

Fuzzing and Symbolic Execution

SUSTech CS 315 Computer Security 2023

Outline

- Program Analysis
- Exploiting **Real World** Programs
- Fuzz Testing
- Symbolic Execution
- Mixed Fuzz and Symbolic Execution (Concolic Execution)
- Summary

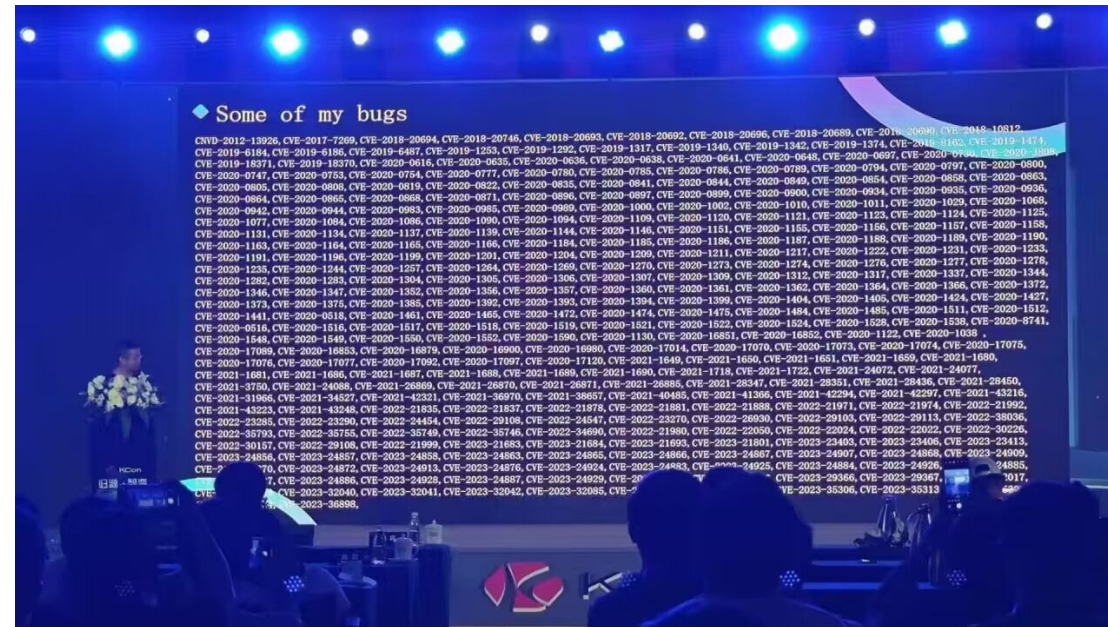
Program Analysis

- Bug hunters reported 1000+ bugs

- They want to find bug efficiently in countless real-world applications

- Large companies like google

- They need to ensure the quality of the **large amounts of software** they release
- They also need to ensure that **the open-source components** they use are safe and stable



KCON 2023

Exploiting Real World Programs

- Real program can be complicated
 - linux kernel has 66000+ files and 2700000000+ lines of code
 - what about your Java class project? 1000~5000 lines?
- How to find bugs and vulnerabilities in large and complicated programs?

当你帮酒吧老板写了一个程序:

- 一个测试工程师走进一家酒吧, 要了一杯啤酒;
- 一个测试工程师走进一家酒吧, 要了一杯咖啡;
- 一个测试工程师走进一家酒吧, 要了0.7杯啤酒;
- 一个测试工程师走进一家酒吧, 要了-1杯啤酒;
- 一个测试工程师走进一家酒吧, 要了一杯蜥蜴;
- 一个测试工程师走进一家酒吧, 要了一份asdfQwer@24dg!&*(@;
- 一个测试工程师走进一家酒吧, 什么也没要;
- 一个测试工程师走进一家酒吧, 又走出去又从窗户进来又从后门出去从下水道钻进来;
- 一个测试工程师走进一家酒吧, 要了一杯烫烫烫的银斤拷;
- 一个测试工程师走进一家酒吧, 要了NaN杯Null;
- 一个测试工程师把酒吧拆了;
- 一个测试工程师化装成老板走进一家酒吧, 要了500杯啤酒并且不付钱;
- 一万个测试工程师在酒吧门外呼啸而过;
- 一个测试工程师走进一家酒吧, "< script >alert("要了一杯酒");< /script >"
- 一个测试工程师走进一家酒吧, 要了一杯啤酒';DROP TABLE 酒吧;

