

Automated Code Repair

Will Klieber



Copyright 2016 Carnegie Mellon University

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense.

NO WARRANTY. THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN “AS-IS” BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.


[Distribution Statement A] This material has been approved for public release and unlimited distribution. Please see Copyright notice for non-US Government use and distribution.

This material may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

Carnegie Mellon® and CERT® are registered marks of Carnegie Mellon University.

DM-0004068

Automated Code Repair – Motivation



Software vulnerabilities constitute a major threat

- A majority arise from common coding errors
- Shown by experience from source code analysis labs at CERT and DoD

Static analysis tools help, but:

- Typically are used late in the development process
- Produce an enormous number of warnings
- The volume of true positives often overwhelms the ability of the development team to fix the code

Huge amount of code in use by DoD

- Billions of lines of C code
- Unknown number of security vulnerabilities



Integer Overflow

This past year (FY16), we developed techniques for automated repair of **integer overflows** that lead to **memory corruption**

Integers in C are represented by a fixed number of bits N (e.g., 32 or 64).

- Overflow occurs when the result cannot fit in N bits
- Modular arithmetic: Only the least significant N bits are kept

How does integer overflow lead to memory corruption?

1. Memory allocation: `malloc(·)`.
2. Bounds checks for an array

Example: Android Stagefright bugs (July 2015)

Benefits to DoD



Eliminate security vulnerabilities at a **much lower cost** than manual repair

Integer overflows are a **very common** type of bug

- In CERT SCALe audits, about 80% of findings were related to fixed-width integers

Our technique:

- **Will not break working code**, provided *inferred specification* is correct (Next slide)
- Typically total slowdown < 5% (Based on theoretical model)
- False positives: Flagged operations that cannot actually overflow
 - Then our 'repair' just adds a little unnecessary overhead



Premises for Automated Repair

1. Many security bugs follow common patterns
 - E.g., “`p = malloc(n * sizeof(T))`” where n is attacker-controlled
 - Integer overflow \Rightarrow too little memory gets allocated
2. By recognizing such a pattern, it is possible to make a reasonable guess of the developer's intention (the *inferred specification*)
 - E.g., “Try to allocate enough memory to hold n objects of type T ”
3. It is possible to repair the code to satisfy this inferred specification
 - Check if overflow occurs; if so, simulate `malloc` failing with `ENOMEM`

Experimental Results



	Overflows (as reported by Kint)	Overflows that are <i>sensitive</i>	Overflows fully repaired	Semi-repair	Unrepaired
OpenSSL (1.0.2g)	969	233	180	28	25
Jasper	481	101	53	32	16

An overflow is ***sensitive*** if it involves variables that are associated with array indices or bounds

Note: Some of the above “repairs” are actually false positives (i.e., operation never overflows). Others are known vulnerabilities with CVEs and patches.

Repair Strategy



Inferred specification: **inequality comparisons** involving array indices or bounds should behave as if normal arithmetic (not modular arithmetic) were used

- Includes `malloc`
- Excludes crypto and hashing functions

Repair: General case is intractable (with bounded memory)

Special case that we handle: non-negative integers with only addition or multiplication (no subtraction or division)


- The value is monotonically non-decreasing (except for multiplication by zero)
- Normal arithmetic can be emulated using **saturation arithmetic**:
 - Replace an overflowed value with the greatest representable value (`SIZE_MAX`)
- If the declared types of variables are smaller than `size_t`, they are promoted up

Arithmetic for Checking Bounds of an Array

Example: copy n bytes from src to $dest$, starting at index $start$ of $dest$.

Repair: `UADD(start, n) /* defined on next slide */`

```
if (start + n <= dest_size) {  
    memcpy(&dest[start], src, n);  
} else {  
    return -EINVAL;  
}
```





wrappers.h

```

1. inline static size_t UADD(size_t lop, size_t rop) {
2.     size_t result;
3.     bool flag = __builtin_add_overflow(lop, rop, &result);
4.     if (flag) {result = SIZE_MAX;}
5.     return result;
6. }

```

Repair: UADD(start, n)

```

if (start + n <= dest_size) {
    memcpy(&dest[start], src, n);
} else {
    return -EINVAL;
}

```

- What if dest_size is SIZE_MAX?
- What if both sides of inequality overflow?
- What if overflow reaches a non-comparison sink?

Semi-Repair

Example adapted from CVE-2015-8370:

```

1.  unsigned cur_len = 0;
2.  while (1) {
3.      key = grub_getkey();
4.      if (key == '\b') {
5.          if (cur_len == 0) {
6.              /* FIXME: Insert error-
7.                 handling code here. */
8.          }
9.          cur_len--;
10.         grub_printf("\b");
11.         continue;
12.     }
13.     if (cur_len + 2 < buf_size) {
14.         buf[cur_len++] = key;
15.         grub_printf("%c", key);
16.     }

```

semi-repair

If a potentially overflowed value is used to index into an array, do a semi-repair.

Tool inserts check for overflow.
User writes error-handling code.

Future Directions



In FY17, we have two Automated Repair projects:

1. **Inference of memory bounds**
 - Buffer overflow (WRITES) and leakage of sensitive information (READs)
2. **Incorrect usage of crypto/security APIs**
 - E.g., incorrect validation of certificate chain using OpenSSL API, leading to MITM

A difficulty we encountered was the **Source ↔ IR mapping problem**:

- Code is most readily analyzed and repaired on an intermediate representation (IR)
- Transformations on the IR aren't unambiguously mappable to the source
- Macros and `#ifdefs` are a further difficulty
 - Prof. Christian Kästner (CMU SCS) has done work on `#ifdefs` as part of this project
- We are further investigating these issues this year (FY17)

Conclusion



Automated code repair (ACR) reduces a system's attack surface and improves its ability to withstand cyber-attacks.

ACR is suitable for problems where many bugs follow a common pattern and have a corresponding pattern for repair.

In FY16, we focused on integer overflows involving memory bounds/indices.

We are continuing work on ACR in FY17.

Contact Information

Will Klieber

Software Security Researcher

Telephone: +1 412.268.9207

Email: weklieber@cert.org

