

**JUNE 23-27, 2024**

MOSCONE WEST CENTER  
SAN FRANCISCO, CA, USA

# PathFuzz: Broadening Fuzzing Horizons with Footprint Memory for CPUs

Yinan Xu, Sa Wang, Dan Tang, Ninghui Sun, Yungang Bao

<sup>1</sup>Institute of Computing Technology, Chinese Academy of Sciences (ICT, CAS)

<sup>2</sup>University of Chinese Academy of Sciences (UCAS)

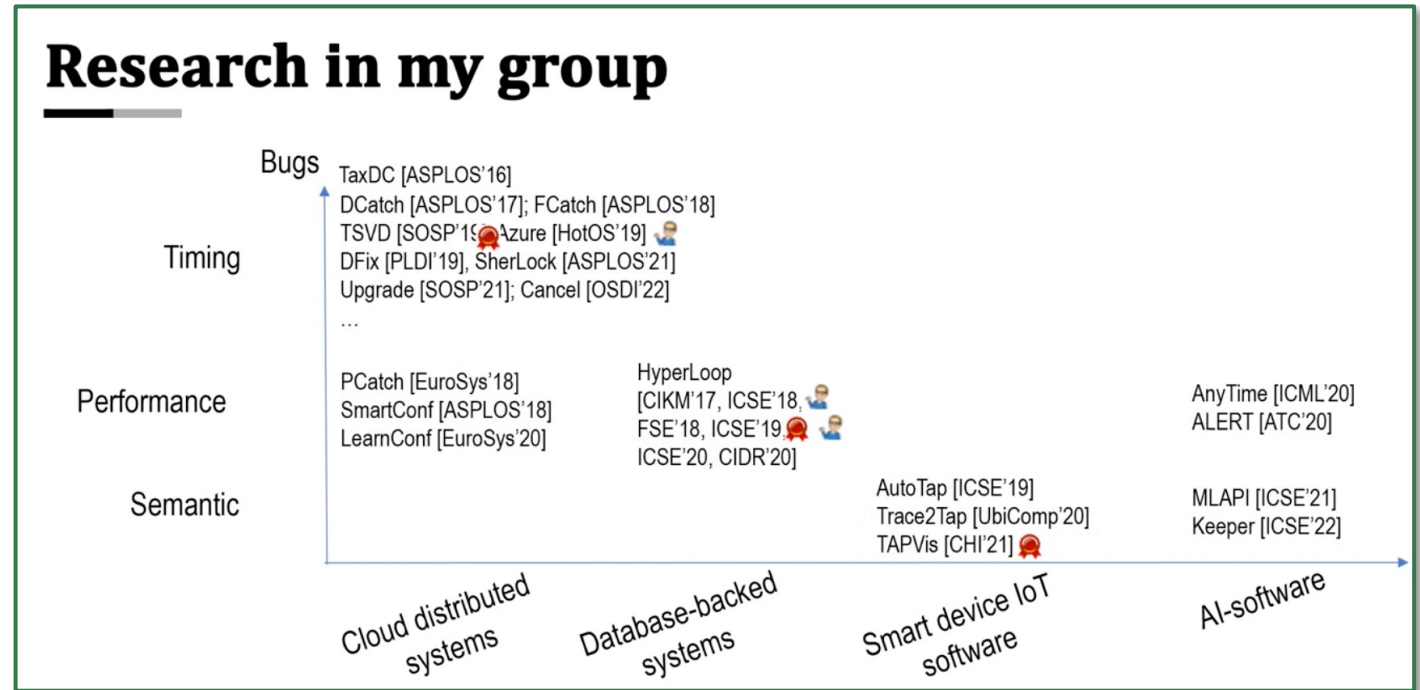
<sup>3</sup>Beijing Institute of Open Source Chip (BOSC)



# Open-Source Software and Software Testing



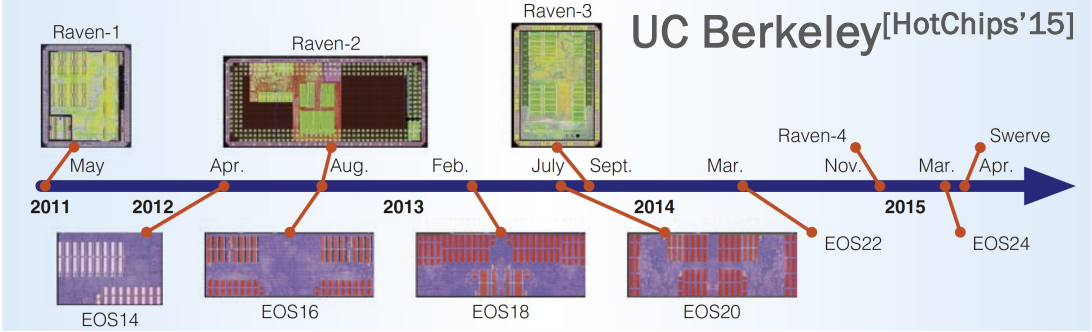
- Software testing research *greatly benefits from* open-source software
  - And the vice versa!

- Linux
- FreeBSD
- MySQL
- PostgreSQL
- Apache
- Mozilla
- OpenOffice
- .....

