

# Zeror: Speed Up Fuzzing with Coverage-sensitive Tracing and Scheduling

**Chijin Zhou**<sup>1</sup>, Mingzhe Wang<sup>1</sup>, Jie Liang<sup>1</sup>,  
Zhe Liu<sup>2</sup>, Yu Jiang<sup>1</sup>

<sup>1</sup>School of Software, Tsinghua University, Beijing, China

<sup>2</sup>Computer Science and Technology, NUAA, Nanjing, China



# Fuzzing is popular

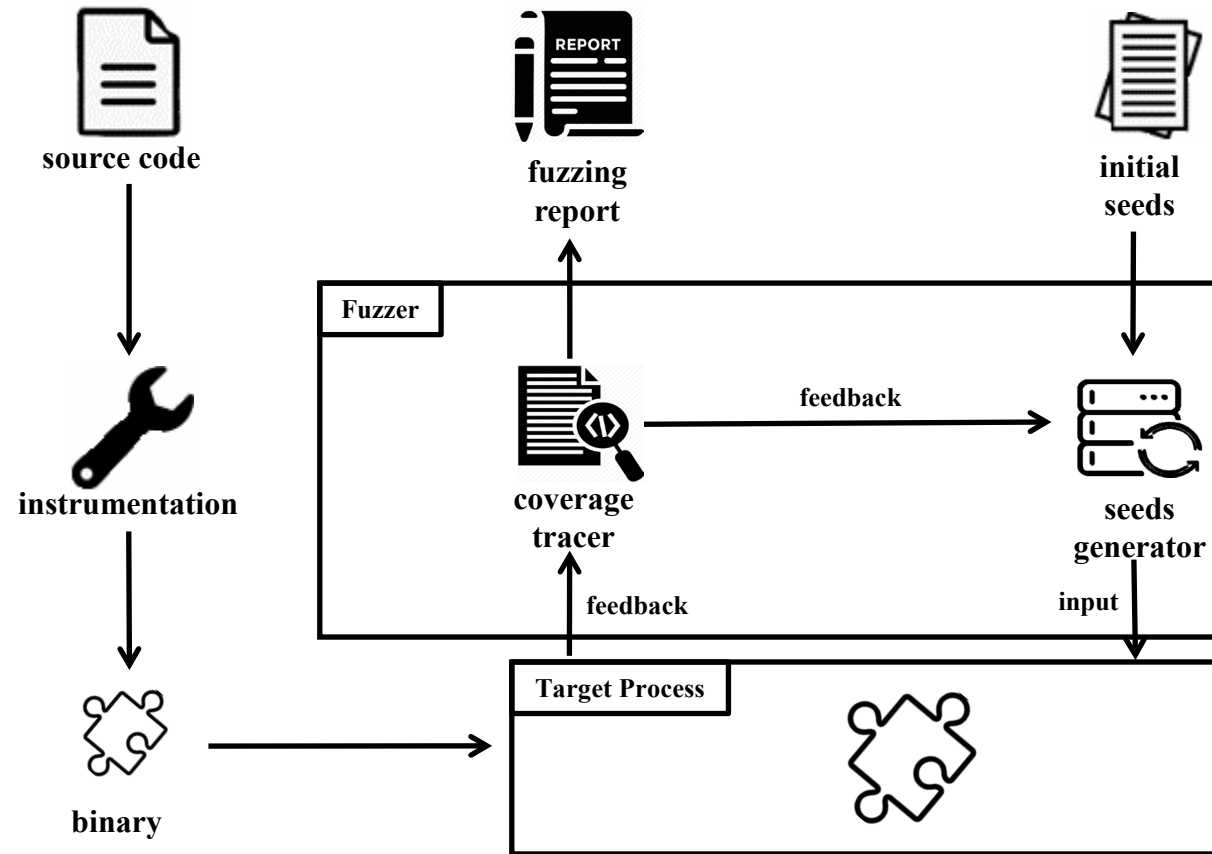
- Fuzzing is widely used for software vulnerability
  - Project Springfield
  - OSS-Fuzz
    - Has found more than 16,000 bugs



