

Designing New Operating Primitives to Improve Fuzzing Performance

Wen Xu, Sanidhya Kashyap, Changwoo Min*, and Taesoo Kim
Georgia Tech, Virginia Tech*

ACM CCS 2017

11.02.2017

Fuzzing becomes popular

```
american fuzzy lop 0.47b (readpng)

process timing
  run time : 0 days, 0 hrs, 4 min, 43 sec
  last new path : 0 days, 0 hrs, 0 min, 26 sec
  last uniq crash : none seen yet
  last uniq hang : 0 days, 0 hrs, 1 min, 51 sec
cycle progress
  now processing : 38 (19.49%)
  paths timed out : 0 (0.00%)
stage progress
  now trying : interest 32/8
  stage execs : 0/9990 (0.00%)
  total execs : 654k
  exec speed : 2306/sec
fuzzing strategy yields
  bit flips : 88/14.4k, 6/14.4k, 6/14.4k
  byte flips : 0/1804, 0/1786, 1/1750
  arithmetics : 31/126k, 3/45.6k, 1/17.8k
  known ints : 1/15.8k, 4/65.8k, 6/78.2k
  havoc : 34/254k, 0/0
  trim : 2876 B/931 (61.45% gain)

overall results
  cycles done : 0
  total paths : 195
  uniq crashes : 0
  uniq hangs : 1


map coverage
  map density : 1217 (7.43%)
  count coverage : 2.55 bits/tuple

findings in depth
  favored paths : 128 (65.64%)
  new edges on : 85 (43.59%)
  total crashes : 0 (0 unique)
  total hangs : 1 (1 unique)

path geometry
  levels : 3
  pending : 178
  pend fav : 114
  imported : 0
  variable : 0
  latent : 0
```



```
1 1 1 0 1
0 0 0 1 0
1 0 1 0 1
1 0 1 1 1
```




The dilemma of fuzzing


- How to produce an input that is more likely to trigger a vulnerability? (*fuzzing strategy*)
- *Our work: How to execute more inputs within a given time? (fuzzing performance)*
 - Save huge cost on computing resources in parallel fuzzing
 - No change in applied fuzzing strategies