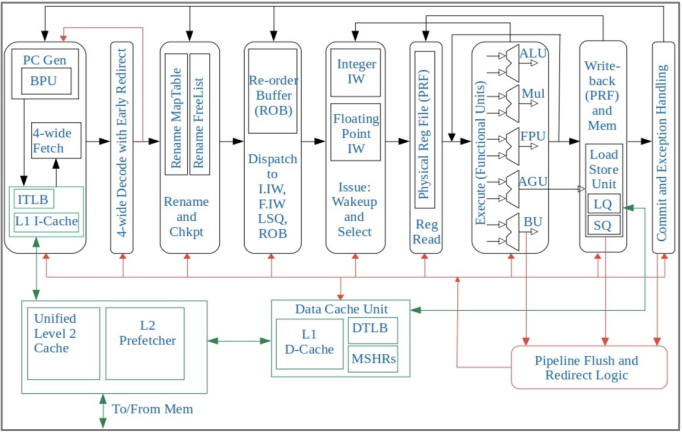
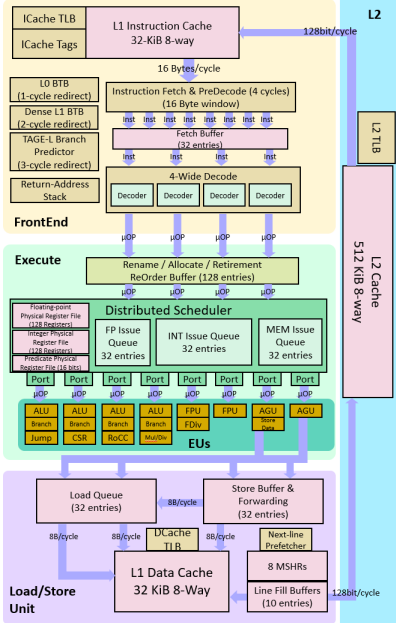


Source: Shan Lu, 15 Years of Learning from Mistakes in Building System Software, 22<sup>nd</sup> ChinaSys Workshop.

# The Era of Open-Source Chip



XiangShan [MICRO'22]



Shakti

SonicBOOM [CARRV'20]



Beihai [Intelligent Computing]

# Hardware Design Verification (DV) is Challenging

≈ (Software) Testing + Verification

50%

Increase in *design*  
engineers since 2007

146%

Increase in *verification*  
engineers since 2007

>50%

Median project time  
spent in *verification*

Source: <https://blogs.sw.siemens.com/verificationhorizons/2022/12/12/part-8-the-2022-wilson-research-group-functional-verification-study/>