- Some fuzzed projects

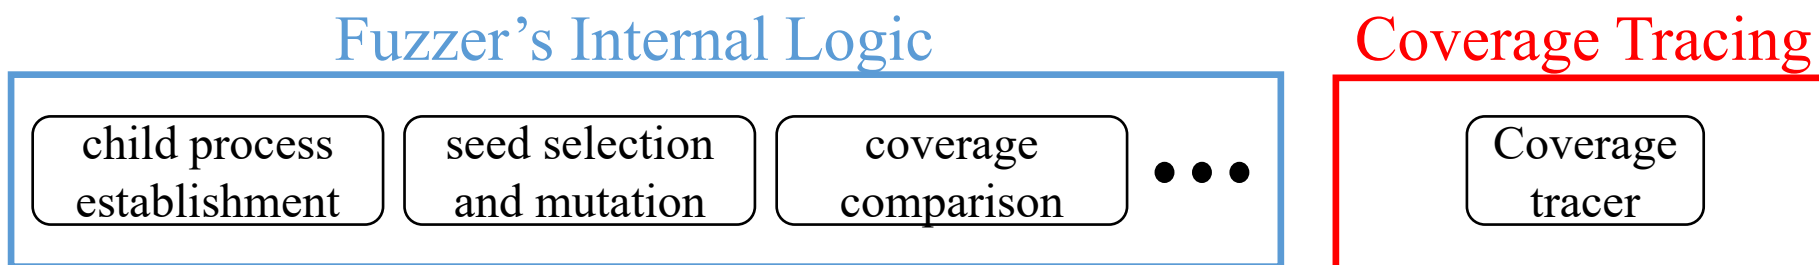


# General Workflow of Coverage-guided Fuzzing



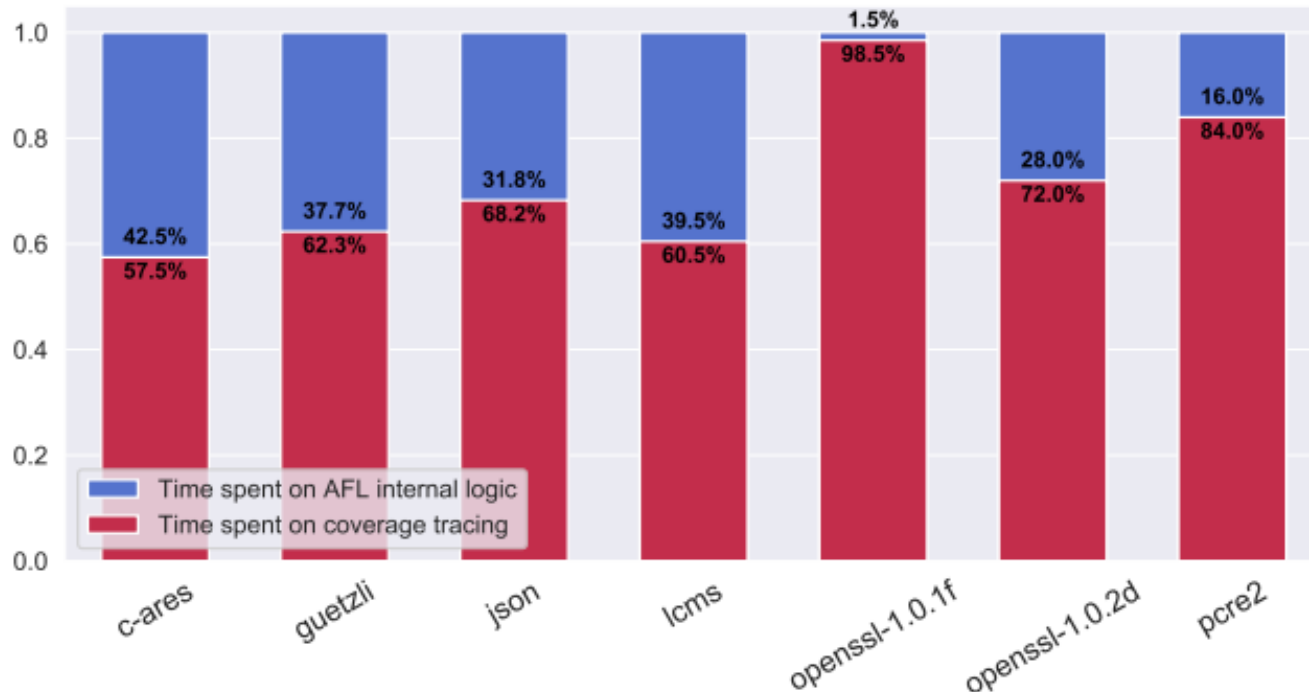
# Runtime of Coverage-guided Fuzzing

Take American Fuzzy Lop (AFL) as example, the fuzzer's runtime consists of two parts:



