the root cause.

5.3 Midnight Commander

Midnight Commander{mc} [23] is an open source file manager for free operating systems. It has high degree of portability and can be compiled and run on a number of operating system including Linux. We used mc version 4.5.55 for our experiment. This version has a known buffer overflow error. According to [24], the bug is triggered when midnight commander is used to process symbolic links in *tgz* archives. Absolute symbolic links in the archives are translated into links relative to the start of the *tgz* file. The buffer that is used to form the relative link is never initialized and hence can be overflowed inside the *strcat* procedure. Unfortunately, our system failed in logging the extended mc execution and thus EFF can not be applied. We used the failure inducing input only to conduct the study. Figure 9 shows the code corresponding to the bug.

```

File : direntry.c
vfs_s_entry * vfs_s_resolve_symlink(...) {
char buf[MC_MAXPATHLEN], *linkname;
...
for(;;p++) {
...
if(!p) {
385 strcat(buf,q);
break;
}
}
...
398 return (MEDATA->find_entry) (...);
}

```

Figure 9: Mc 4.5.55

We use our infrastructure to determine the root cause of the bug. A segment fault occurs at line number 398. Now, when we look at the slice at this point we find an abnormal data dependence between line 398 and line 385. We conclude that a stack buffer overflow happened at line 385, which is the root cause of the bug, such that it corrupted one of the variables used at line 398. We just needed to inspect 2 static statements to get to the error.

5.4 Squid

Squid [28, 6] is a fully featured web proxy cache that supports proxying and caching of HTTP, FTP and other URLs. It is designed to run on Unix based systems. We use squid version 2.3 for our experiment. It has a known heap buffer overflow error. When an input request contains some special characters, squid miscalculates the length of the heap buffer that is used to hold the request. As a result, the buffer is overflowed and then the server crashes. [29] explains it in more details. Figure 10 shows the portion of the code that contains the bug.

On running squid using our infrastructure we find that there is a heap buffer overflow at line number 1024. Inspecting the slice at this point leads us to the root cause of the bug at line number 1005, at which the extra padding space of size 64 is not enough to

```

File : ftp.c
static void
ftpBuildTitleUrl(FtpStateData * ftpState) {
    ...
1005 len = 64
1006     + strlen(ftpState->user) ...
    ...
1021 t=xcalloc(len,1);
    ...
    if(strcmp(...)) {
1024     strcat(t,rfc1738_escape_part(ftpState->user));
    ...
}

```

Figure 10: Squid 2.3

accommodate the special characters. We had gone through 5 static statements before we reached the root cause.

5.5 bc

Bc [1] is a numeric processing language that supports arbitrary precision numbers. It is generally distributed along with the Linux operating system and is a part of the GNU project. We used bc-1.06 for our experiment. This version has a known heap overflow error.

[6] describes the bug that is triggered when bc is used. A certain heap buffer is not declared wide enough and overflows. The code corresponding to the error is shown in Figure 11. The heap array *arrays* declared at line number 167 is overflowed.

```

File : storage.c
void
more_arrays() {
    ...
167 arrays=(bc_var_array **) bc_malloc(
    a_count * sizeof(bc_var_array *);
    ...
176 for(; indx < v_count; indx++)
177     arrays[indx] = NULL;
    ...
}

```

Figure 11: Bc-1.06

Our tracking infrastructure detects a heap memory violation at line number 177. Looking at the slice at this point we see that the root cause of the bug is at line number 167. This is because *a_count* entries have been declared but *v_count* entries are accessed. We needed to inspect just these 3 statements to find the root cause.

From the studies we find that the microthese bugs are not as mysterious as they appear, under the micro-inspection of dynamic slicing. They usually require examining a few static statements before the root cause is located. Two conclusions can be drawn: dynamic slicing is very effective to handle memory type of bugs even in the long running programs examined; the real challenge is to isolate the part of the execution that is relevant to the error and hence dynamic slicing can be applied. The EFF technique is designed for the purpose. According to our experience in [18, 3, 19], most non-memory bugs still have very good locality even though not as

apparent as memory bugs. We firmly believe EFF plus dynamic slicing will still be highly effective for non-memory bugs in long running programs. Unfortunately, most bugs that are reported and studied for those programs are memory bugs. We plan to mine some software repositories of long running programs to get more interesting non-memory bugs in the future.

6. CONCLUSIONS

We have enabled dynamic slicing on a set of long running programs by developing a novel execution fast forwarding technique. Fast forwarding can be achieved by driving the replay with a reduced event log file. Given a desired execution region, a large portion of the events are not relevant to replaying it. Meta slicing is designed to eliminate this redundancy in the log file. With the execution fast forwarding technique, the replayed execution becomes substantially shorter and yet the wanted execution region is precisely reproduced. Hence dynamic slicing can be practically applied to isolate the cause effect chain leading to the failure. Our studies show that most of the reported memory bugs for long running programs are trivial to locate with dynamic slicing once the execution has been shortened to an affordable level.