# The Lockstep Between Design and Verification

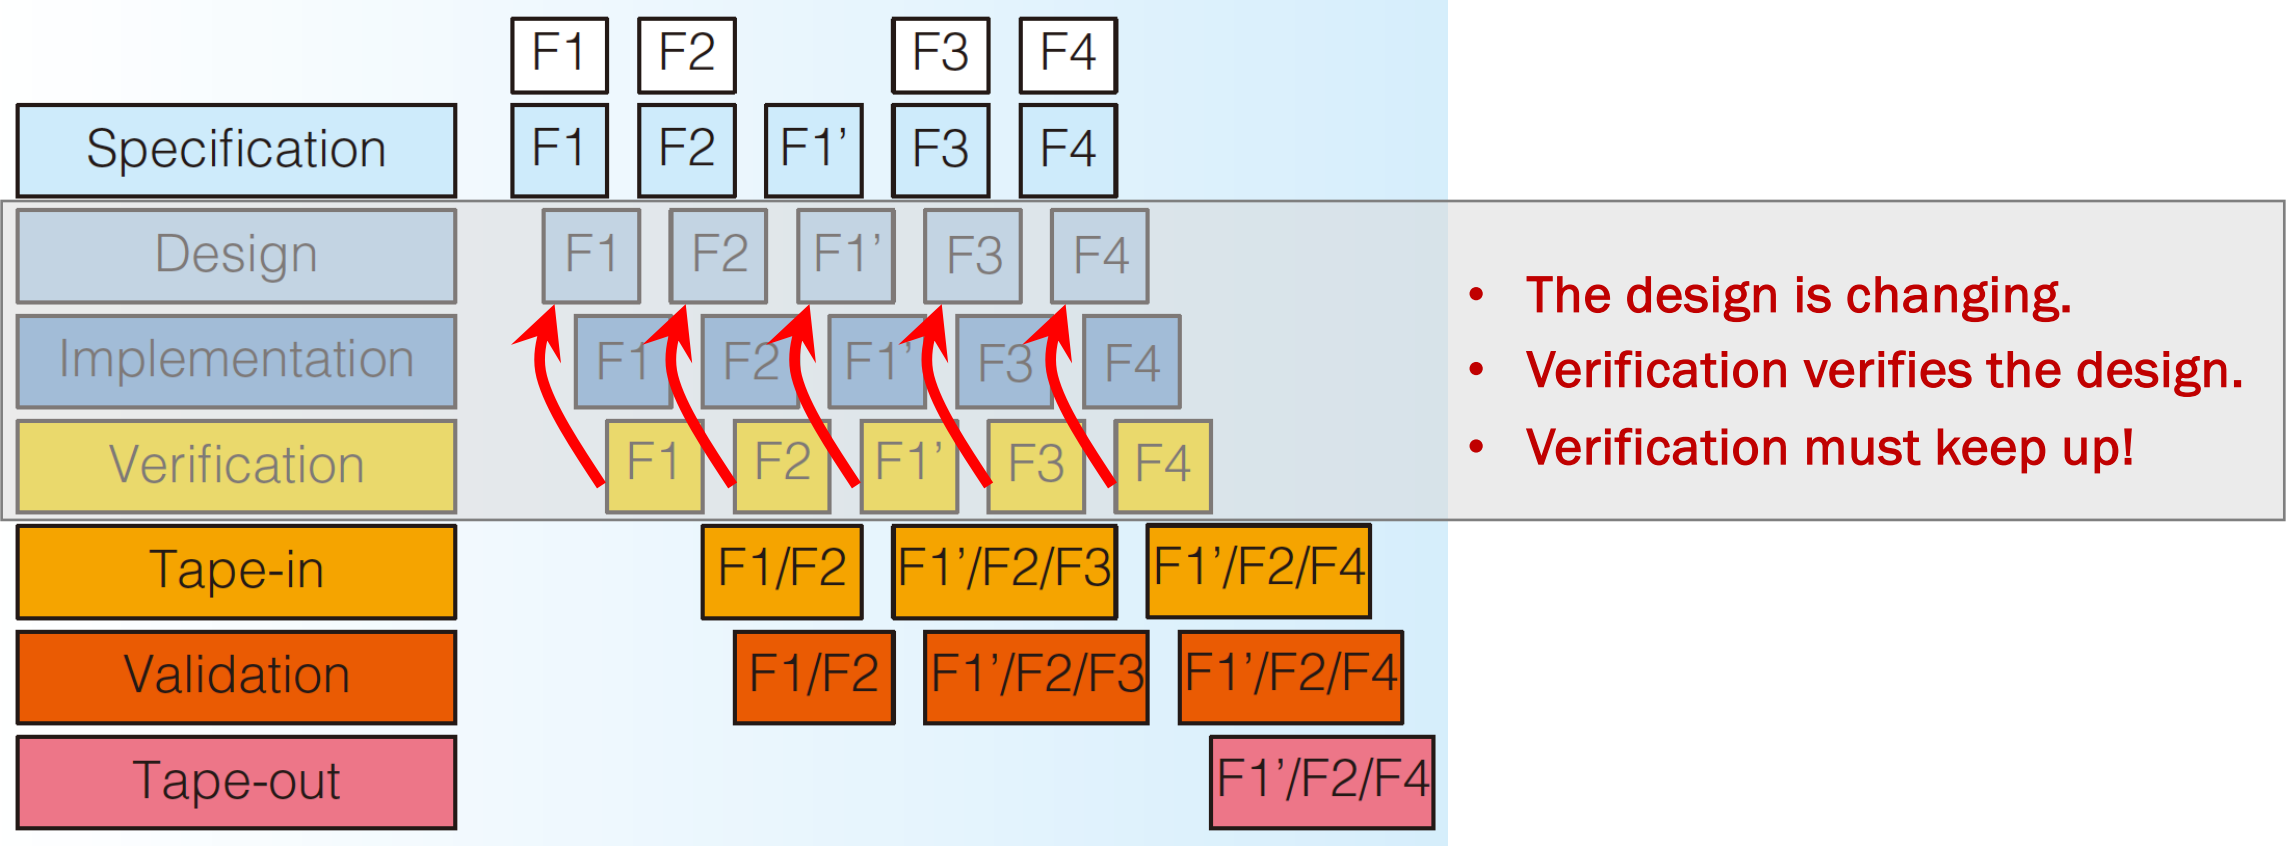
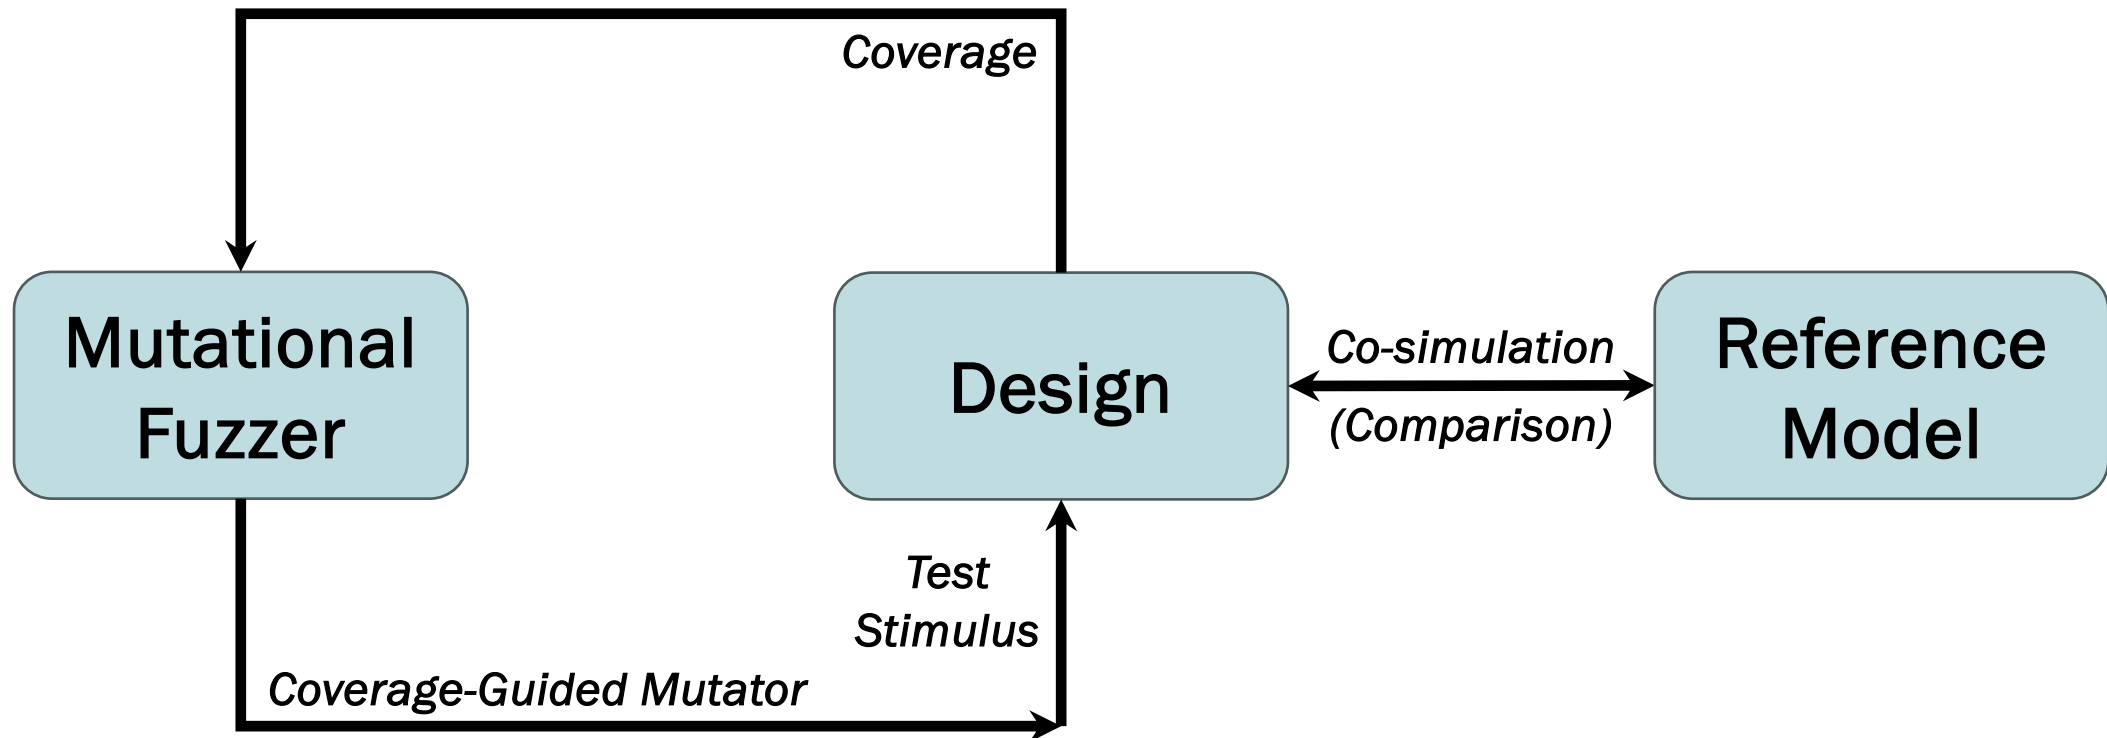