测试工程师们满意地离开了酒吧。

然后一名顾客点了一份炒饭, 酒吧炸了。

Manually constructing test cases is not enough

Fuzz Testing


- Recall vulnerable program in last lab

- Easily crashed by some long input

```
student@ubuntu:~/Desktop/lab3$ ./bof
Welcome back to 2023 CS315, let's have some fun!
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
Segmentation fault (core dumped)
student@ubuntu:~/Desktop/lab3$
```

- How can we find bug in thousands of real world programs?

- Hire some monkeys
- "randomly" generate some inputs



Fuzz Testing

- Recall vulnerable program in last lab

- Easily crashed by some long input


```
student@ubuntu:~/Desktop/lab3$ ./bof
Welcome back to 2023 CS315, let's have some fun!
aaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaaa
Segmentation fault (core dumped)
student@ubuntu:~/Desktop/lab3$
```

- How can we find bug in thousands of real-world programs?

- Hire some monkeys
- “randomly” generate some inputs

- Fuzzer: hire some clever monkeys

- automatically generate and mutate inputs



Fuzz Testing

- AFL & AFL++ (American fuzzy lop)
 - A modern fuzzing tool
 - employs **genetic algorithms** to efficiently increase code coverage
 - <https://www.usenix.org/system/files/woot20-paper-fioraldi.pdf>
 - Best used with address sanitization (ASAN)
 - Flag *all* invalid memory accesses
 - Shadow memory tracking which memory areas are valid
 - Finds out of bounds access and use after free bugs
 - AFL+ASAN combination is gold standard of fuzzing
 - we will try AFL+ASAN in our lab

```
american fuzzy lop 1.86b (test)
-----
process timing | overall results
-----|-----
run time : 0 days, 0 hrs, 0 min, 2 sec | cycles done : 0
last new path : none seen yet | total paths : 1
last uniq crash : 0 days, 0 hrs, 0 min, 2 sec | uniq crashes : 1
last uniq hang : none seen yet | uniq hangs : 0
-----|-----
cycle progress | map coverage
-----|-----
now processing : 0 (0.00%) | map density : 2 (0.00%)
paths timed out : 0 (0.00%) | count coverage : 1.00 bits/tuple
-----|-----
stage progress | findings in depth
-----|-----
now trying : havoc | favored paths : 1 (100.00%)
stage execs : 1464/5000 (29.28%) | new edges on : 1 (100.00%)
total execs : 1697 | total crashes : 39 (1 unique)
exec speed : 626.5/sec | total hangs : 0 (0 unique)
-----|-----
fuzzing strategy yields | path geometry
-----|-----
bit flips : 0/16, 1/15, 0/13 | levels : 1
byte flips : 0/2, 0/1, 0/0 | pending : 1
arithmetics : 0/112, 0/25, 0/0 | pend fav : 1
known ints : 0/10, 0/28, 0/0 | own finds : 0
dictionary : 0/0, 0/0, 0/0 | imported : n/a
havoc : 0/0, 0/0 | variable : 0
trim : n/a, 0.00% |
-----|-----
[cpu: 92%]
```

American fuzzy lop's afl-fuzz running on a test program

Fuzz Testing

- AFL & AFL++ (American fuzzy lop)
 - A modern fuzzing tool
 - employs **genetic algorithms** to efficiently increase code coverage
 - <https://www.usenix.org/system/files/woot20-paper-fioraldi.pdf>
- VUzzer, SYMcc, etc...

