

Automated Memory Leak Fixing on Value-Flow Slices for C Programs

Hua Yan*[†] Yulei Sui* Shiping Chen[†] Jingling Xue*
*School of Computer Science and Engineering, UNSW, Australia
[†]Digital Productivity Flagship (DPF), CSIRO, Australia

ABSTRACT

C is the dominant programming language for developing embedded software, operating systems, and device drivers. Unlike programs written in managed languages like Java, C programs rely on explicit memory management and are prone to memory leaks. Existing (static or dynamic) detectors only report leaks, but fixing them often requires considerable manual effort by inspecting a list of reported true and false alarms. How to develop on-demand lightweight techniques for automated leak fixing without introducing new memory errors remains challenging. In this paper, we introduce AUTOFIX, a fully automated leak-fixing approach for C programs by combining static and dynamic program analysis. Given a leaky allocation site reported by a static memory leak detector, AUTOFIX performs a graph reachability analysis to identify leaky paths on the value-flow slices of the program, and then conducts a liveness analysis to locate the program points for inserting fixes (i.e., the missing `free` calls) on the identified leaky paths. We have implemented AUTOFIX in LLVM-3.5.0 and evaluated it using five SPEC2000 benchmarks and three open-source applications. Experimental results show that AUTOFIX can safely fix all the leaks reported by a state-of-the-art static memory leak detector with small instrumentation overhead.

CCS Concepts

•Software and its engineering → Memory management; Software performance; Software reliability;

Keywords

Memory Leaks; Bug Fixing; Value-Flow Analysis

1. INTRODUCTION

The two central tasks facing programmers during software testing are finding bugs and fixing them. Both tasks are

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC 2016, April 04-08, 2016, Pisa, Italy

Copyright 2016 ACM 978-1-4503-3739-7/16/04...\$15.00

<http://dx.doi.org/10.1145/2851613.2851773>

| | |
|--------------------------------|--------------------------------|
| 1 Foo () { | 1 Bar () { |
| 2 char* p = malloc (...); // o | 2 char* p = malloc (...); // o |
| 3 char* q = "on stack"; | 3 fgets(p, ...); |
| 4 fgets(p, ...); | 4 if (C2) { |
| 5 free (p); ✗ | 5 Use (p); free (p); |
| 6 if (C1) { | 6 } |
| 7 char * t = p; | 7 else { |
| 8 p = q; | 8 Use (p); |
| 9 q = t; | 9 free (p); ✓ |
| 10 } | 10 } |
| 11 printf ("%s", q); | 11 } |
| 12 free (p); ✗ | 12 Use (char* x) { |
| 13 if (C1) free (q); | 13 printf ("%s", x); |
| 14 else free (p); ✓ | 14 free (x); ✗ |
| 15 } | 15 } |

(a) Fixing an intra-procedural leak (b) Fixing an inter-procedural leak

Figure 1: Incorrect (✗) and correct (✓) fixes.

expensive in dollar terms and time-consuming due to the ever-increasing complexity of modern software systems. A large number of existing program analysis and testing techniques focus on automatic bug detection. However, finding bugs is only the first step. Once reported, bugs must still be repaired. Manually fixing bugs can be a non-trivial and error-prone process, especially for large-scale software.

Recently, a few approaches on automatic bug fixing have been proposed to reduce maintenance costs by producing candidate patches for validation and repairing [1, 10, 14, 19, 21, 31]. For example, CLEARVIEW [21] enforces violated invariants to correct buffer overflow and illegal control flow errors by creating patches for binaries. AUTOFIX-E [31] relies on user specifications and generates repairs using contracts. PACHIKA [8] infers object behavior models to propose candidate fixes for bugs like null dereferences. GENPROG [10] uses genetic programming to repair bugs in legacy code.

Most of the existing automatic approaches for fixing bugs in C programs are related to spatial memory errors such as buffer overflows and null pointer dereferences. Such bugs can be validated by inserting assertions and repaired by adding conditional checks to avoid executing the code segment that leads to undesired behaviors (e.g., program crashes) [16].

Memory leaks represent another major category of temporal memory errors that are more complicated to fix automatically. Unlike a spatial error that can be fixed by adding a conditional check to bypass the point where the spatial error occurs, every leaky path from a leaky allocation site needs to be tracked by inserting an appropriate fix (i.e., a `free` call) without introducing new memory errors.

Figure 1 illustrates how intra- and inter-procedural leaks are fixed correctly and incorrectly. Suppose the memory allocated in line 2 in Figure 1(a) is never freed. Adding a fix, `free(p)`, too early in line 5 can cause a use-after-free error in line 11, but adding `free(p)` in line 12 at the end of program without considering path correlation may introduce an invalid `free` site for a stack object (when the `if` branch is executed). A correct fix is provided in lines 13 – 14, with the underlying path correlation accounted for correctly. Let us consider now an inter-procedural leak shown in Figure 1(b), where the memory allocated in line 2 is leaked partially along the `else` branch (lines 7 – 10). A simple-minded fix, `free(x)`, which is inserted in line 14 in the function `Use`, is incorrect. Without considering correlated calling contexts, this fix may introduce a double-free in line 5 when the `if` branch is executed. A correct fix is given in line 9, ensuring that only the leak along the `else` branch is fixed.

Existing (static and dynamic) memory leak detectors for C programs only report leaks, but fixing them along every leaky path remains to be done manually by programmers. Dynamic detectors [4, 18] find leaks by instrumenting and tracking memory accesses at runtime, incurring high overhead. By instrumenting a program under some inputs, dynamic detectors produce no false positives but potentially miss many bugs. In contrast, static detectors [5, 15, 26, 27, 32], which approximate runtime behaviors of the program without executing it, can soundly pinpoint all the leaks in the program, but at the expense of some false positives.