Figure. The agile model of hardware design

Source: Yunsup Lee et al., "An Agile Approach to Building RISC-V Microprocessors," in *IEEE Micro*, vol. 36, no. 2, pp. 8-20, Mar.-Apr. 2016.

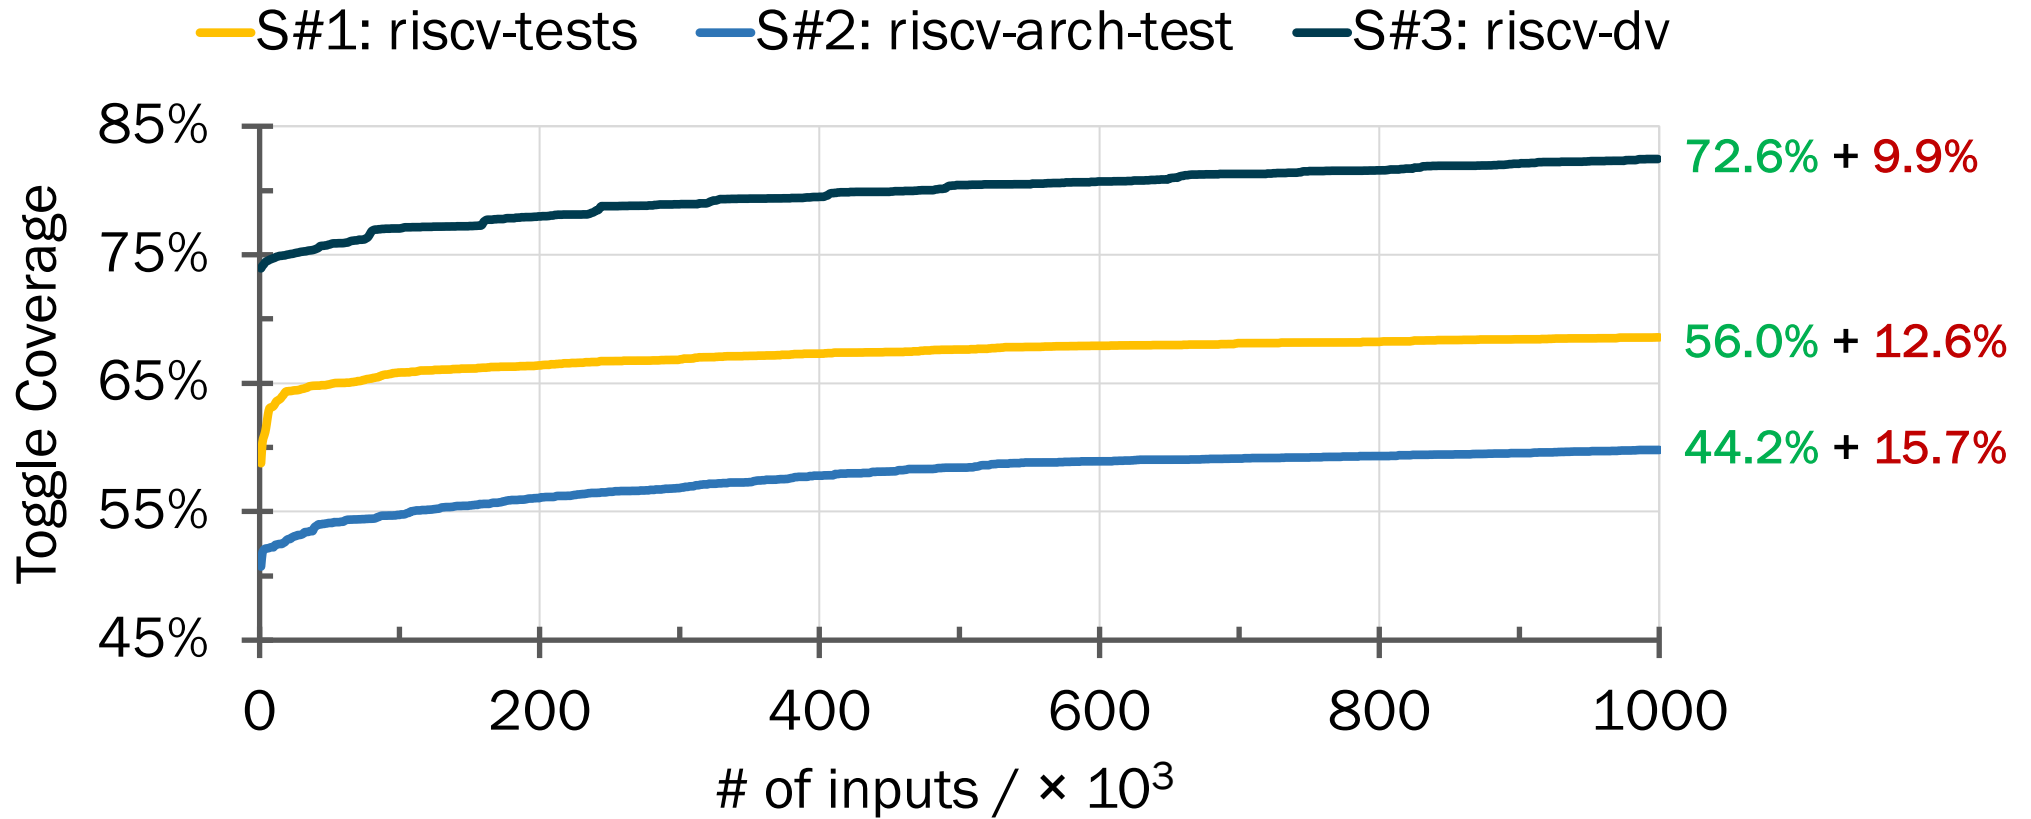
# Fuzzing: Automated Design-Directed Verification