Poor scalability of AFL

total executions/second



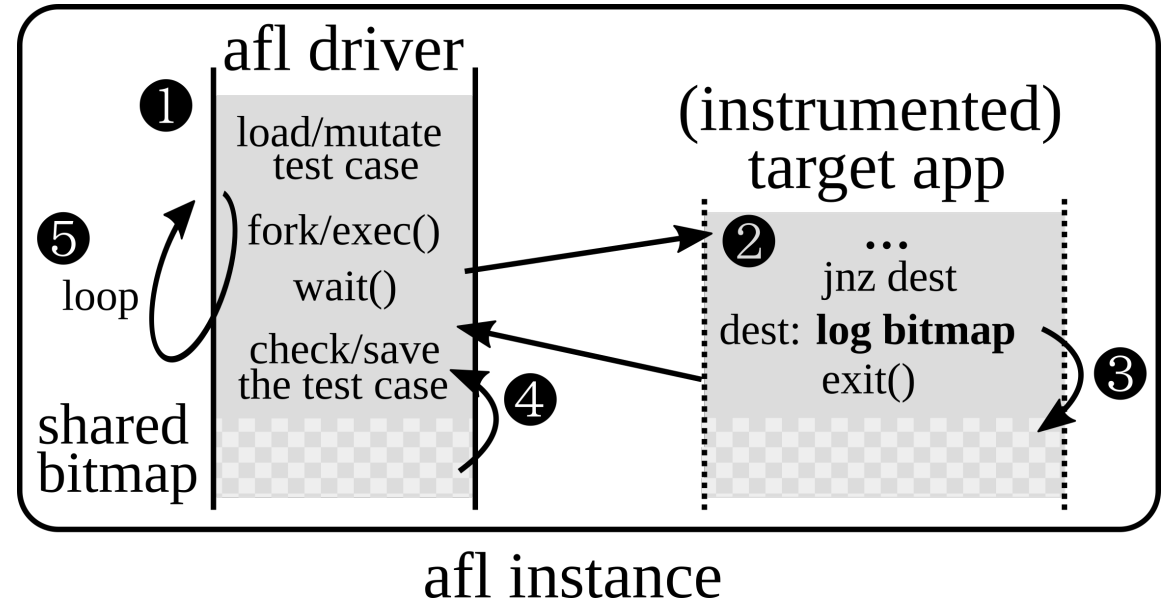
% time spent



AFL explained – single instance

Repeating

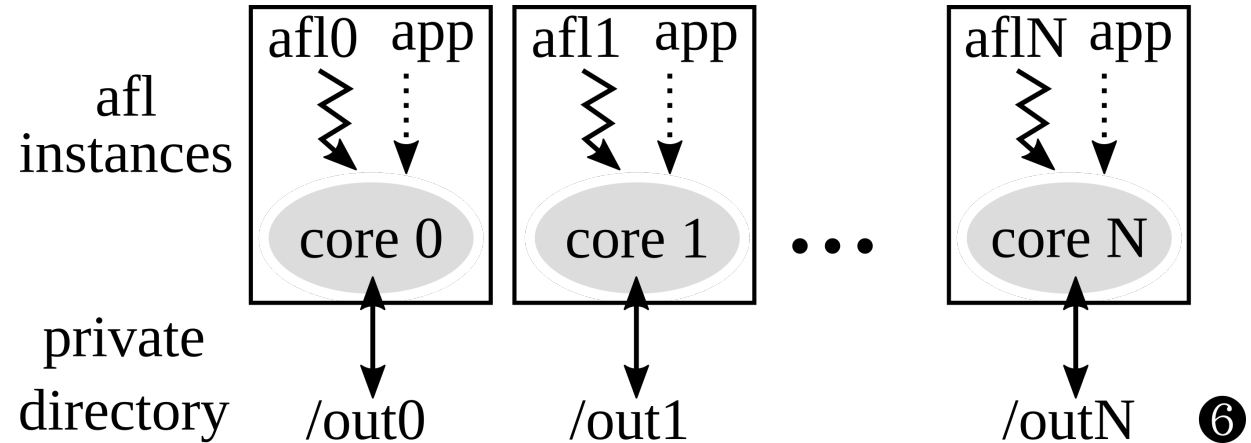
- (1) Reading and mutating inputs
- (2) Launching the target application
- (3) Executing and recording runtime coverage
- (4) Bookkeeping results



AFL explained – parallel fuzzing

Syncing phase

- (1) Scanning the private directories of other fuzzer instances
- (2) Executing unseen test cases
- (3) Copying to own directory if interesting



Fuzzers rely on non-scalable OS primitives

- Launching the target application
 - `fork()`
- Reading or writing test cases on the disk
 - **typical disk file system operations on small files**
- Syncing test cases from other fuzzer instances
 - **Directory scanning**
 - **`fork()` for test case re-executions**

I. fork() to clone new target instances

- fork() is generally designed to duplicate the state of any running process
- In terms of fuzzing on multicores, fork() involves

Redundant operations

- Duplicating virtual memory space
- Duplicating files, sockets, credentials

...

Non-scalable operations

- Updating the reverse mapping
- Stressing the global memory allocator
- Scheduling the new task

...

II. Managing test cases through the disk file system

open/creat

generate test cases
on the disk



creating files in
a private directory

write

flush interesting test cases
to the disk



writing a small files in
a private directory

heavy modifications on the file
system metadata in critical sections
which are **not scalable**

III. Syncing test cases from other instances

Directory enumeration

- **non-linearly increase**

number of fuzzers × number of test cases

- **interfere with the running fuzzer**

Directory read and write cannot be performed concurrently

Test case re-execution

- **redundant**

The runtime coverage information was achieved before

Solutions

- *General* operating primitives specialized for fuzzers
 - Snapshot() system call
 - *A lightweight, scalable fork() substitute for fuzzing*
 - Dual file system service
 - Shared in-memory test case log