# Limitation of Coverage-guided Fuzzing

**Observation:** tracing coverage is costly




AFL spends an average of **71.85%** and up to **98.5%** of its runtime to trace coverage

**Figure 3: Percentage of internal logic execution time and edge level coverage tracing time in AFL.**


# Limitation of Coverage-guided Fuzzing

**Observation:** tracing coverage is costly

## Fuzzer's Internal Logic



## Coverage Tracing



Average Cost:

28.15%


71.85%

# Focus of This Paper


**Target:** boost fuzzing speed while preserve fine-grained coverage collection



## Fuzzer's Internal Logic



## Coverage Tracing



Average Cost:

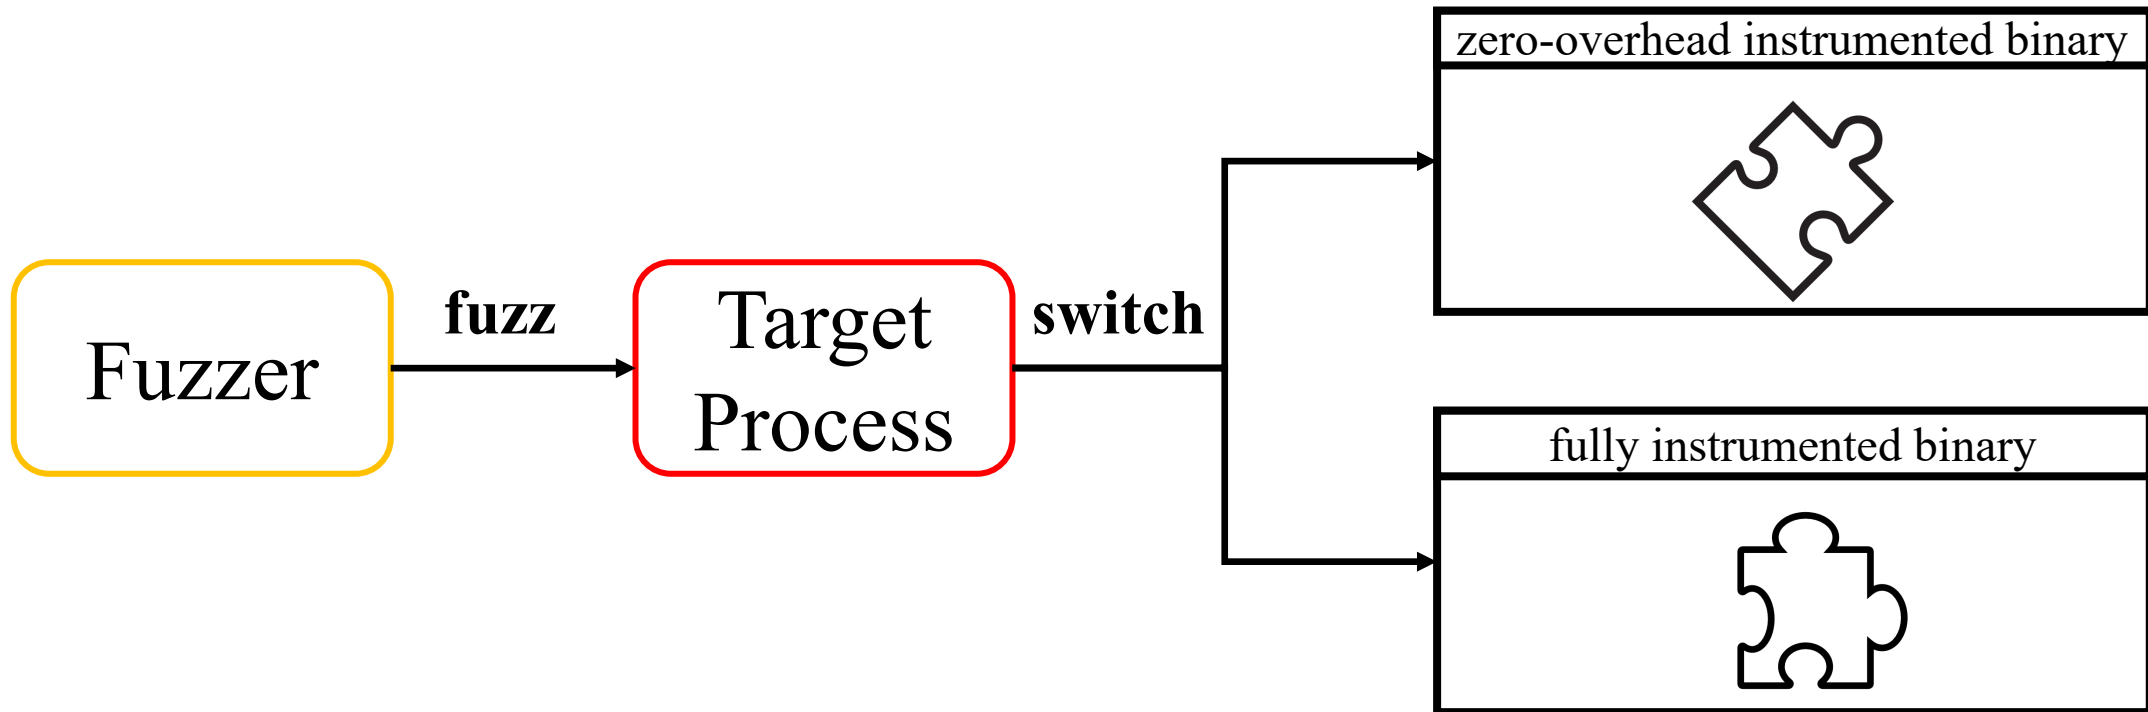
28.15%

71.85%

# Our Solution: Zeror

Main idea:

Switching between diversely-instrumented binaries





# Our Solution: Zeror

## Main idea:

Switching between diversely-instrumented binaries.

## Approaches:

- A **self-modifying tracing** mechanism to provide a zero-overhead instrumentation for coverage collection
- A **real-time scheduling** mechanism to support adaptive switch between the zero-overhead instrumented binary and the fully instrumented binary for better vulnerability detection

# Self-modifying Tracing

## Key insight:

Instrument program in edge level, and dynamically remove visited instrumentation points during fuzzing process

## Problems:

- (1) How to instrument program