# Fuzzing: Hardware vs. Software

- **Software fuzzing has been widely accepted and adopted**
  - Highly automated and efficient; significant return on investment (ROI)
  - [AFLplusplus/AFLplusplus] 4 steps: instrumenting, preparing, fuzzing, managing
  - [google/oss-fuzz] As of August 2023, OSS-Fuzz has helped identify and fix over 10,000 vulnerabilities and 36,000 bugs across 1,000 projects.
- **Hardware fuzzing is *more challenging* due to various issues**
  - Sophisticated binary-level input/output semantic
  - High design complexity; similar to highly concurrent programs
  - Low simulation speed/throughput
  - Lack of open-source practice (designs, corpus, crash detection, ...)

# Fuzzing CPUs: Coverage Increase



***Fact: the fuzzer achieves limited coverage increase from the start points***

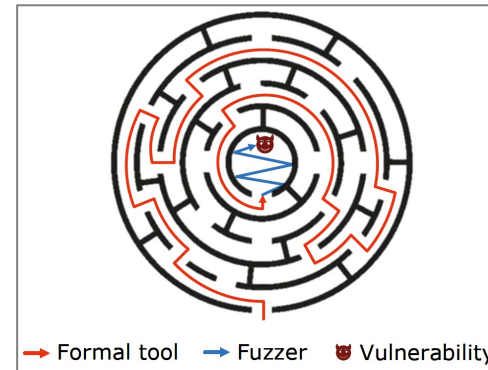
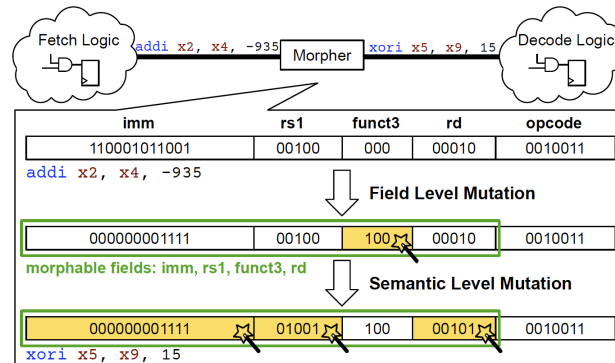
Note: data collected by the LibAFL fuzzer with havoc mutator and toggle coverage feedback from rocket-chip.



# Insight<sup>1</sup>: Exploitation and Exploration

- Literally, mutational fuzzers are *very good at exploitation*
  - Mutators generally create a large number of input cases
- However, *the exploration is inefficient* when applying fuzzing to CPUs
  - Reason: instruction set architectures (ISAs) are complicated and sophisticated
  - Recent works improve it with domain-specific knowledge or formal methods