CCS Hawkeye: Towards a Desired Directed Grey-box Fuzzer

CCS Revery: from Proof-of-Concept to Exploitable (One Step towards Automatic Exploit Generation)

S&P T-Fuzz: fuzzing by program transformation

(and many many other interesting recent works)

```
american fuzzy lop 1.86b (test)
-----
process timing | overall results
-----|-----
run time : 0 days, 0 hrs, 0 min, 2 sec | cycles done : 0
last new path : none seen yet | total paths : 1
last uniq crash : 0 days, 0 hrs, 0 min, 2 sec | uniq crashes : 1
last uniq hang : none seen yet | uniq hangs : 0
-----|-----
cycle progress | map coverage
-----|-----
now processing : 0 (0.00%) | map density : 2 (0.00%)
paths timed out : 0 (0.00%) | count coverage : 1.00 bits/tuple
-----|-----
stage progress | findings in depth
-----|-----
now trying : havoc | favored paths : 1 (100.00%)
stage execs : 1464/5000 (29.28%) | new edges on : 1 (100.00%)
total execs : 1697 | total crashes : 39 (1 unique)
exec speed : 626.5/sec | total hangs : 0 (0 unique)
-----|-----
fuzzing strategy yields | path geometry
-----|-----
bit flips : 0/16, 1/15, 0/13 | levels : 1
byte flips : 0/2, 0/1, 0/0 | pending : 1
arithmetics : 0/112, 0/25, 0/0 | pend fav : 1
known ints : 0/10, 0/28, 0/0 | own finds : 0
dictionary : 0/0, 0/0, 0/0 | imported : n/a
havoc : 0/0, 0/0 | variable : 0
trim : n/a, 0.00% |
-----|-----
[cpu: 92%]
```


American fuzzy lop's afl-fuzz running on a test program

Fuzz Testing

- Problem:

- Still hard to reach some branches

hard to pass this
check by fuzz



```
1
2 #define KEY_SIZE 95
3 int sc_decompress(int infd, int outfd) {
4     unsigned char keys[KEY_SIZE];
5     unsigned char data[KEY_SIZE];
6     char *header = read_header(infd)
7     // C1: check for hardcoded values
8     if (strcmp(header, "SECO") != 0)
9         return ERROR;
10    read(infd, keys, KEY_SIZE);
11    memset(data, 0, sizeof(data));
12    // C2: range check and duplicate check for keys
13    for (int i = 0; i < sizeof(data); ++i) {
14        if (keys[i] < 32 || keys[i] > 126)
15            return ERROR;
16        if (data[keys[i] - 32]++ > 0)
17            return ERROR;
18    }
19    unsigned int in_len = read_len(infd);
20    char *in = (char *) malloc(in_len);
21    read(infd, in, in_len);
22    unsigned int crc = read_checksum(infd);
23    // C3: check the crc of the input
24    if (crc != compute_crc(in, in_len)) {
25        free(in);
26        return ERROR;
27    }
28    char *out;
29    unsigned int out_len;
30    // Bug: function with stack buffer overflow
31    decompress(in, in_len, keys, &out, &out_len);
32    write(outfd, out, out_len);
33    return SUCCESS;
34 }
```

Listing 1: An example containing various sanity checks


S&P'18 T-Fuzz: fuzzing by program transformation

Fuzz Testing

- **Problem:**


- Still hard to reach some branches

hard to pass this check by fuzz



if fuzz is "clever" enough and applied some genetic algorithms, it may pass the magic number check after a while

hard to pass crc check even after mutations



but the fuzz is nearly impossible to pass CRC check
CRC is a kind of hash algorithm, it often used to check the integrity of data (e.g. a zip archive can store CRC to check weather uncompressed file is correct)

```
1
2 #define KEY_SIZE 95
3 int sc_decompress(int infd, int outfd) {
4     unsigned char keys[KEY_SIZE];
5     unsigned char data[KEY_SIZE];
6     char *header = read_header(infd)
7     // C1: check for hardcoded values
8     if (strcmp(header, "SECO") != 0)
9         return ERROR;
10    read(infd, keys, KEY_SIZE);
11    memset(data, 0, sizeof(data));
12    // C2: range check and duplicate check for keys
13    for (int i = 0; i < sizeof(data); ++i) {
14        if (keys[i] < 32 || keys[i] > 126)
15            return ERROR;
16        if (data[keys[i] - 32]++ > 0)
17            return ERROR;
18    }
19    unsigned int in_len = read_len(infd);
20    char *in = (char *) malloc(in_len);
21    read(infd, in, in_len);
22    unsigned int crc = read_checksum(infd);
23    // C3: check the crc of the input
24    if (crc != compute_crc(in, in_len)) {
25        free(in);
26        return ERROR;
27    }
28    char *out;
29    unsigned int out_len;
30    // Bug: function with stack buffer overflow
31    decompress(in, in_len, keys, &out, &out_len);
32    write(outfd, out, out_len);
33    return SUCCESS;
34 }
```