This paper presents AUTOFIX, a fully automated approach for fixing memory leaks in C programs by combining static and dynamic analysis. Given a list of leaky allocation sites reported by a static detector, AUTOFIX automatically fixes all the reported leaks by inserting appropriate fixes (i.e., `free` calls) along all the leaky paths. There are two main challenges. First, a detector reports a leaky allocation site as long as it discovers one leaky path from the site without necessarily reporting all the leaky paths. AUTOFIX is designed to fully repair the leak for all its leaky paths. Second, some reported leaks are false positives. AUTOFIX must guarantee memory safety by ensuring that the fixes are correct regardless of whether the reported leak is a true bug or a false positive. Note that AUTOFIX certainly cannot fix any leaks that are missed (i.e., not reported) by a static detector.

AUTOFIX applies to a large class of real-world C programs where memory management is explicitly orchestrated by programmers without resorting to garbage collection (GC) and/or reference counting (RC). Compared to the GC and RC approaches, our approach is lightweight as only small instrumentation overhead is incurred. To safely reclaim a leaked memory object o from an allocation site without any programmer intervention, all the memory allocation and deallocation sites reachable from o on the value-flow slices of the program are instrumented to keep track of the liveness of o in shadow memory, thereby enforcing correct leak fixing inside a memory-safe execution sandbox at runtime.

Figure 2 highlights the basic idea behind AUTOFIX. Given a leaked object o from an allocation site, AUTOFIX builds from the program a sparse value-flow graph (S1), on which a graph reachability analysis is first performed to locate the candidate functions for inserting appropriate fixes, i.e., `free`

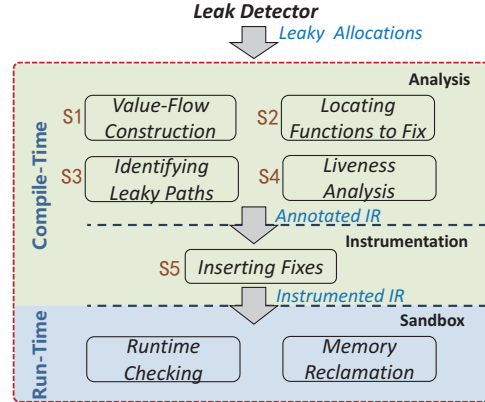


Figure 2: The AutoFix framework.

calls (S2). For each candidate function f , AUTOFIX then identifies the leaky paths in f for o by computing the value-flow guards with respect to its existing deallocation sites found in the program (S3). Next, a liveness analysis is applied inside f on the value-flow slice of the identified leaky paths for o to determine every program point P where a fix is needed with path correlation considered (S4). Finally, the fixes are inserted immediately after the last use sites of o on all its leaky paths (S5).

This paper makes the following contributions:

- We present AUTOFIX, a fully automated approach to memory leak fixing for C programs that can safely reclaim all the leaked objects reported by a leak detector.
- We propose a lightweight memory leak fixing approach by combining static and dynamic analysis so that all the missing `free` calls for a leaked object are first identified at compile-time and then inserted into the program to prevent the object from being leaked at runtime based on the shadow information recorded.
- We have implemented AUTOFIX in LLVM-3.5.0 and evaluated it using five SPEC2000 benchmarks and three open-source applications. Our experimental results show that AUTOFIX can safely fix all the leaks reported by the state-of-the-art static leak detector, SABER [26, 27], with runtime overhead averaged at under 2%. For the long-running server application `redis` evaluated, AUTOFIX has successfully reduced its memory usage by more than 300MB in a three-hour continuous run after having fixed its leaks.

2. A MOTIVATING EXAMPLE

We use an example in Figure 3 to describe our approach by going through its five key steps shown in Figure 2. An allocation site in line 4 of `fun` in Figure 3(a) is partially leaky along the `else` branch in line 10 inside a `while` loop. Given a leaked object o detected from this allocation site by a static leak detector, AUTOFIX first constructs a value-flow graph shown in Figure 3(b) for o . Based on this graph, AUTOFIX inserts one fix immediately after line 16 on the leaky `else` branch, as shown in Figure 3(c). The leaky allocation site

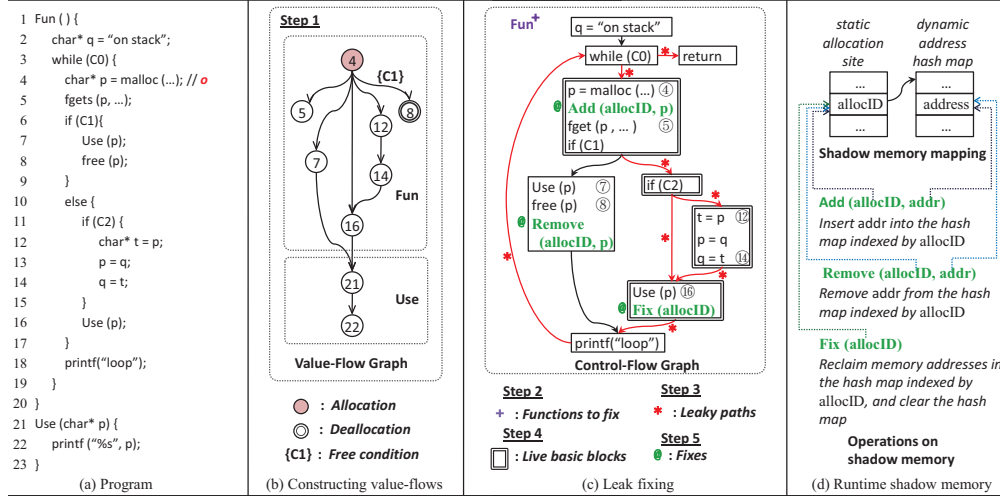


Figure 3: A motivating example.

(in line 4) and the deallocation site (in line 8) in the non-leaky `if` branch are instrumented with dynamic checks to ensure safe fixing at runtime, as illustrated in Figure 3(d).