[1] MorFuzz proposes RISC-V specific instruction mutations

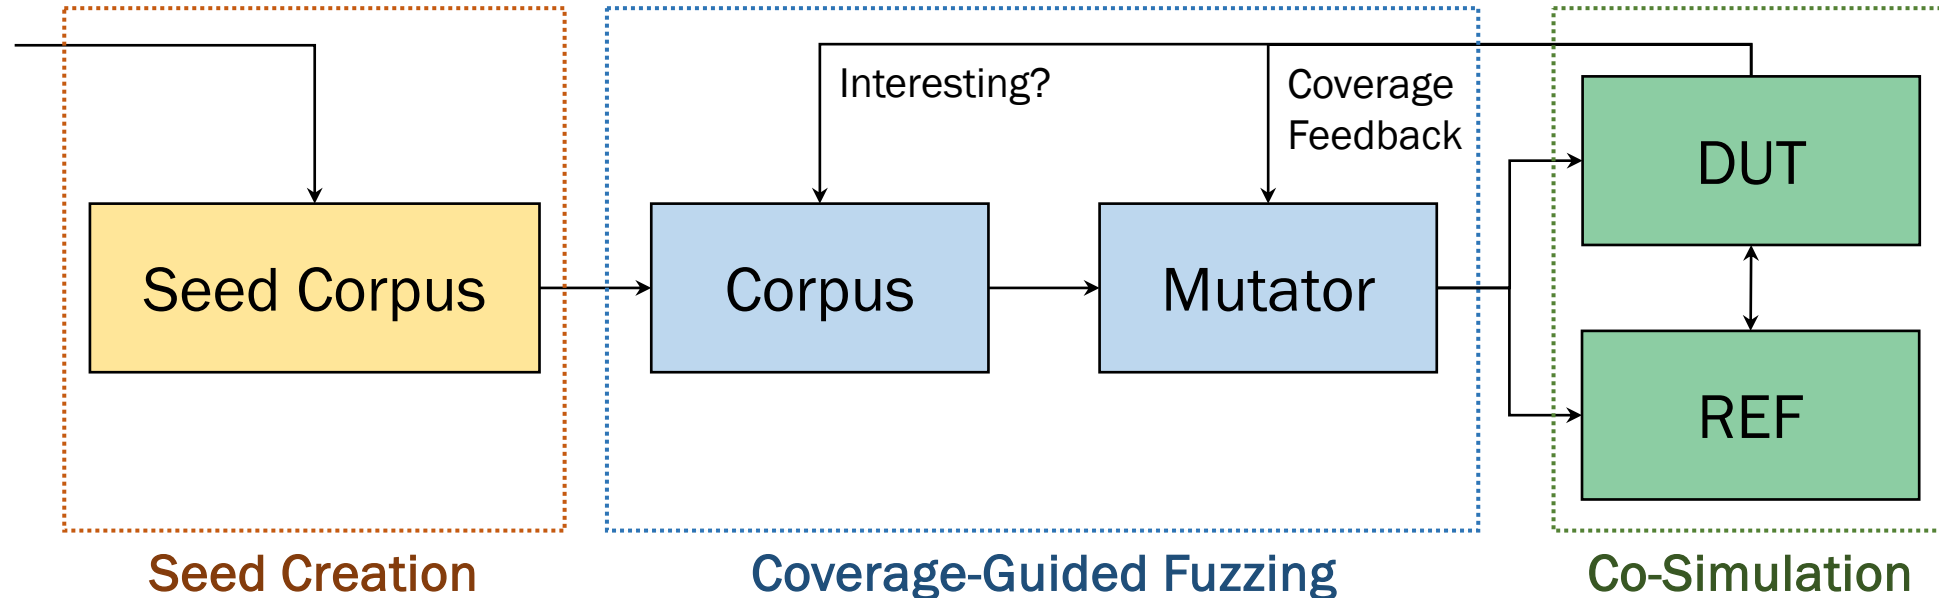


[2] HyPFuzz proposes formal-assisted state exploration

[1] Jinyan Xu, Yiyuan Liu, Sirui He, Haoran Lin, Yajin Zhou, and Cong Wang. 2023. MorFuzz: fuzzing processor via runtime instruction morphing enhanced synchronizable co-simulation. In *Proceedings of the 32nd USENIX Conference on Security Symposium (SEC '23)*. USENIX Association, USA, Article 74, 1307–1324.

[2] Chen Chen, Rahul Kande, Nathan Nguyen, Flemming Andersen, Aakash Tyagi, Ahmad-Reza Sadeghi, and Jeyavijayan Rajendran. 2023. HyPFuzz: formal-assisted processor fuzzing. In *Proceedings of the 32nd USENIX Conference on Security Symposium (SEC '23)*. USENIX Association, USA, Article 77, 1361–1378.

# Fuzzing CPUs: SOTAs have done good jobs



Seed Creation

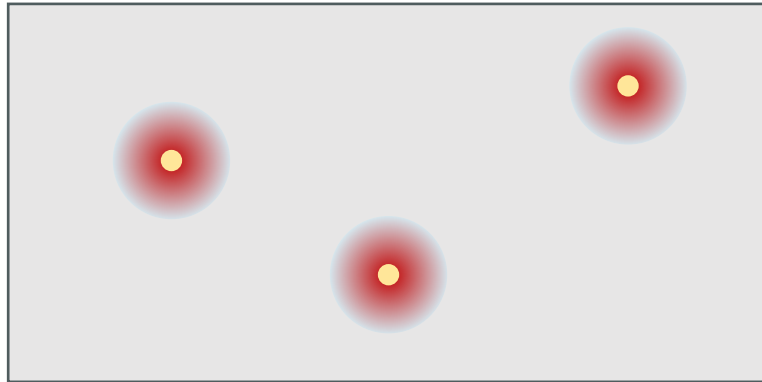
Coverage-Guided Fuzzing

Co-Simulation


- [SP'21] DifuzzRTL
- [USENIX Security'22] TheHuzz
- [GLSVLSI'22] CFG for Processor
- [USENIX Security'23] HyPFuzz
- [USENIX Security'23] MorFuzz
- [DATE'23] SoCFuzzer

- [MICRO'21] Dromajo
- [MICRO'22] DiffTest

# Insight<sup>2</sup>: Ways to Exploring the State Space



Fuzzing Basics



Wider Search Scope

*Better mutations*



More Start Points


*Richer seeds*

- Start point
- Search range

# Observations: Fuzzing Horizons are Constrained

## (1) *Mutations* are not that effective

Given the seeds S#3 (riscv-dv):