Listing 1: An example containing various sanity checks


Fuzz Testing

- Problem:


- Still hard to reach some branches

- Another technique: **symbolic execution**

hard to pass this
check by fuzz



hard to pass crc check
even after mutations



```
1
2 #define KEY_SIZE 95
3 int sc_decompress(int infd, int outfd) {
4     unsigned char keys[KEY_SIZE];
5     unsigned char data[KEY_SIZE];
6     char *header = read_header(infd)
7     // C1: check for hardcoded values
8     if (strcmp(header, "SECO") != 0)
9         return ERROR;
10    read(infd, keys, KEY_SIZE);
11    memset(data, 0, sizeof(data));
12    // C2: range check and duplicate check for keys
13    for (int i = 0; i < sizeof(data); ++i) {
14        if (keys[i] < 32 || keys[i] > 126)
15            return ERROR;
16        if (data[keys[i] - 32]++ > 0)
17            return ERROR;
18    }
19    unsigned int in_len = read_len(infd);
20    char *in = (char *) malloc(in_len);
21    read(infd, in, in_len);
22    unsigned int crc = read_checksum(infd);
23    // C3: check the crc of the input
24    if (crc != compute_crc(in, in_len)) {
25        free(in);
26        return ERROR;
27    }
28    char *out;
29    unsigned int out_len;
30    // Bug: function with stack buffer overflow
31    decompress(in, in_len, keys, &out, &out_len);
32    write(outfd, out, out_len);
33    return SUCCESS;
34 }
```

Listing 1: An example containing various sanity checks

S&P'18 T-Fuzz: fuzzing by program transformation

Symbolic Execution

- ① Use symbol to represent variables
- ② Simulate program execution
- ③ Extract and solve constraints in execution path
- generate constraints in every path
- constraint solver: z3, cvc5

```
void test_me(int x) {  
    if (x == 94389) {  
        ERROR;  
    }  
}
```

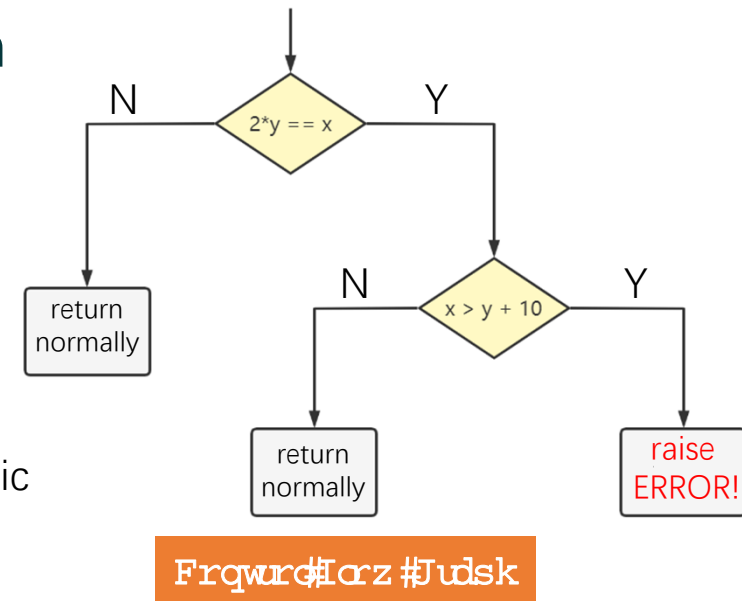
Probability of **ERROR**:
 $1/2^{32} \approx 0.0000000023\%$

some path is hard to reach by randomly fuzzing
If we can extract the constraint ($x==94389$), we can
easily find the input to reach those path