- (2) How to remove visited instrumentation points


# Self-modifying Tracing

How to instrument program


**Solution:** add dummy block to critical edge

```
void foo(int *a)
{
  if (a)
    *a = 0;
}
```

(a) code



(b) basic-block level



(c) edge level

# Self-modifying Tracing

How to remove visited instrumentation points


**Solution:** self-modifying its instructions during fuzzing process

## Compilation Stage

① Construct CFG

② Generate binary file and record target addresses

```
int foo(int a, int b)
{
    if (a > b)
        return 1;
    else if (a < b)
        return 2;
    else
        return 3;
}
```



Address	Binary Codes	Assembly
0x2b1980	55	push %rbp
0x2b1981	48 89 e5	mov %rsp,%rbp
0x2b1984	b8 01 00 00 00	mov \$0x1,%eax
0x2b1989	39 f7	cmp %esi,%edi
0x2b198b	7f 0c	jg 2b1999
0x2b198d	b8 02 00 00 00	mov \$0x2,%eax
0x2b1992	7c 05	jl 2b1999
0x2b1994	b8 03 00 00 00	mov \$0x3,%eax
0x2b1999	5d	pop %rbp
0x2b199a	c3	retq

## Runtime Stage

③ Execute binary and inject breakpoints

④ Recover binary code when trigger interrupt

Address	Binary Codes	Assembly
0x2b1980	cc	int3 ...
0x2b1981	48 89 e5	mov %rsp,%rbp
0x2b1984	b8 01 00 00 00	mov \$0x1,%eax
0x2b1989	39 f7	cmp %esi,%edi
0x2b198b	7f 0c	jg 2b1999
0x2b198d	cc 02 00 00 00	int3 ...
0x2b1992	7c 05	jl 2b1999
0x2b1994	cc 03 00 00 00	int3 ...
0x2b1999	cc	int3 ...
0x2b199a	c3	retq