Snapshot() system call

Before fuzzing,

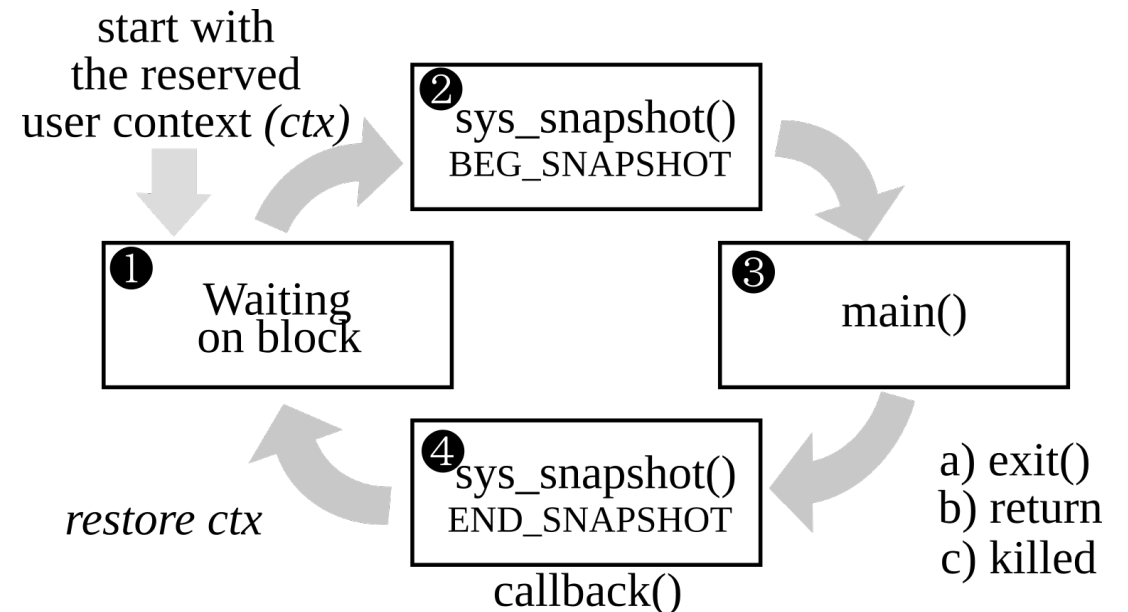
- Saving memory and file information
- Manipulating PTE of rw pages to ro (CoW)

During fuzzing,

- Demanding page copy when memory write triggers page fault handler

After fuzzing,

- Recovering memory and file information
- Recovering copied page data
- Returning to the starting point



Snapshot() system call

- Compared with fork()
 - No copies of numerous kernel data structures
 - No new stack area allocation for the new process
 - No stress on the kernel memory allocator
 - No scheduling cost