Symbolic Execution

- ① Use symbol to represent variables
- ② Simulate program execution
- ③ Extract and solve constraints in execution path
- generate constraints in every path
- constraint solver: z3, cvc5

```
int main(){
  x = read_int();
  y = read_int();
  z = 2 * x;
  if (z == x){
    if (x > y+10)
      ERROR;
  }
}
```

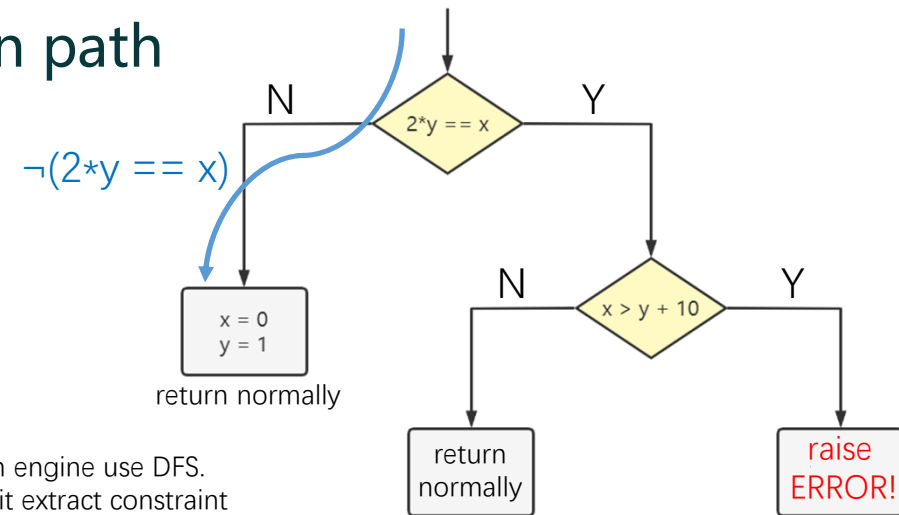


let's see how symbolic execution works ->

Symbolic Execution

- ① Use symbol to represent variables
- ② Simulate program execution
- ③ Extract and solve constraints in execution path
- generate constraints in every path
- constraint solver: z3, cvc5

```
int main(){
  x = read_int();
  y = read_int();
  z = 2 * x;
  if (z == x){
    if (x > y+10)
      ERROR;
  }
}
```

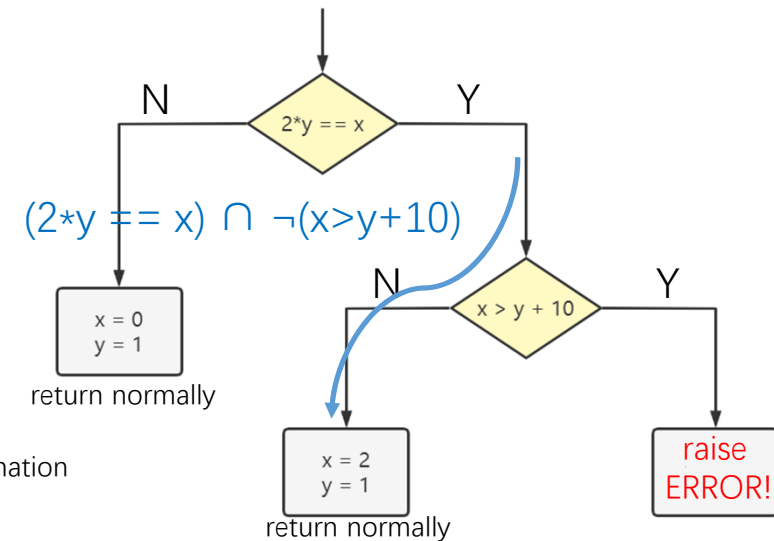


assume simulation engine use DFS.
For the first path, it extract constraint
and get a satisfied input

Symbolic Execution

- Use symbol to represent variables
- Simulate program execution
- Extract and solve constraints in execution path
- generate constraints in every path
- constraint solver: z3, cvc5

```
int main(){
  x = read_int();
  y = read_int();
  z = 2 * x;
  if (z == x){
    if (x > y+10)
      ERROR;
  }
}
```