**Effective only if targeting 6.2% of input bytes**

Further decreased to 2.5% after 1M mutations

## (2) *Sources of seeds* are limited

Given the seeds S#4 (force-riscv):

```
0000000080000000 <text0>: # base
.....
0000000080000100 <text1>: # base+0.25KB
.....
0000000080011000 <text2>: # base+68KB
.....
000000008477fff8 <text3>: # base+71.5MB
.....
00000000a6411d80 <text77>: # base+612MB
.....
```

**OOM crashes for in-memory fuzzers**


Significant slow down fuzzers with corpus on disk

# Why: How CPU Fetches and Executes

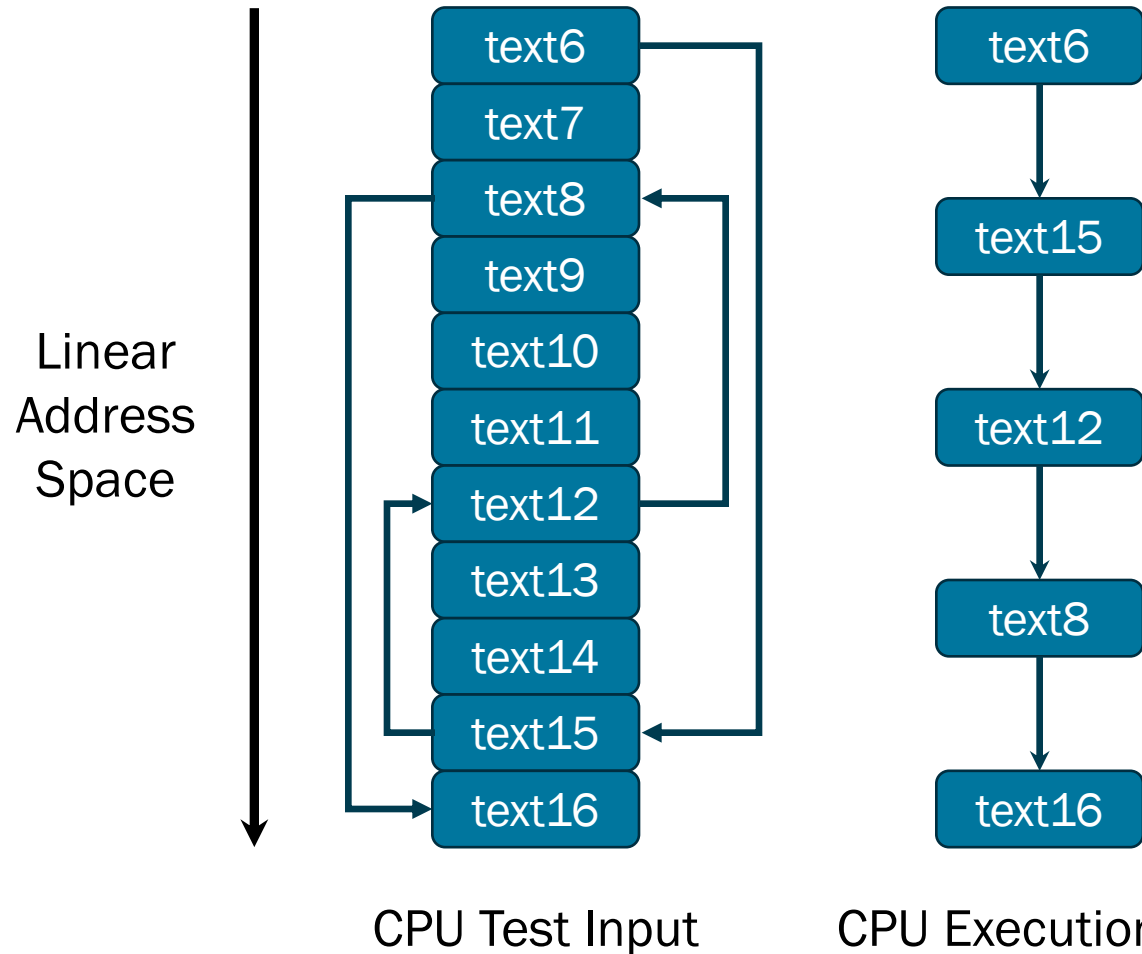
```
00000000880ca508 <text6>:
.....
880ca58c: 1d03d6ef jal a3,0x8810775c
.....
0000000088106b20 <text8>:
.....
88106b28: 62da92e3 bne s5,a3,0x8810794c
.....
0000000088107068 <text12>:
.....
88107088: a8de7ce3 bgeu t3,a3,0x88106b20
.....
0000000088107758 <text15>:
.....
88107888: fec65263 bge a2,a2,0x8810706c
.....
0000000088107948 <text16>:
.....
```

This is a case from the seeds S#4 (force-riscv)