Address	Binary Codes	Assembly
0x2b1980	55	push %rbp
0x2b1981	48 89 e5	mov %rsp,%rbp
0x2b1984	b8 01 00 00 00	mov \$0x1,%eax
0x2b1989	39 f7	cmp %esi,%edi
0x2b198b	7f 0c	jg 2b1999
0x2b198d	cc 02 00 00 00	int3 ...
0x2b1992	7c 05	jl 2b1999
0x2b1994	cc 03 00 00 00	int3 ...
0x2b1999	cc	int3 ...
0x2b199a	c3	retq

# Binary-switching Scheduling

## Key insight:

Estimate fuzzing efficiencies of diversely-instrumented binaries, and switch to high-efficiency binary at set interval

## Problem:

How to estimate fuzzing efficiencies

# Binary-switching Scheduling

For a binary, the efficiency at time period  $t$  is defined as

$$e_t = \frac{I_t}{T_t}$$

the number of interesting seeds

time spent on fuzzing

$$= \frac{I_t}{M_t} * \frac{M_t}{T_t} = r_t * S$$

the number of executions

average execution speed (constant)

# Binary-switching Scheduling

For a binary, the efficiency at time period  $t$  is defined as

$$\begin{aligned} e_t &= \frac{I_t}{T_t} \\ &= \frac{I_t}{M_t} * \frac{M_t}{T_t} = r_t * S \end{aligned}$$

the number of interesting seeds

time spent on fuzzing

the number of executions

average execution speed (constant)

# Binary-switching Scheduling

## Procedure:

1. Collect statistical data (i.e. the number of interesting seeds, the number of executions and the time spent on fuzzing) of each binary
2. Use empirical Bayesian to estimate interesting rate  $\hat{r}_t$
3. Calculate the efficiency of each binary
4. Choose high-efficiency binary as optimal target



# Binary-switching Scheduling

Specially, to smooth time-varying observed data, we leverage **exponential smoothing** to calculate the smoothed number of interesting seeds:

$$I_i = \begin{cases} I'_i & i = 1 \\ \gamma I'_i + (1 - \gamma)I_{i-1} & i > 1 \end{cases}$$

# Evaluation

## 1. Efficiency of Zeror

Project	average execution time for each test case ( $\mu$ s)				number of covered branches			
	AFL	AFL+INSTRIM	AFL+Untracer	AFL+Zeror	AFL	AFL+INSTRIM	AFL+Untracer	AFL+Zeror
boringssl	96.69	69.68	N/A	33.05	2661	<b>2694</b>	N/A	2549
c-ares	43.34	25.42	13.95	16.32	<b>57</b>	<b>57</b>	55	<b>57</b>
freetype2	44.68	25.17	25.13	20.33	8255	9268	7007	<b>10059</b>
guetzli	99.92	67.98	45.80	41.00	4757	4845	4748	<b>4987</b>
harfbuzz	149.82	80.36	66.06	55.73	8148	8048	7195	<b>9168</b>
json	145.82	100.03	64.33	98.39	1315	1333	1152	<b>1346</b>
lcms	97.71	70.92	44.18	63.96	2115	<b>2244</b>	1436	2077
libarchive	193.44	112.50	112.90	112.72	1208	1119	1082	<b>1618</b>
libjpeg	1469.47	668.96	261.30	337.36	2364	2564	2399	<b>2857</b>
libpng	15.34	5.48	5.27	7.54	1092	1096	1029	<b>1140</b>
libssh	638.00	340.52	309.62	309.29	<b>867</b>	<b>867</b>	<b>867</b>	<b>867</b>
libxml2	268.07	135.05	N/A	88.13	4063	4318	N/A	<b>4745</b>
llvm-libcxxabi	137.61	81.61	43.75	42.04	6488	6005	6000	<b>7012</b>
openssl-1.0.1f	3418.66	1998.27	N/A	1948.43	4748	6745	N/A	<b>7372</b>
openssl-1.0.2d	161.09	92.48	N/A	63.23	1825	<b>1828</b>	N/A	1769
openssl-1.1.0c	210.70	89.74	N/A	50.60	<b>1712</b>	1711	N/A	1658
openthread	145.51	91.17	64.80	85.16	3561	3537	3279	<b>3591</b>
pcre2	199.12	102.21	53.86	49.11	<b>6890</b>	6888	6597	<b>6890</b>
proj4	23.22	14.24	8.47	7.86	2541	2584	2347	<b>3886</b>
re2	640.24	391.97	260.19	235.40	4608	4647	4533	<b>4725</b>
sqlite	221.18	160.84	136.01	141.40	1892	<b>1997</b>	1986	1972
vorbis	96.14	58.08	36.45	25.48	2035	<b>2152</b>	1817	2079
woff2	31.55	20.12	11.80	8.67	2119	2152	1453	<b>2157</b>
wpantund	1921.02	2019.62	1544.89	1789.23	7959	7892	7802	<b>8781</b>
<b>Zeror improvement</b>	<b>+159.80%</b>	<b>+50.70%</b>	<b>-0.46%</b>		<b>+10.14%</b>	<b>+6.82%</b>	<b>+20.84%</b>	