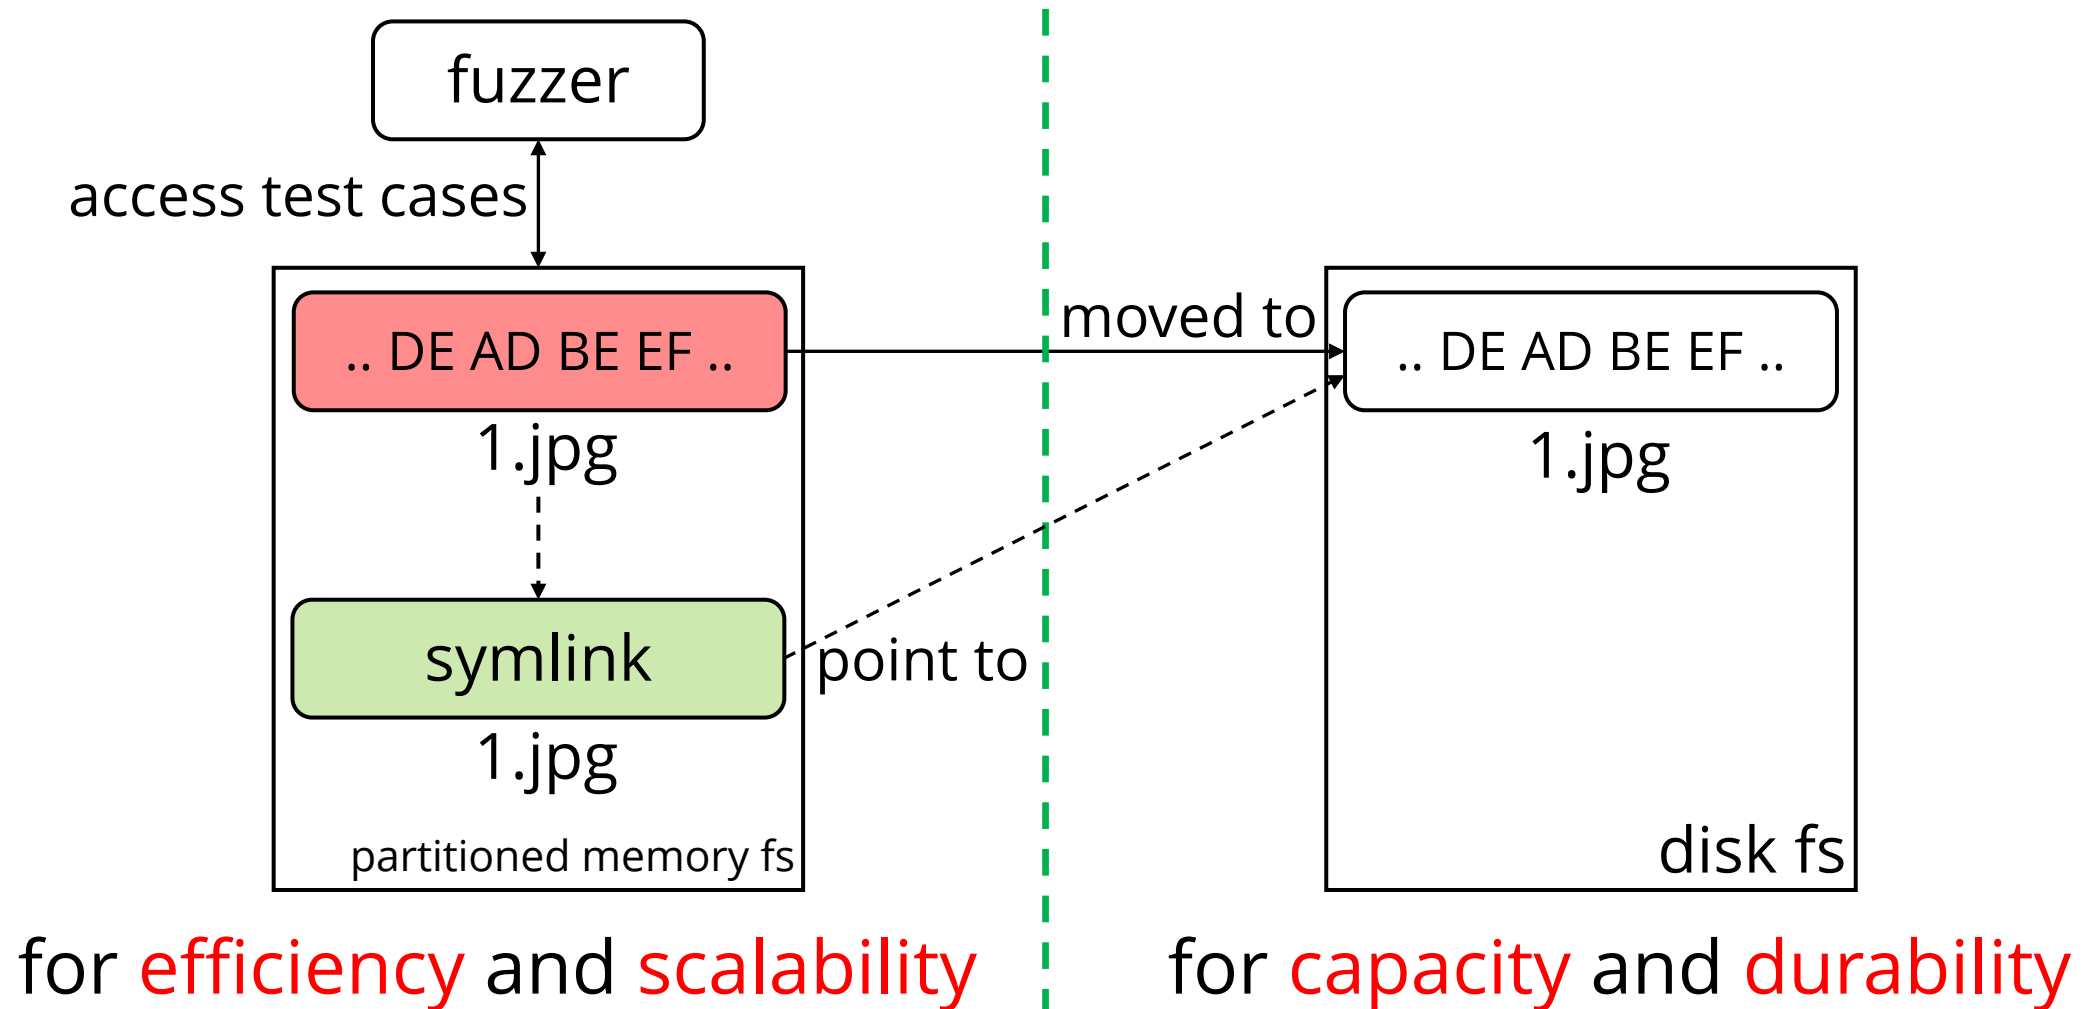
Solutions

- *General* operating primitives specialized for fuzzers
 - Snapshot() system call
 - **Dual file system service**
 - *A two-level tiering of file systems ensuring efficiency and deferred durability*
 - Shared in-memory test case log

Dual file system service

- *Observation:* neither the fuzzer instance nor the target instances requires strong consistency provided by the disk file system