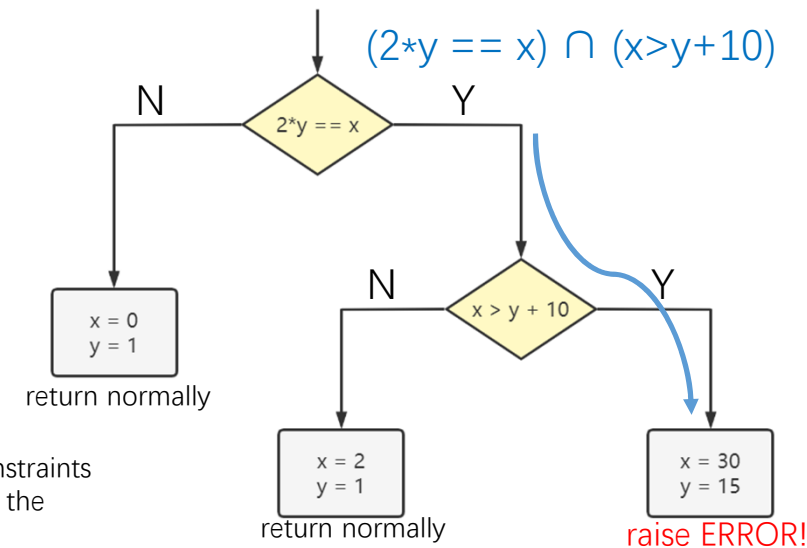


for second one it is a combination of two constraints

Symbolic Execution

- Use symbol to represent variables
- Simulate program execution
- Extract and solve constraints in execution path
- generate constraints in every path
- constraint solver: z3, cvc5

```
int main(){
  x = read_int();
  y = read_int();
  z = 2 * x;
  if (z == x){
    if (x > y+10)
      ERROR;
  }
}
```



similarly, it extract simple constraints
and can get a result to reach the
path that contains bug
linear problems can be solved by
constraint solver easily

Symbolic Execution

- angr: platform-agnostic binary analysis framework

- convert input to bit vector, simulate program instructions

Shoshitaishvili, Yan, et al. "Sok:(state of) the art of war: Offensive techniques in binary analysis." *2016 IEEE symposium on security and privacy (SP)*. IEEE, 2016.

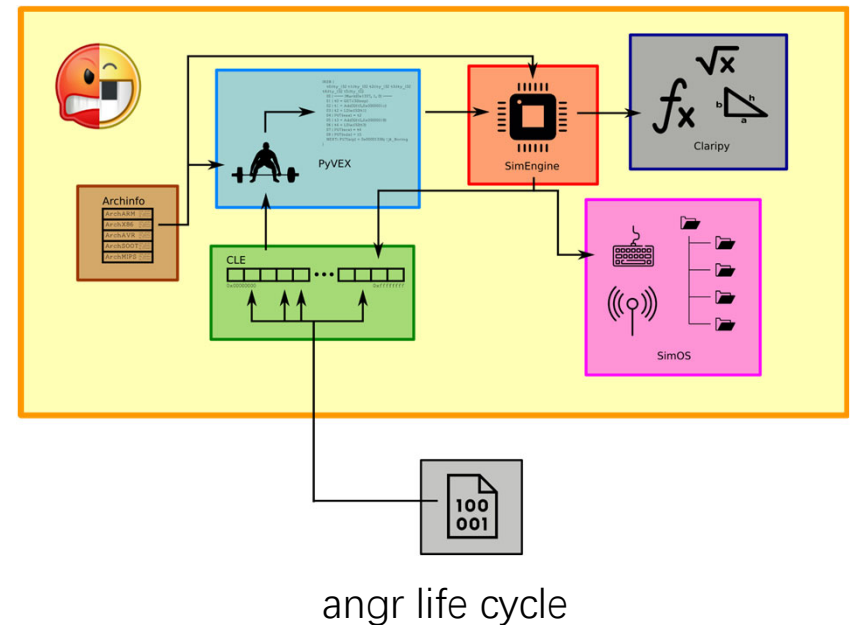
- KLEE: source code analysis framework

- compile from source code and make instrumentation

- <http://klee.doc.ic.ac.uk/> (try it online)

Cadar, Cristian, Daniel Dunbar, and Dawson R. Engler. "Klee: Unassisted and automatic generation of high-coverage tests for complex systems programs." *OSDI*. Vol. 8. 2008.