7. REFERENCES

- [1] GNU bc. <http://www.gnu.org/software/bc>
- [2] S. Bhansali, W-K. Chen, S. de Jong, A. Edwards, R. Murray, M. Drinic, D. Mihocka, and J. Chau, "Framework for instruction-level tracing and analysis of program executions," *Virtual Execution Environments Conference*, Ottawa, Canada, June 2006.
- [3] N. Gupta, H. He, X. Zhang, and R. Gupta, "Locating Faulty Code Using Failure-Inducing Chops," *20th IEEE/ACM International Conference on Automated Software Engineering*, pages 263-272, Long Beach, California, Nov. 2005.
- [4] T. Gyimothy, A. Beszedes, I. Forgacs, "An efficient relevant slicing method for debugging," *7th European Software Engineering Conference/ 7th ACM SIGSOFT International Symposium on Foundations of Software Engineering*, Toulouse, France, 1999.
- [5] B. Korel and J. Laski, "Dynamic program slicing," *Information Processing Letters*, Vol. 29, No. 3, pages 155-163, 1988.
- [6] S. Lu, Z. Li, F. Qin, L. Tan, P. Zhou, and Y. Zhou, "BugBench: a benchmark for evaluating bug detection tools", *Workshop on the Evaluation of Software Defect Detection Tools*, 2005.
- [7] R.H.B. Netzer and M.H. Weaver, "Optimal Tracing and Incremental Reexecution for Debugging Long-Running Programs", *ACM SIGPLAN Conference on Programming Language Design and Implementation*, Orlando, FL, USA, pages 313-325, June 1994.
- [8] D.Z. Pan and M.A. Linton, "Supporting reverse execution of parallel programs," *ACM workshop on parallel and distributed debugging*, Madison, WI, USA, May 1988.
- [9] F. Qin, J. Tucek, J. Sundaresan, and Y. Zhou, "Rx: treating bugs as allergies - a safe method to survive software failures", *the 20th ACM Symposium on Operating Systems Principles* Brighton, UK, pages 235-248, Oct. 2005
- [10] M.C. Rinard, C. Cadar, D. Dumitran, D.M. Roy, T. Leu, and W.S. Beebe, "Enhancing Server Availability and Security Through Failure-Oblivious Computing", *the Sixth Symposium on Operating System Design and Implementation* San Francisco, California, pages 303-316, 2004
- [11] M. Ronsse, K. De Bosschere, M. Christiaens, J.C. de Kergommeaux, and D. Kranzlmler, "Record/replay for nondeterministic program executions", *Communication of the ACM* 46(9), pages 62-67, 2003
- [12] M. Ronsse, K. De Bosschere, and J.C. de Kergommeaux, "Execution replay and debugging", *Fourth Workshop on Automated and Analysis-Driven Debugging*, Munich, Germany, August 2000.
- [13] Y. Saito, "Jockey: a user-space library for record-replay debugging", *Sixth International Symposium on Automated and Analysis-Driven Debugging*, Monterey, California, September 2005.
- [14] J. Seward et al. "Valgrind: A GPL'd system for debugging and profiling x86-linux programs", <http://valgrind.ked.org/>, 2004.

- [15] S.M. Srinivasan, S. Kandula, C.R. Andrews, and Y. Zhou, "Flashback: a lightweight extension for rollback and deterministic replay for software debugging", *USENIX Annual Technical Conference*, Boston, MA, USA, June 1994.
- [16] L.D. Wittie. "Debugging distributed C programs by real time replay," *ACM workshop on parallel and distributed debugging*, pages 57-67, Madison, WI, USA, May 1988.
- [17] X. Zhang and R. Gupta, "Whole Execution Traces," *IEEE/ACM 37th International Symposium on Microarchitecture*, pages 105-116, 2004.
- [18] X. Zhang, H. He, N. Gupta and R. Gupta, "Experimental evaluation of using dynamic slices for fault location," *Sixth International Symposium on Automated and Analysis-Driven Debugging*, Monterey, California, September 2005.
- [19] X. Zhang, N. Gupta, and R. Gupta "Locating Faults Through Automated Predicate Switching," *IEEE/ACM International Conference on Software Engineering*, Shanghai, China, May 2006
- [20] X. Zhang, N. Gupta, and R. Gupta "Pruning Dynamic Slices With Confidence," *ACM SIGPLAN Conference on Programming Language Design and Implementation*, Ottawa, Canada, June 2006
- [21] Mutt Website. www.mutt.org
- [22] Mutt Buffer Overflow.
<http://www.securiteam.com/unixfocus/5FP0T0U9FU.html>
- [23] Midnight Commander. www.ibiblio.org/mc
- [24] Midnight Commander exploit. www.securityfocus.com/bid/8658
- [25] Pine Website. www.washington.edu/pine/
- [26] Pine Stack Buffer Overflow Error.
<http://www.xatrix.org/advisory.php?s=7408>
- [27] Pine Heap Buffer Overflow Error.
<http://www.securityfocus.com/bid/6120>
- [28] Squid Website. <http://www.squid-cache.org/>
- [29] Squid Buffer Overflow.
<http://www.securiteam.com/unixfocus/5BPOP2A6AY.html>