# Evaluation

## 1. Efficiency of Zeror

**Table 3: Time to expose known bugs,  $\infty$  denotes the fuzzer cannot expose the known bugs in 6 hours and the projects whose bugs can not be triggered by any fuzzer are removed.**

Project	AFL	AFL+INSTRIM	AFL+Untracer	AFL+Zeror
c-ares	<b>8</b>	26	842	<b>8</b>
guetzli	$\infty$	$\infty$	16257	<b>6001</b>
json	<b>5</b>	<b>5</b>	<b>5</b>	<b>5</b>
lcms	20679	$\infty$	11827	<b>10953</b>
llvm-libcxxabi	788	2197	2347	<b>709</b>
openssl-1.0.1f	<b>19</b>	<b>19</b>	$\infty$	21
openssl-1.0.2d	8716	6877	$\infty$	<b>6013</b>
pcre2	822	1375	6095	<b>439</b>
re2	$\infty$	$\infty$	$\infty$	<b>8194</b>
woff2	3565	<b>1535</b>	$\infty$	3260

# Evaluation

## 2. Scalability of Zeror

**Table 4: Time to expose known bugs, and the projects whose bugs cannot be triggered by them in 6 hours are removed.**


Project	MOPT	MOPT+Zeror
c-ares	<b>8</b>	<b>8</b>
json	<b>5</b>	<b>5</b>
llvm-libcxxabi	1818	<b>761</b>
openssl-1.0.1f	31	<b>21</b>
openssl-1.0.2d	1633	<b>1320</b>
pcre2	1944	<b>968</b>
woff2	3767	<b>3196</b>

# Evaluation


## 3. Evaluation of Individual Components

self-modifying tracing:

- 13.74x faster than Untracer when erasing instrumentation points
- helps fuzzer cover more branches compared with Untracer



(a) Average time taken for different methods to erase instrumentation points (lower is better).



(b) Relative covered branches improvement of Zeror- compared with Untracer.

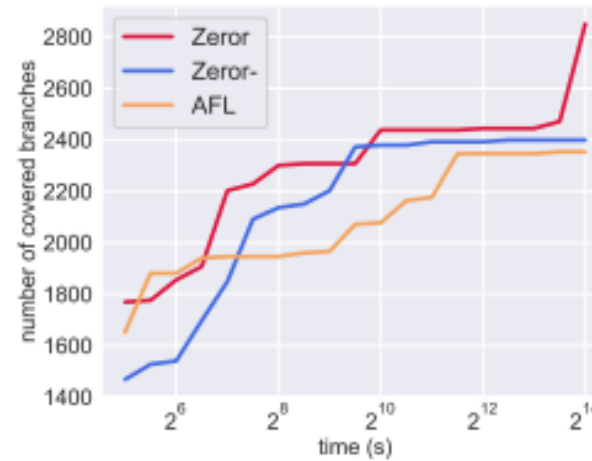
Figure 8: Comparison between Zeror- and Untracer.

# Evaluation


## 3. Evaluation of Individual Components

binary-switching scheduling:

- help fuzzer cover more branches



(a) libjpeg



(b) harfbuzz

Figure 9: Branches covered over time with different configurations. The x-axis is on a logarithmic scale.

# Conclusion

We propose a novel speed-up fuzzing framework Zeror:

It is made up of two parts: (1) zero-overhead instrumentation (2) real-time scheduling.

It helps fuzzers speed up fuzzing process, further increase covered branches and discovered bugs.

It is easy to be complemented to other orthogonal fuzzing optimizations.

# Thank You

If you have any questions, please send emails to

[zcj18@mails.tsinghua.edu.com](mailto:zcj18@mails.tsinghua.edu.com)