Step 1: Constructing Value-Flows. Following [26, 27], we construct an inter-procedural value-flow graph (VFG) for every leaky allocation, with the one for the example given in Figure 3(b). Note that its nodes are numbered by their corresponding line numbers in Figure 3(a). In a VFG, its nodes represent the definitions of variables and its edges capture their def-use relations.

Step 2: Locating Functions to Fix. Given a leaked object o from an allocation site, AUTOFIX first determines the functions where the leaks of o should be fixed. A candidate function f is chosen if f allocates o (directly in itself or indirectly in its callee functions) such that o is never returned to any caller of f . Note that the existence of a candidate function for o is guaranteed since the `main` function will be the last resort. In our example, `Fun` is selected as a candidate function since it contains an allocation site of o (line 4) and o is never returned to `Fun`'s callers based on its value-flows.

Step 3: Identifying Leaky Paths. AUTOFIX identifies the leaky paths for o in `Fun` by reasoning about value-flow guards, which are Boolean formulas capturing branch conditions between defs and uses in the control flow graph (CFG). The *free* condition $C1$ under which the `free` site in line 8 is reached is computed by performing a guarded reachability analysis from the `malloc` source ④ to its `free` sink ⑧. Thus, the *leak* condition $\neg C1$ encodes the leaky paths in the `else` branch as highlighted in red in Figure 3(c).

Step 4: Liveness Analysis. An intra-procedural liveness analysis is performed for `Fun` to mark the live basic blocks for o (shown as double-framed boxes in Figure 3(c)) that are reachable from the allocation site of o on its leaky paths. As a result, node ⑩ is identified as the last-use site of o .

Step 5: Inserting Fixes. As shown in Figure 3(c), a deallocation `Fix()` is inserted immediately after ⑩ (i.e., line 16), where the last use of o is found. In addition, the instrumentation code (in green) also includes the metadata-manipulating functions inserted (after the `malloc` source ④

and the `free` sink ⑧) to maintain runtime shadows for o to ensure safe fixing for both leaky and non-leaky paths.

Figure 3(d) shows that the shadow memory simply maps an allocation site with its unique ID, `allocID`, to a hash map that records (start) addresses of the dynamically allocated objects that are not yet freed. Consider Figure 3(c) again. Every address p that points to an object allocated at ④ is recorded in the shadow memory by calling function `Add(allocID, p)` instrumented immediately after ④. The deallocation site ⑧, which is reachable from ④ via value-flows, is instrumented by calling function `Remove(allocID, p)` to delete the address p from the shadow hash map since its pointed-to object has been released along the non-leaky `if` branch. On reaching `Fix(allocID)` during program execution, all the objects identified by the addresses corresponding to the allocation site `allocID` in the shadow memory are freed (as they would be leaked otherwise).

3. AUTOMATED LEAK FIXING

AUTOFIX is a compile-time transformation for inserting runtime checks to reclaim leaked memory for C programs, which keeps track of potential leaked memory addresses via an efficient shadow metadata structure. In this section, we first present the five steps of AUTOFIX's compile-time transformation (§3.1 – §3.5), and then describe the design of AUTOFIX's metadata structure (§3.6).

3.1 Step 1: Constructing a Value-Flow Graph

An inter-procedural sparse value-flow graph (VFG) [11, 25, 26, 27, 28, 29, 34, 35] for a program is a multi-edged directed graph that captures the def-use chains of both top-level and address-taken variables. Top-level variables are the variables whose addresses are not taken. The def-use chains for top-level variables are readily available once they have been put in SSA form. Address-taken variables are accessed indirectly at loads and stores. Their def-use chains are built in several steps following [6, 26, 27]. First, the points-to information for the program is computed by using, e.g., Andersen's analysis. Second, a load $p = *q$ is annotated with a function $\mu(x)$ for each variable x that may be pointed to by q to represent

Algorithm 1: Liveness Analysis (for a leaked object o)

```
1 Let  $F$  be the set of candidate functions for fixing  $o$ 
2 foreach  $f \in F$  do
3   Let  $subCFG$  be the subgraph of  $f$ 's CFG that contains only
   the leaky paths for  $o$  in  $f$ 
4   foreach basic block  $b$  in  $subCFG$  do
5      $isLive(b) \leftarrow false$ 
6     if  $b$  contains an  $o$ -reachable variable then
7        $isLive(b) \leftarrow true$ 
8   while  $isLive$  has changed do
9     foreach basic block  $b$  in  $subCFG$  do
10       $isLive(b) \leftarrow \bigvee_{s \in succ(b)} isLive(s)$ 
11      //  $succ(b)$  is the set of successors of  $b$ 
```

a potential use of x at the load. Similarly, a store $*p = q$ is annotated with $x = \chi(x)$ for each variable x that may be pointed to by p to represent a potential def and use of x at the store. A callsite cs is also annotated with $\mu(x)$ and $x = \chi(x)$ for each variable x to capture inter-procedural reference and modification of x in a callee of cs . Third, all the address-taken variables are converted to SSA form, with each $\mu(x)$ being treated as a use of x and each $x = \chi(x)$ as both a def and use of x . Finally, the value-flows are constructed by connecting the def-uses for each converted SSA variable.

3.2 Step 2: Locating Functions to Fix

DEFINITION 1 (VALUE-FLOW REACHABILITY). *A variable v is o -reachable if there exists a value-flow path from the allocation site of o to the definition site of v on the VFG of the program. A callsite $p = call(\dots)$ is o -reachable if either variable p or x in any $x = \chi(x)$ function annotated at the callsite is o -reachable.*

Given a leaked object o from an allocation site reported by a static detector, **AUTOFIX** first determines the candidate functions where the leaks of o will be fixed. A function f is a candidate function to insert fixes for o if (1) f contains at least one o -reachable callsite and (2) there is no o -reachable variable in any caller of f .

