

Fuzzing project

- Steer clear of FFmpeg
- Look for evidence of fuzzing in the code repo

Fuzzing - last week

1. Basic fuzzing with random/long inputs
2. 'Dumb' mutational fuzzing
example: OCPP
3. Generational fuzzing aka grammar-based fuzzing
example: GSM
4. Whitebox fuzzing with SAGE
using symbolic execution

Still left for today

1. Code-coverage guided evolutionary fuzzing with afl
aka grey box fuzzing or 'smart' mutational fuzzing

**Coverage-guided evolutionary fuzzing
with afl
(American Fuzzy Lop)**



Evolutionary Fuzzing

Use **evolution**:

try random input mutations, and

observe the effect on some form of coverage, and

let only the interesting mutations evolve further

where “interesting” = resulting in ‘new’ execution paths

Aka **coverage-guided evolutionary greybox fuzzing**,

but terminology is a bit messy/non-standard.

afl

[<http://lcamtuf.coredump.cx/afl>]

- **Code instrumented** to observe execution paths:
 - if source code is available, by using modified compiler
 - if source code is not available, by running code in an emulator
- **Code coverage represented as a 64KB bitmap:**
each control flow jumps is mapped to a change in this bitmap
 - different executions could result in same bitmap, but chance is small
- **Mutation strategies include:** bit flips, incrementing/decrementing integers, using pre-defined interesting values (eg. 0, -1, MAX_INT,...) or user-supplied dictionary, deleting/combining/zeroing input blocks, ...
- The fuzzer forks the SUT to speed up the fuzzing
- **Big win:** no need to specify the input format, but still good coverage

afl's instrumentation of compiled code

Code is injected at every branch point in the code

```
cur_location = <SOME_RANDOM_NUMBER_FOR_THIS_CODE_BLOCK>;  
shared_mem[cur_location ^ prev_location]++;  
prev_location = cur_location >> 1;
```

where `shared_mem` is a 64 KB memory region

Intuition: for every jump from L_1 to L_2 a different byte in `shared_mem` is changed (increased).

Which byte is determined by random values chosen at compile time inserted at source and destination of every jump

american fuzzy lop 2.52b (dnsmasq)

```
process timing
  run time : 0 days, 20 hrs, 31 min, 27 sec
  last new path : 0 days, 0 hrs, 48 min, 28 sec
  last uniq crash : 0 days, 2 hrs, 22 min, 39 sec
  last uniq hang : none seen yet
cycle progress
  now processing : 3138* (92.05%)
  paths timed out : 0 (0.00%)
stage progress
  now trying : user extras (insert)
  stage execs : 509k/1.38M (36.79%)
  total execs : 29.4M
  exec speed : 464.9/sec
fuzzing strategy yields
  bit flips : 151/1.22M, 104/1.22M, 47/1.22M
  byte flips : 0/152k, 2/61.4k, 4/59.8k
  arithmetics : 133/3.47M, 0/1.04M, 0/286k
  known ints : 32/264k, 29/1.62M, 10/2.55M
  dictionary : 103/2.43M, 48/5.49M, 176/1.58M
  havoc : 1060/6.14M, 0/0
  trim : 40.91%/56.3k, 58.16%
map coverage
  map density : 0.34% / 4.51%
  count coverage : 2.92 bits/tuple
findings in depth
  favored paths : 686 (20.12%)
  new edges on : 1022 (29.98%)
  total crashes : 363 (12 unique)
  total tmouts : 54 (18 unique)
path geometry
  levels : 17
  pending : 2326
  pend fav : 7
  own finds : 1887
  imported : n/a
  stability : 100.00%
overall results
  cycles done : 3
  total paths : 3409
  uniq crashes : 12
  uniq hangs : 0
^C [cpu000:150%]
```

+++ Testing aborted by user +++

[+] We're done here. Have a nice day!

american fuzzy lop 2.52b (dnsmasq)

```
process timing
  run time : 0 days, 20 hrs, 31 min, 27 sec
  last new path : 0 days, 0 hrs, 48 min, 28 sec
  last uniq crash : 0 days, 2 hrs, 22 min, 39 sec
  last uniq hang : none seen yet
cycle progress
  now processing : 3138* (92.05%)
  paths timed out : 0 (0.00%)
stage progress
  now trying : user extras (insert)
  stage execs : 509k/1.38M (36.79%)
  total execs : 29.4M
  exec speed : 464.9/sec
fuzzing strategy yields
  bit flips : 151/1.22M, 104/1.22M, 47/1.22M
  byte flips : 0/152k, 2/61.4k, 4/59.8k
  arithmetics : 133/3.47M, 0/1.04M, 0/286k
  known ints : 32/264k, 29/1.62M, 10/2.55M
  dictionary : 103/2.43M, 48/5.49M, 176/1.58M
  havoc : 1060/6.14M, 0/0
  trim : 40.91%/56.3k, 58.16%
overall results
  cycles done : 3
  total paths : 3409
  uniq crashes : 12
  uniq hangs : 0
map coverage
  map density : 0.34% / 4.51%
  count coverage : 2.92 bits/tuple
findings in depth
  favored paths : 686 (20.12%)
  new edges on : 1022 (29.98%)
  total crashes : 363 (12 unique)
  total tmouts : 54 (18 unique)
path geometry
  levels : 17
  pending : 2326
  pend fav : 7
  own finds : 1887
  imported : n/a
  stability : 100.00%
^C [cpu000:150%]
```

+++ Testing aborted by user +++

[+] We're done here. Have a nice day!

afl statistics

- **total execs**
- **total paths**
- **(unique) crashes**
- **(unique) hangs**
- **cycles**


Cool example: learning the JPG file format

Fuzzing a program that expects a JPG as input, starting with 'hello world' as initial test input, afl can learn to produce legal JPG files

along the way producing/discovering error messages such as

- Not a JPEG file: starts with 0x68 0x65
- Not a JPEG file: starts with 0xff 0x65
- Premature end of JPEG file
- Invalid JPEG file structure: two SOI markers
- Quantization table 0x0e was not defined

and then JPGs like



[Source <http://lcamtuf.blogspot.nl/2014/11/pulling-jpegs-out-of-thin-air.html>]

Other strategies in evolutionary fuzzing

Instead of maximizing path/code coverage, we can also let inputs evolve to **maximize some other variable or property**

- Code may need to be instrumented to let fuzzer observe that property



Eg the **x-coordinate of Super Mario**

[Aschermann et al., *IJON: Exploring Deep State Spaces via Fuzzing*, IEEE S&P 2020]

<https://www.youtube.com/watch?v=3PyhXIHDkNI>

Conclusions

- Fuzzing is great technique to find (a certain class of) security flaws!
- If you ever write or use C(++) code, you should fuzz it.
- Challenge: getting good coverage fuzzing without too much effort

Successful approaches include

- White-box fuzzing based on symbolic execution with **SAGE**
- Evolutionary fuzzing aka coverage guided greybox fuzzing with **afl**
- *Does fuzzing makes sense for code in other programming languages?*
Yes, even if the kind of bugs found may have lower security impact.
- A more ambitious generation of tools not only tries to find security flaws, but also to then build exploits, eg. **angr**

To read (see links on the course page)

- Section 1 of technical white paper for afl
- Patrice Godefroid, *Fuzzing: Hack, Art, and Science* CACM 2020