- Mayhem, Triton, etc...



Symbolic Execution

- Problem:

- state explosion when handling many branches
 - With **each if statement**, the number of possible branches might **double**. The growth of the problem is **exponential** with respect to the size of the program.
- cost many time solving large constraint (often a constraint will be solved many times)




Image source: <http://www.icodeguru.com/vc/10book/books/book3/chap6.htm>


Symbolic Execution

- Problem:
 - state exploitation when handling many branches
 - cost many time solving large constraint (often a constraint will be solved many times)
- Combination of two techniques: **Concolic Execution**

Concolic Execution

- Skip solving unnecessary constraints

- use fuzzing to accelerate
- use symbolic execution to reach deeper branches



- An example: Driller




Fig. 1. The nodes initially found by the fuzzer.




Fig. 2. The nodes found by the first invocation of concolic execution.




Fig. 3. The nodes found by the fuzzer, supplemented with the result of the first Driller run.





Fig. 4. The nodes found by the second invocation of concolic execution.

Concolic Execution

- Skip solving unnecessary constraints

- use fuzzing to accelerate
- use symbolic execution to reach deeper branches



- An example: Diller

```
// C1: check for hardcoded values
if (strcmp(header, "SECO") != 0)
    return ERROR;
check magic (fuzzer is hard to pass this check)
```

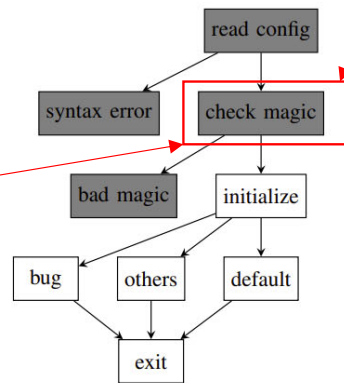


Fig. 1. The nodes initially found by the fuzzer.

after fuzz is hard to continue, only do symbolic execution on this node




Fig. 2. The nodes found by the first invocation of concolic execution.




Fig. 3. The nodes found by the fuzzer, supplemented with the result of the first Driller run.




Fig. 4. The nodes found by the second invocation of concolic execution.

Concolic Execution

- Skip solving unnecessary constraints

- use fuzzing to accelerate
- use symbolic execution to reach deeper branches

- An example: Diller





Fig. 1. The nodes initially found by the fuzzer.

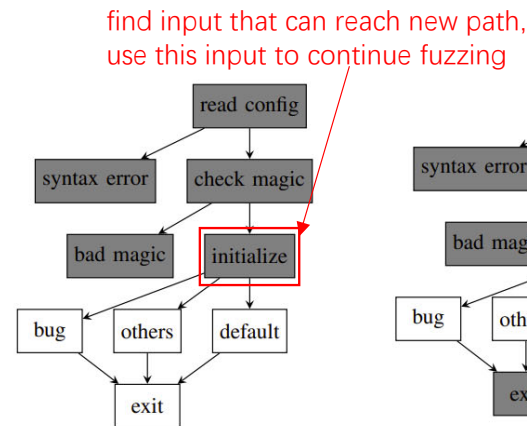


Fig. 2. The nodes found by the first invocation of concolic execution.

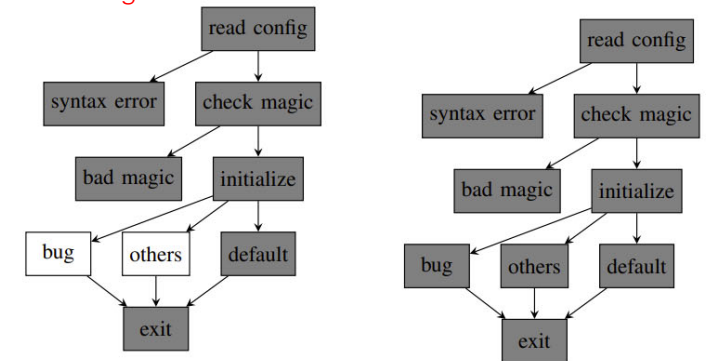


Fig. 3. The nodes found by the fuzzer, supplemented with the result of the first Driller run.