In the case of recursion, if there is no data dependence on the leaked object o between any two function calls in the recursive cycle, o can be fixed in the recursive functions; otherwise fixes for o must be put outside the recursive functions to ensure safe fixing. In the case of global variables, since they are reachable for every function, leaked objects pointed to by global pointers can only be fixed in **main**. However, global variables are generally not considered as leaks in existing leak detectors [5, 15, 26, 27].

3.3 Step 3: Identifying Leaky Paths

To identify the leaky paths in a candidate function f , **AUTOFIX** performs a forward analysis on the VFG from an o -reachable callsite src to construct a value-flow slice S_{src} that includes all the nodes reachable from src but confined in f .

If no **free** sites are reachable from src , then all paths in f are leaky paths. Otherwise, for a **free** site snk corresponding to

S_{src} , let $vfp(src, snk)$ be the set of all value-flow paths from src to snk on the VFG, and $vfe(P)$ be the set of all value-flow edges in a single value-flow path $P \in vfp(src, snk)$. Thus, we can obtain the value-flow guards from src to snk :

$$VFGuard(src, snk) = \bigvee_{P \in vfp(src, snk)} \bigwedge_{(\hat{s}, \hat{d}) \in vfe(P)} CFGuard(\hat{s}, \hat{d})$$

$$CFGuard(\hat{s}, \hat{d}) = \bigvee_{Q \in cfp(\hat{s}, \hat{d})} \bigwedge_{e \in Q} boolCond(e)$$

where $CFGuard(\hat{s}, \hat{d})$ is a boolean formula that encodes the set of control-flow paths, denoted as $cfp(\hat{s}, \hat{d})$, from program point \hat{s} to \hat{d} on f 's CFG. Each branch edge e on a control flow path Q is uniquely assigned a boolean variable $boolCond(e)$. In the presence of loops, guards can grow unboundedly. To avoid unbounded conjunctions that describe all loop iterations, we follow [5, 26, 27] and bound loops to one iteration. Finally, the leak condition for src is obtained by computing guards from src to all its reachable **free** sites in f :

$$LeakCond = \bigvee_{snk \in S_{src}} VFGuard(src, snk)$$

Any path from src to the end of function f that satisfies $LeakCond$ is a leaky path for the leaked object o .

3.4 Step 4: Liveness Analysis

For a candidate function f , a subgraph is extracted from f 's CFG by excluding the control flow edges that are not on any leaky path. Then, a liveness analysis is applied to this subgraph to determine the basic blocks where o may be live.

As shown in Algorithm 1, a backward data-flow analysis is performed in f , starting from blocks containing o -reachable variables (lines 6 – 7) and iteratively marking the liveness of each block until a fix-point is reached (lines 8 – 11). For a leaked object o allocated inside a loop, if there is no data dependence on o between different loop iterations, then fixes for o can be inserted inside the loop; otherwise fixes for o must be put outside the loop to ensure safe fixing.

3.5 Step 5: Instrumentation

Based on the liveness information, Algorithm 2 performs instrumentation to insert fixes for every leaked object o in each of its candidate functions f . A call to **Fix()** is inserted either at the end of b where b is the last live basic block for o (lines 6 – 7) or at the beginning of n where n is a newly created basic block between a live block and a non-live block of o (lines 9 – 11). As shown earlier in Figure 3(d), these fixes serve to reclaim the dynamically allocated memory at the allocation site of o that would otherwise be leaked. In addition, a call to **Add()** is inserted after the allocation site of o (line 12) to track all its allocated objects in the shadow memory. A call to **Remove()** is inserted after each **free** site reachable from its corresponding allocation site (line 13), so that the freed objects are removed from the shadow memory.

Our instrumentation is safe even if the leaked object o is a false positive, for two reasons. First, the VFG of a program over-approximates its def-use chains. Thus, the last-use sites of o in its candidate functions for fixing o are conservatively

Algorithm 2: Instrumentation (for a leaked object o)

```

1 Let  $F$  be the set of candidate functions for fixing  $o$ 
2 Let  $allocID$  be the unique ID of  $o$ 's allocation site
3 foreach  $f \in F$  do
4   Let  $liveBBs$  be the set of basic blocks that are marked live
   for  $o$  in  $f$ 's CFG
5   foreach  $b \in liveBBs$  do
6     if  $\nexists s \in succ(b), isLive(s) = true$  then
7       Insert  $Fix(allocID)$  at the end of  $b$ 
8     else
9       foreach  $s \in succ(b)$  s.t.  $isLive(s) = false$  do
10        Insert a new block  $n$  between  $b$  and  $s$ 
11        Insert  $Fix(allocID)$  at the beginning of  $n$ 
12 Insert  $Add(allocID, p)$  immediately after the allocation site
     $p = malloc(\dots)$  for  $o$ 
13 Insert  $Remove(allocID, q)$  immediately after each free site
     $free(q)$  where  $q$  is  $o$ -reachable
  
```

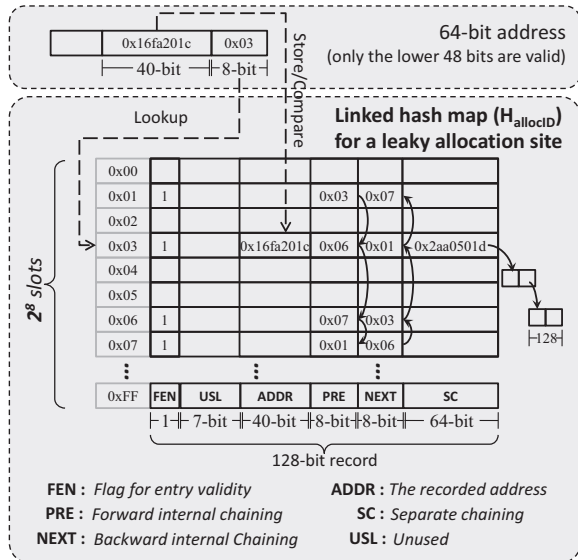


Figure 4: The metadata structure design.

found, ensuring safety by avoiding any use-after-free. Second, `Add()` and `Remove()` maintain valid memory addresses in the shadow memory, ensuring safety by avoiding any potential double free along any program path.