 - Fuzzers can always reproduce lost test cases upon unexpected failures
 - We introduce memory file system and trade off between *consistent storage* and *fuzzing performance*

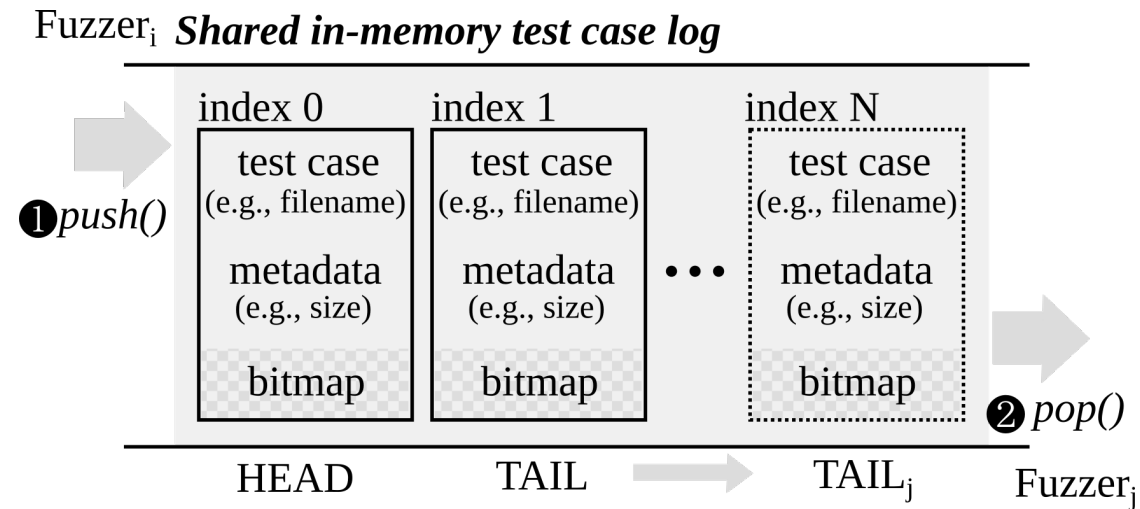
Dual file system service



Solutions

- *General* operating primitives specialized for fuzzers
 - Snapshot() system call
 - Dual file system service
 - **Shared in-memory test case log**
 - *A circular queue for efficient collaborative fuzzing*