Fig. 4. The nodes found by the second invocation of concolic execution.

Concolic Execution

- Skip solving unnecessary constraints

- use fuzzing to accelerate
- use symbolic execution to reach deeper branches



- An example: Driller




Fig. 1. The nodes initially found by the fuzzer.

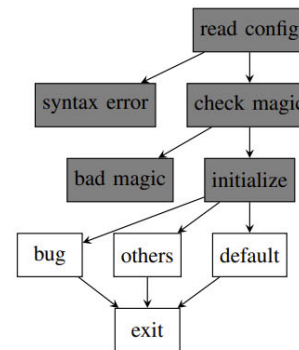


Fig. 2. The nodes found by the first invocation of concolic execution.

after fuzz for a while, do symbolic execution on this edge and try find input that can reach another path

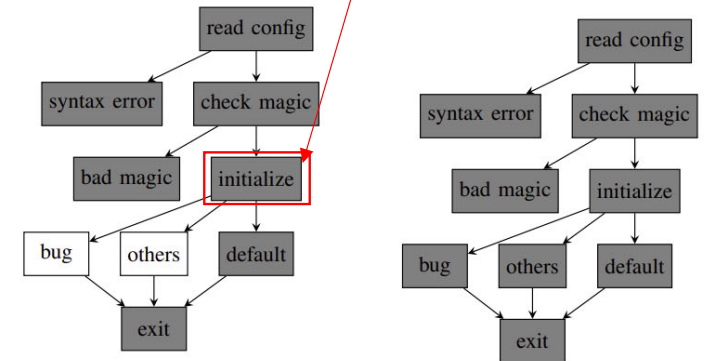


Fig. 3. The nodes found by the fuzzer, supplemented with the result of the first Driller run.





Fig. 4. The nodes found by the second invocation of concolic execution.

Concolic Execution

- Skip solving unnecessary constraints

- use fuzzing to accelerate
- use symbolic execution to reach deeper branches



- An example: Driller




Fig. 1. The nodes initially found by the fuzzer.




Fig. 2. The nodes found by the first invocation of concolic execution.

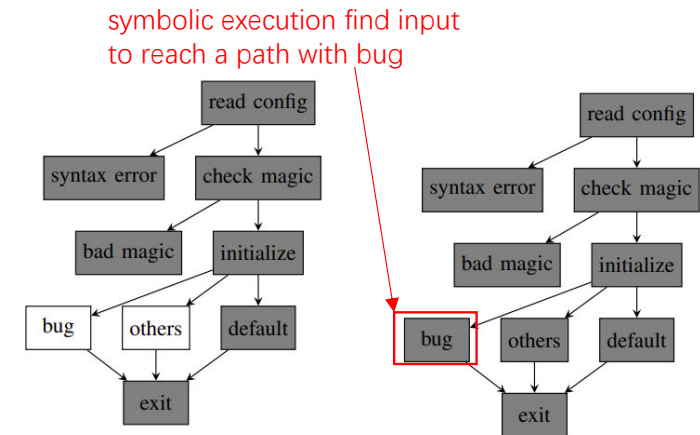


Fig. 3. The nodes found by the fuzzer, supplemented with the result of the first Driller run.

Fig. 4. The nodes found by the second invocation of concolic execution.

Summary:

- Fuzz can quickly find some bug of large real world program, but fuzz is hard to reach some complicated path
- We use symbolic execution to find more complicated bugs, but it may produce too many unnecessary states
- We can combine fuzzing and symbolic execution

Happy exploiting!

Extra Notes:

- Fuzz target not only contains source code and binary
- browser, blockchain, compiler, kernel... everything can be tested!
- There are some other program analyzing techniques like module checking