3.6 The Metadata Structure

The design philosophy behind our metadata structure is to enable a judicious tradeoff between time and space, which aims to support fast lookup, insertion and removal operations with reasonable space overhead. As shown in Figure 4, AUTOFIX maintains a closed hash map $H_{allocID}$ for every leaky allocation site (with its unique ID, $allocID$) to keep track of all the dynamic allocated memory addresses. The size of $H_{allocID}$ can be user-defined, with larger hash maps consuming more space and smaller ones potentially imposing higher slowdown due to hash collisions. To achieve a reasonable tradeoff (as evaluated in §4.3), the default hash map size is set to 2^8 , with 128 bits for each slot, resulting in a total of 4 KB consumed for the hash map.

Without loss of generality, our shadow mechanism is supported on a 64-bit x86_64 architectures with 48-bit virtual address space and word-aligned pointers. For a 64-bit memory address allocated at a leaky site, AUTOFIX uses its lower 8 bits as the index to the corresponding entry in the hash map, and maps its middle 40 bits to the field ADDR of the entry. A linked list is implemented to handle hash collisions in each hash slot, with the field SC recording the head of the list. Due to the sparsity of the hash map, it is expensive to retrieve all the valid entries to reclaim leaked memory by performing a full scan for the map. To speed up the search, we have used a doubly linked list (similar to Java’s `LinkedHashMap`) with the two 8-bit fields, PRE and NEXT, to record the previous and the next valid entry indexes. The one-bit field FEN indicates whether an entry is valid or not. It is set to 0 when the entry is removed from the hash map. The 7-bit field USL is preserved for future use.

Figure 5 gives the implementation of our shadow metadata functions i.e. `Add()`, `Remove()`, `Fix()`. The *lookup*, *insert*, *remove* and *clear* are standard operations that are similar to those in Java’s `LinkedHashMap` and are thus omitted. Function `Add()`, which is instrumented immediately after the leaky allocation site `allocID`, first finds a slot in $H_{allocID}$ to create an entry for the allocated memory object o (line 1), then maps bits 8 through 47 of o ’s address $addr_o$ to the 40-bit field ADDR using a simple shift operation (line 2), and finally inserts the entry into $H_{allocID}$ (line 3). Function `Remove()`, which is instrumented after a deallocation site of o , checks whether the deallocated address is recorded in $H_{allocID}$ (lines 4 and 5). If so, the corresponding entry is removed (line 6). Function `Fix()`, which is instrumented after the last-use site of o , first traverses $H_{allocID}$ using its internal linked list via *getNext* (lines 7), then frees all the recorded addresses (line 8) and clears the hash map (line 9).

4. EVALUATION

We have implemented AUTOFIX on top of LLVM (version 3.5.0). Eight C programs are used for evaluation as shown in Table 1, including five SPEC2000 benchmarks and three popular open-source applications. The SPEC2000 suite is widely used for evaluating static leak detectors [5, 15, 26, 27]. However, the SPEC2000 benchmarks that have less than two reported leaks (e.g. `parser` and `gap`) are excluded from our evaluation. The five selected SPEC2000 benchmarks are `ammp` (contains many leaks), `gcc` (large-sized and contains many leaks), `perlbnkm` (allocation-intensive), `twolf` (allocation-intensive) and `mesa` (deallocation-intensive). We also include three open-source applications: `a2ps-4.14` (a postscript filter) containing a relative large number of leaks, and two long-running server programs: `h2o-1.2` (an http server) and `redis-2.8` (a NoSQL database).

All our experiments are conducted on a platform consisting of a 3.0 GHZ Intel Core2 Duo processor with 16 GB memory, running RedHat Enterprise Linux 5 (kernel version 2.6.18). The source code of each program is compiled into bit-code files using `clang` and then merged together using LLVM Gold Plugin at link-time (LTO) to produce a whole-program bit-code file. The compiler option `mem2reg` is turned on to promote memory into registers. Andersen’s pointer analysis is used to build the VFG for the program [11]. We use

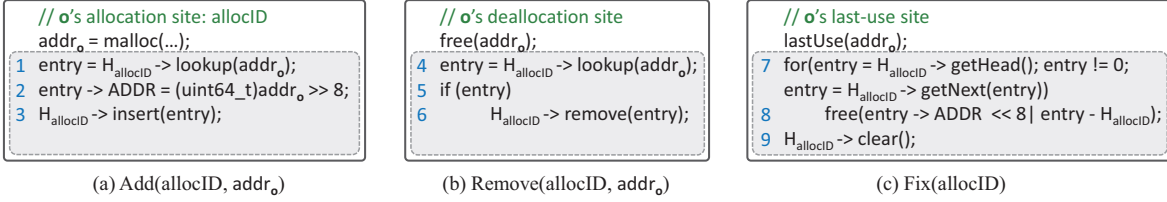


Figure 5: Implementation of shadow metadata operations.