Shared in-memory test case log



- No directory enumeration: `pop()` to examine test cases from neighbors
- No test case re-execution: direct reference on the bitmap
- No contention: a lock-free design

Applicability of techniques


Fuzzers	Snapshot	Dual FS	In-memory log
AFL	✓	✓	✓
AFLFast	✓	✓	✓
Driller	✓	✓	✓
LibFuzzer	-	✓	✓
Honggfuzz	-	✓	✓
VUzzer	✓	✓	✓
Choronzon	✓	✓	✓
IFuzzer	✓	✓	✓
jsfunfuzz	✓	✓	-
zzuf	✓	-	-

Implementation


- A new x86_64 system call snapshot() (750 LoC)
- A library for the shared in-memory test case log (100 LoC)
- A dual file system service daemon (100 LoC)
- We applied our new primitives to AFL (400 LoC) and LibFuzzer (200LoC)

Evaluation – shared test case log

total fuzzing executions/second




% time spent




Evaluation – snapshot() system call

total fuzzing executions/second




Process spawns/second (log scale)

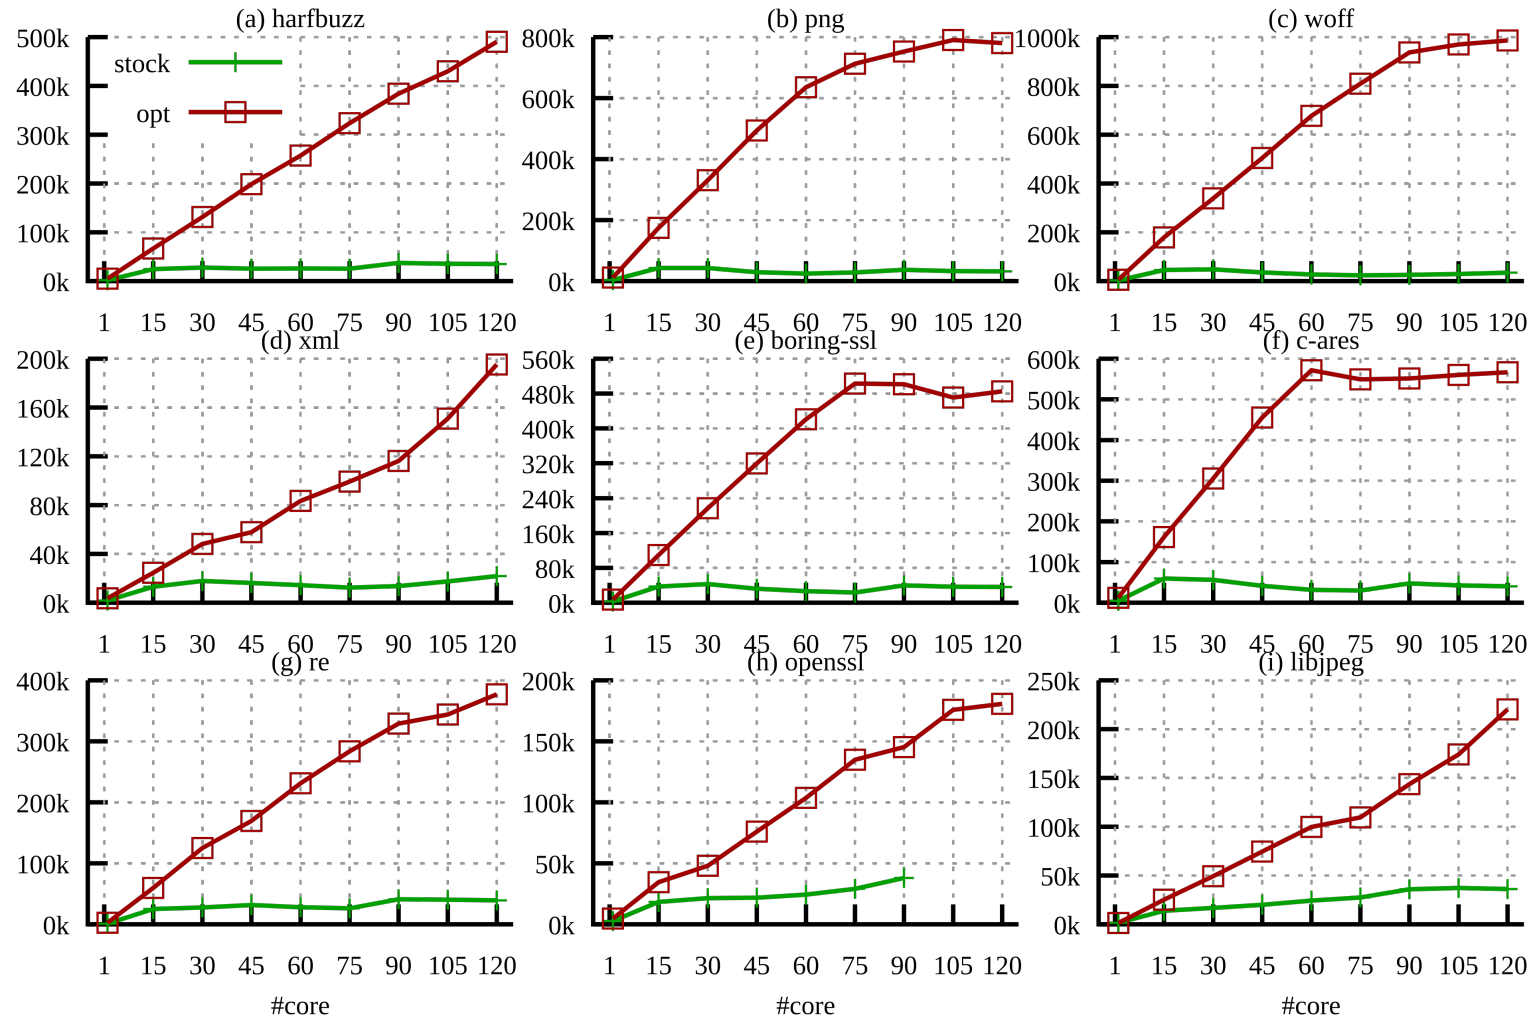


Evaluation – file system service in fuzzing

total fuzzing executions/second



Evaluation - AFL

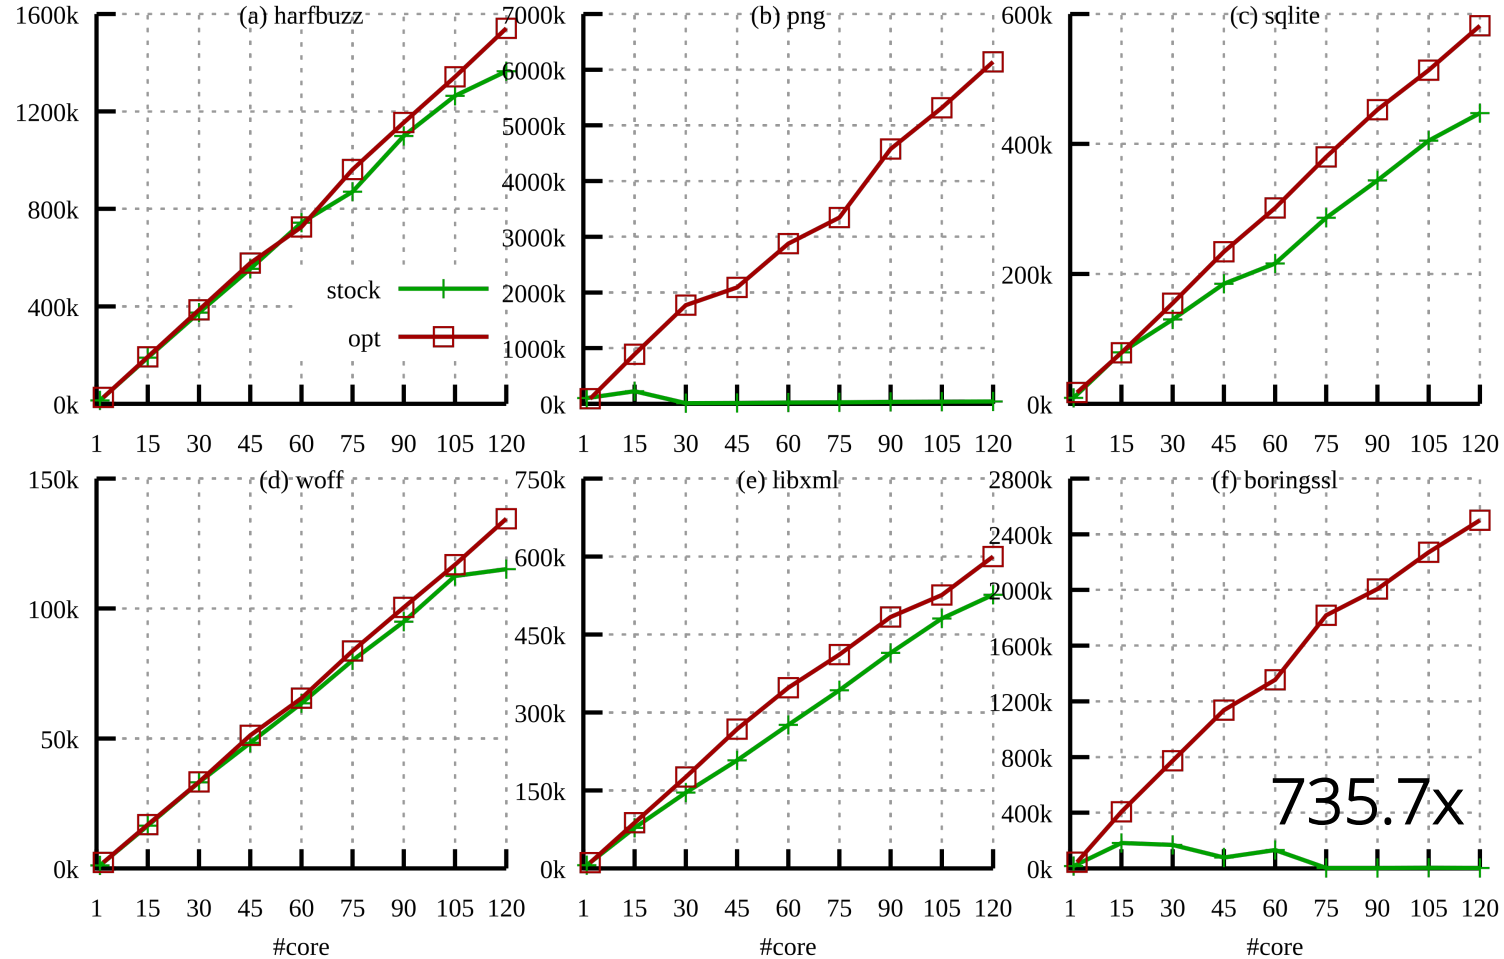


7.7x on 30 cores

25.9x on 60 cores

28.9x on 120 cores

Evaluation - LibFuzzer



Conclusion

- Current fuzzers are not at all scalable on modern OSes with manycore architectures.
- The underlying system components heavily relied on by the fuzzer degrade its scalability.
- New operating primitives specially designed for fuzzing can largely improve the performance and scalability for the state-of-the-art fuzzers.

Open source at <https://github.com/sslab-gatech/perf-fuzz>

Supported by 

Thanks for listening!
Wen Xu (wen.xu@gatech.edu)