Table 1: Benchmark characteristics

| Program | Size (KLOC) | #Alloc Sites | #Free Sites | #Reported Leaky Alloc Sites | #True Leaks |
|---------|-------------|--------------|-------------|-----------------------------|-------------|
| ammp | 13.4 | 37 | 30 | 20 | 20 |
| gcc | 230.4 | 161 | 19 | 45 | 40 |
| perlbmk | 87.1 | 148 | 2 | 12 | 8 |
| mesa | 61.3 | 82 | 76 | 7 | 3 |
| twolf | 20.5 | 185 | 1 | 5 | 5 |
| a2ps | 41.8 | 295 | 161 | 39 | 28 |
| h2o | 18.2 | 95 | 123 | 27 | 26 |
| redis | 61.8 | 47 | 62 | 24 | 20 |
| Total | 534.5 | 1050 | 474 | 179 | 150 |

Table 2: Compile-time statistics of AutoFix

| Program | #Fix() | #Add() | #Remove() | Analysis Time(s) |
|---------|--------|--------|-----------|------------------|
| ammp | 20 | 20 | 0 | 0.9 |
| gcc | 74 | 45 | 13 | 81.7 |
| perlbmk | 131 | 12 | 28 | 32.0 |
| mesa | 19 | 7 | 9 | 15.1 |
| twolf | 7 | 5 | 0 | 3.9 |
| a2ps | 51 | 39 | 48 | 17.0 |
| h2o | 61 | 27 | 2 | 9.5 |
| redis | 30 | 24 | 7 | 56.0 |
| Total | 393 | 179 | 107 | 216.1 |

the leak warnings (leaky allocation sites) reported by the state-of-the-art leak detector, SABER [26, 27], as input to AUTOFIX.

We evaluate AUTOFIX based on three criteria: (1) efficiency (number of fixes generated and the analysis time taken to do so), (2) effectiveness (ability to fix memory leaks and reduce memory usage at runtime), and (3) performance degradation (instrumentation overhead at runtime).

4.1 Efficiency of AutoFix

The compile-time results of AUTOFIX are summarized in Table 2. Given a total of 179 leaky allocations reported in the eight programs, AUTOFIX fixes them all by inserting 393 calls to Fix() (Column 2), 179 calls to Add() (Column 3) and 107 calls to Remove() (Column 4). On average, a leaky allocation results in only 2.2 fixes. This shows that AUTOFIX is able to precisely place fixes along the identified leaky paths with lightweight instrumentation. As shown in Table 2 (Column 5), it takes 216.1 seconds to analyze the 534.5 KLOC for the eight C programs altogether. In particular, AUTOFIX spends 81.7 seconds on gcc, the largest program (230.4 KLOC) studied. The analysis times for the other seven programs are all within one minute.

4.2 Effectiveness of AutoFix

To evaluate the effectiveness of AUTOFIX in fixing leaks at runtime, we compare the memory usage of each program before and after automated fixing using VALGRIND [18]. For

the five SPEC2000 benchmarks, their *reference* inputs are used. For the three open-source applications, their own regression test suites are used. For a total of 67 real leaks triggered by the inputs (Column 2 in Table 3), AUTOFIX is able to reclaim all the leaked memory at runtime, which is verified by VALGRIND [18].

In our experiments, we observed that a substantial number of leaks are inter-procedural, involving path correlation. These leaks are ignored and cannot be fixed by the pure static approach LEAKFIX [9] due to the over-approximative nature of static analysis. In contrast, AUTOFIX combines static analysis with runtime checking to enable precise fixing for all leaks including those involving path correlations.

Figure 6 shows a leak pattern in gcc and its fixed code. The pointer *p* at the callsite `use` may point to either a heap object when the `if` branch is taken or a stack object otherwise. A leak happens in the former case, while the code is leak-free in the latter case. AUTOFIX tracks the memory allocation by instrumenting an `Add` after `malloc` and reclaims only truly leaked memory by performing runtime checks in `Fix`.

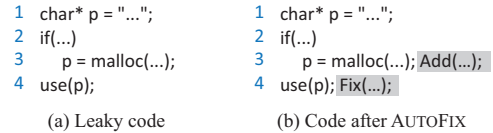


Figure 6: Fixing path correlated leaks in AutoFix.

To further evaluate the effectiveness of AUTOFIX in fixing leaks for long-running programs, we reproduce a real leak which causes memory exhaustion in redis with its corresponding regression tests [17]. As shown in Figure 7, a loop (line 281) is used to query the IPs of slave servers. If a slave sever is dead, reconnection attempts are repeatedly made by calling `getaddrinfo` (line 291), which allocates memory chunks that are never freed, resulting in leaks inside the `for` loop. As shown in Figure 8, the memory consumption of the leaked version of redis increases around 31.7 KB per second, and over 300 MB are leaked after three hours. AUTOFIX can fix this leak effectively, enabling redis’s memory consumption to remain below 35 MB in the fixed version.

4.3 Runtime Overhead

To measure runtime overhead, each program is executed five times before and after automated fixing respectively, and the average overhead is reported in Table 3. The experimental results show that AUTOFIX only introduces negligible overhead for all the eight programs, 1.06% on average, with the maximum 1.82% observed in a2ps. This confirms that our instrumentation is lightweight, achieved by identifying the required deallocation fixes on the value-flow slices of leaky allocations and tracking leaked objects with simple shadow

```

256: static int _redisContextConnectTcp(...) { ...
281:     for (p = servinfo; p != NULL; p = p->ai_next) { ...
291:         rv = getaddrinfo(source_addr, NULL, &hints, &bservinfo);
++         Add (bservinfo, allocID); // instrumented by AutoFix
297:         for (b = bservinfo; b != NULL; b = b->ai_next) { ...
302:             }
++         Fix (allocID); // instrumented by AutoFix
329:     }
342: }

```

Figure 7: The leaky code (in net.c of redis-2.8) fixed by AutoFix.




Figure 8: Memory footprint of redis-2.8 before and after fixing the leak.

operations at runtime.

To evaluate the impact of the metadata structure (as described in §3.6) on runtime instrumentation overhead, we choose four different sizes for the hash map used in order to demonstrate the time and space tradeoffs made: 1 (with the hash map degenerating into one linked list), 2^6 , 2^8 and 2^{16} . The results are shown in Figure 9.

For the five benchmarks, `ammp`, `gcc`, `mesa`, `twolf` and `redis`, the four configurations yield similar overheads. However, for the other three benchmarks, `perlbnk`, `a2ps` and `h2o`, much higher overheads are incurred when their underlying hash maps have degenerated into a single linked list. In this degenerate case, the lookup operations become too expensive, especially when a large number of memory objects are present. When the other three hash map sizes are used, lookup operations can be performed more efficiently. The hash map with 2^6 slots is not very space-consuming, costing 1 KB for each leaky allocation. However, due to its high collision rates, this hash map still results in high overheads for `perlbnk`, `a2ps` and `h2o`. As shown in Figure 9, the hash maps with 2^8 slots (4 KB per leaky allocation) and 2^{16} slots (1 MB per leaky allocation) suffer from similar overheads. This indicates that 2^8 slots are already sufficient to guarantee low collision rates, and more slots cannot provide any noticeable performance benefit. For more complicated applications beyond our evaluation, it is still possible that 2^8 slots are not enough to ensure low hash collision rates. In this situation, AUTOFIX allows users to allocate more slots for the shadow hash map to achieve better performance.

5. RELATED WORK

Leak Detection: Memory leak detection has been extensively studied using static [15, 26, 27, 32] or dynamic [12, 18, 24] analysis. Static detectors examine the source code at compile-time without executing the program. SATURN [32]

Table 3: Run-time statistics of AutoFix

| Program | #Triggered Leaks | Overhead (%) |
|---------|------------------|--------------|
| ammp | 1 | 1.36 |
| gcc | 13 | 0.75 |
| perlbnk | 12 | 0.88 |
| mesa | 3 | 0.76 |
| twolf | 2 | 0.89 |
| a2ps | 12 | 1.82 |
| h2o | 15 | 1.58 |
| redis | 9 | 0.66 |

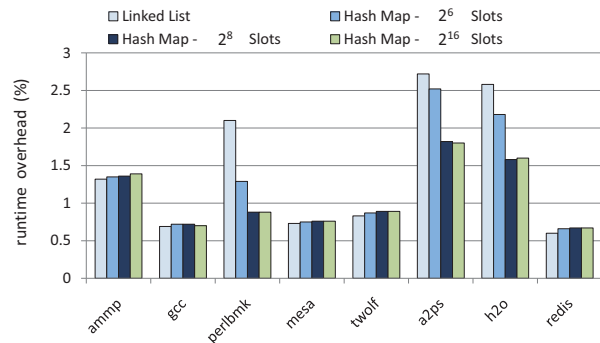


Figure 9: Comparing runtime overheads for different metadata structures used in AutoFix.

detects memory leaks by solving a Boolean satisfiability problem. SPARROW [15] is based on abstract interpretation, and uses function summaries. FASTCHECK [5] and SABER [26, 27] find memory leaks on the value-flow graph of the program. Dynamic detectors, which find leaks by executing the program, track the memory allocation and deallocation via either binary instrumentation as in VALGRIND [18] or source code instrumentation as in ADDRESSSANITIZER [24].

Leak Tolerance: Another line of research focuses on tolerating leaks at runtime [3, 20, 30]. The basic idea is to delay out-of-memory crashes at runtime by offloading stale objects (regarded as likely leaked) to disks and reclaiming their virtual memory. Upon accessing a mistakenly swapped-out object, the object will be swapped back into the memory, thereby guaranteeing safety. AUTOFIX, which aims at fixing leaks, is orthogonal to tolerating leaks.

Garbage Collection: Garbage collection can eliminate most memory leaks. However, in type-unsafe languages such as C and C++, it is theoretically impossible to implement sound garbage collectors to automatically manage memory. A few unsound solutions for C and C++ [2, 13, 22] have been shown empirically to be effective with low space and time overheads, in which memory allocations (e.g. `malloc` sites) are replaced by special allocators, and memory deallocations (e.g. `free` sites) are removed from the program, at the expense of the prompt low-cost reclamation provided by explicit memory management. Compared to AUTOFIX’s static fixing on value-flow slices, garbage collection uses runtime object reachability to over-approximate object liveness. In addition, garbage collectors for C and C++ typically need to monitor all static data areas, stacks, registers and heap. In contrast, AUTOFIX only monitors potential leaky allocations reported by leak detectors, which makes AUTOFIX much more lightweight than garbage collectors.

Leak Fixing: Memory leaks can be fixed manually or automatically. LEAKPOINT [7] is a dynamic taint analysis that identifies last-use sites of leaked objects by tracking pointers and presents programmers the identified sites as candidate locations for leak fixing. LEAKCHASER [33] relies on user annotations to improve the relevance of bug reports, thereby assisting programmers to diagnose and fix memory leaks. Object ownership profiling has also been applied to assisting manual leak detection and fixing [23]. LEAKFIX [9] is a pure static approach to automatically fixing leaks in C programs. Because it cannot handle false positives produced by other state-of-the-art leak detectors, LEAKFIX relies on its own dedicated leak detector and can fix only some but not all reported leaks. In contrast, our approach combines static and dynamic analysis, and is able to automatically fix all the true leaks reported by a detector with small overhead.

6. CONCLUSION

This paper presents AUTOFIX, a fully automated approach to memory leak fixing for C programs by combining static and dynamic analysis. Given a leaky allocation reported by a leak detector, AUTOFIX performs a graph reachability analysis to identify the leaky paths on the value-flow graph of the program and then a liveness analysis to locate the program points for instrumenting the required fixes on the identified leaky paths at compile-time. Our evaluation shows that AUTOFIX is capable of fixing reported memory leaks with small instrumentation overhead.

7. ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers for their valuable comments. This work is supported by ARC grants, DP130101970 and DP150102109.

8. REFERENCES

- [1] A. Arcuri and X. Yao. A novel co-evolutionary approach to automatic software bug fixing. In *IEEE World Congress on Computational Intelligence*, pages 162–168, 2008.
- [2] H.-J. Boehm. Bounding space usage of conservative garbage collectors. In *POPL '02*, pages 93–100, 2002.
- [3] M. D. Bond and K. S. McKinley. Tolerating memory leaks. In *OOPSLA '08*, pages 109–126, 2008.
- [4] D. Bruening and Q. Zhao. Practical memory checking with dr. memory. In *CGO '11*, pages 213–223, 2011.
- [5] S. Cherem, L. Princehouse, and R. Rugina. Practical memory leak detection using guarded value-flow analysis. In *PLDI '07*, pages 480–491, 2007.
- [6] F. Chow, S. Chan, S. Liu, R. Lo, and M. Streich. Effective representation of aliases and indirect memory operations in SSA form. In *CC '96*, pages 253–267, 1996.
- [7] J. Clause and A. Orso. Leakpoint: pinpointing the causes of memory leaks. In *ICSE '10*, pages 515–524, 2010.
- [8] V. Dallmeier, A. Zeller, and B. Meyer. Generating fixes from object behavior anomalies. In *ASE '09*, pages 550–554, 2009.
- [9] Q. Gao, Y. Xiong, Y. Mi, L. Zhang, W. Yang, Z. Zhou, B. Xie, and H. Mei. Safe memory-leak fixing for C programs. In *ICSE'15*, 2015.
- [10] C. L. Goues, T. Nguyen, S. Forrest, and W. Weimer. Genprog: A generic method for automatic software repair. *IEEE Trans. on Software Engineering*, 38(1):54–72, 2012.
- [11] B. Hardekopf and C. Lin. Flow-sensitive pointer analysis for millions of lines of code. In *CGO '11*, pages 289–298, 2011.
- [12] M. Hauswirth and T. M. Chilimbi. Low-overhead memory leak detection using adaptive statistical profiling. In *ASP-LOS '04*, pages 156–164, 2004.
- [13] M. Hirzel, A. Diwan, and J. Henkel. On the usefulness of type and liveness accuracy for garbage collection and leak detection. *ACM Transactions on Programming Languages and Systems*, 24(6):593–624, 2002.
- [14] G. Jin, W. Zhang, D. Deng, B. Liblit, and S. Lu. Automated concurrency-bug fixing. In *OSDI '12*, pages 221–236, 2012.
- [15] Y. Jung and K. Yi. Practical memory leak detector based on parameterized procedural summaries. In *ISMM'08*, pages 131–140, 2008.
- [16] C. Le Goues, S. Forrest, and W. Weimer. Current challenges in automatic software repair. *Software Quality Journal*, 21(3):421–443, 2013.
- [17] D. Mezzatto. Sentinel 2.8 branch memory leak in redis, <https://github.com/antirez/redis/issues/2012>, 2014.
- [18] N. Nethercote and J. Seward. Valgrind: a framework for heavyweight dynamic binary instrumentation. In *PLDI '07*, pages 89–100, 2007.
- [19] H. D. T. Nguyen, D. Qi, A. Roychoudhury, and S. Chandra. Semfix: Program repair via semantic analysis. In *ICSE '13*, pages 772–781, 2013.
- [20] G. Novark, E. D. Berger, and B. G. Zorn. Plug: automatically tolerating memory leaks in C and C++ applications. *Technical Report UM-CS-2008-009*, University of Massachusetts, 2008.
- [21] J. H. Perkins, S. Kim, S. Larsen, S. Amarasinghe, J. Bachrach, M. Carbin, C. Pacheco, F. Sherwood, S. Sidiroglou, G. Sullivan, et al. Automatically patching errors in deployed software. In *SOSP '09*, pages 87–102, 2009.
- [22] J. Raffkind, A. Wick, J. Regehr, and M. Flatt. Precise garbage collection for c. In *ISMM '09*, pages 39–48, 2009.
- [23] D. Rayside and L. Mendel. Object ownership profiling: a technique for finding and fixing memory leaks. In *ASE'07*, pages 194–203, 2007.
- [24] K. Serebryany, D. Bruening, A. Potapenko, and D. Vyukov. AddressSanitizer: A fast address sanity checker. In *USENIX Annual Technical Conference*, pages 309–318, 2012.
- [25] Y. Sui, P. Di, and J. Xue. Sparse flow-sensitive pointer analysis for multithreaded c programs. In *CGO '16*.
- [26] Y. Sui, D. Ye, and J. Xue. Static memory leak detection using full-sparse value-flow analysis. In *ISSTA '12*, pages 254–264, 2012.
- [27] Y. Sui, D. Ye, and J. Xue. Detecting memory leaks statically with full-sparse value-flow analysis. *IEEE Trans. Software Eng.*, 40(2):107–122, 2014.
- [28] Y. Sui, S. Ye, J. Xue, and P. Yew. SPAS: Scalable path-sensitive pointer analysis on full-sparse SSA. In *APLAS '11*, pages 155–171.
- [29] Y. Sui, S. Ye, J. Xue, and J. Zhang. Making context-sensitive inclusion-based pointer analysis practical for compilers using parameterised summarisation. *SPE '13*, 2013.
- [30] Y. Tang, Q. Gao, and F. Qin. Leaksurvivor: towards safely tolerating memory leaks for garbage-collected languages. In *USENIX Annual Technical Conference*, pages 307–320, 2008.
- [31] Y. Wei, Y. Pei, C. A. Furia, L. S. Silva, S. Buchholz, B. Meyer, and A. Zeller. Automated fixing of programs with contracts. In *ISSTA '10*, pages 61–72, 2010.
- [32] Y. Xie and A. Aiken. Context- and path-sensitive memory leak detection. In *FSE'05*, pages 116–125, 2005.
- [33] G. Xu, M. D. Bond, F. Qin, and A. Rountev. Leakchaser: helping programmers narrow down causes of memory leaks. In *PLDI '11*, pages 270–282, 2011.
- [34] D. Ye, Y. Sui, and J. Xue. Accelerating dynamic detection of uses of undefined variables with static value-flow analysis. In *CGO '14*, pages 154–164.
- [35] S. Ye, Y. Sui, and J. Xue. Region-based selective flow-sensitive pointer analysis. In *SAS '14*, pages 319–336. Springer, 2014.