# Input Format: The Linear Memory



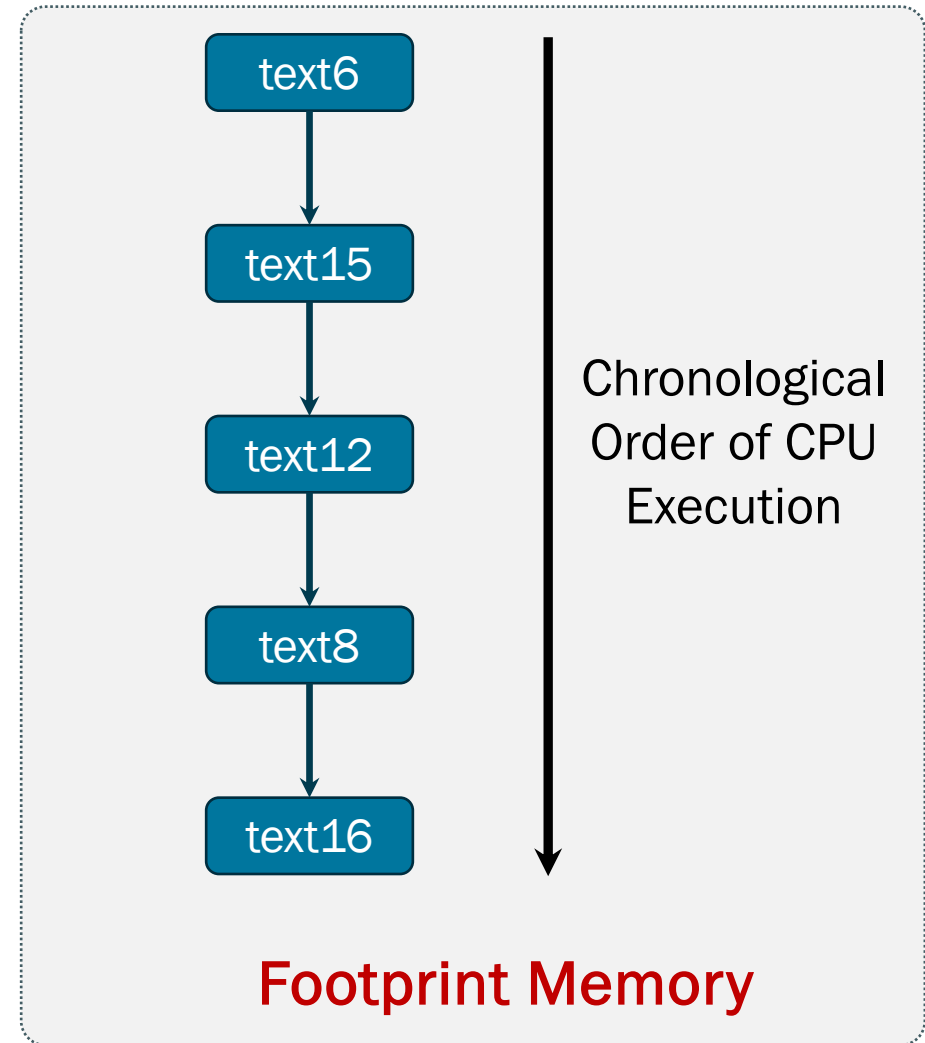

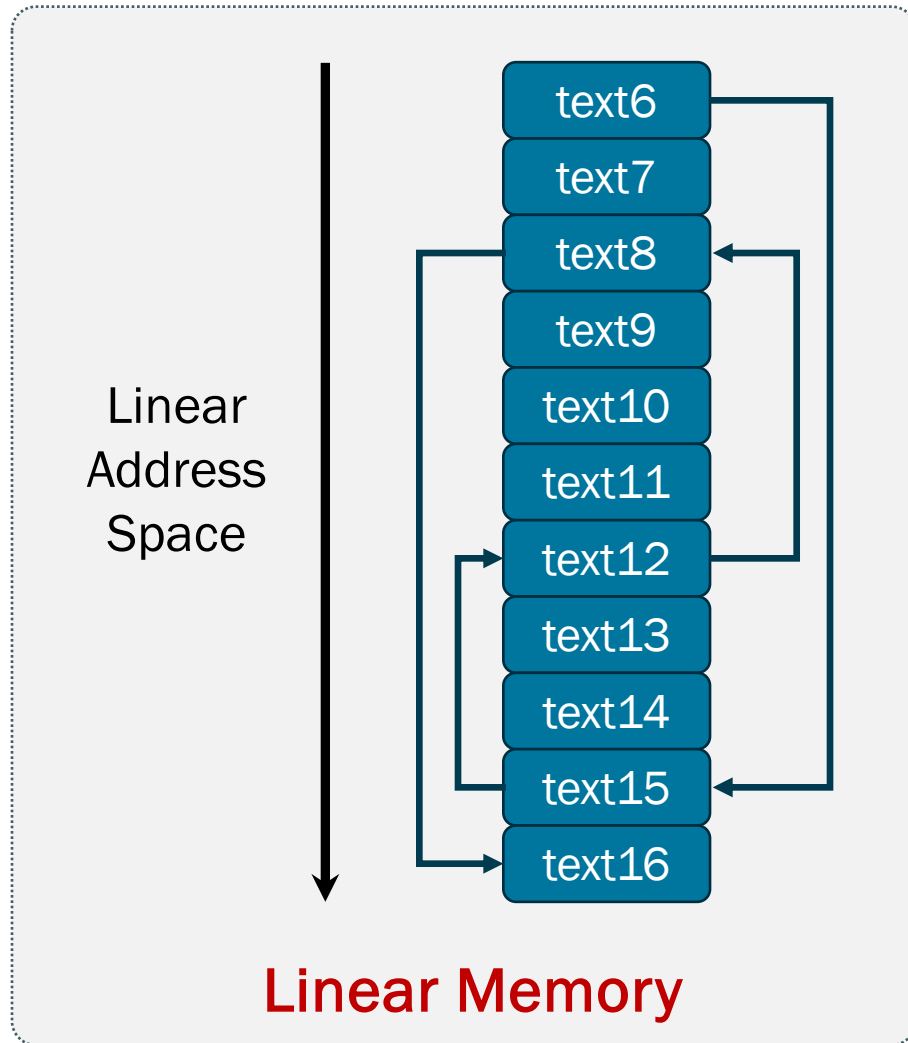
# Insight<sup>3</sup>: Linear Memory Hides Execution Paths



*If removing untouched memory contents ...*

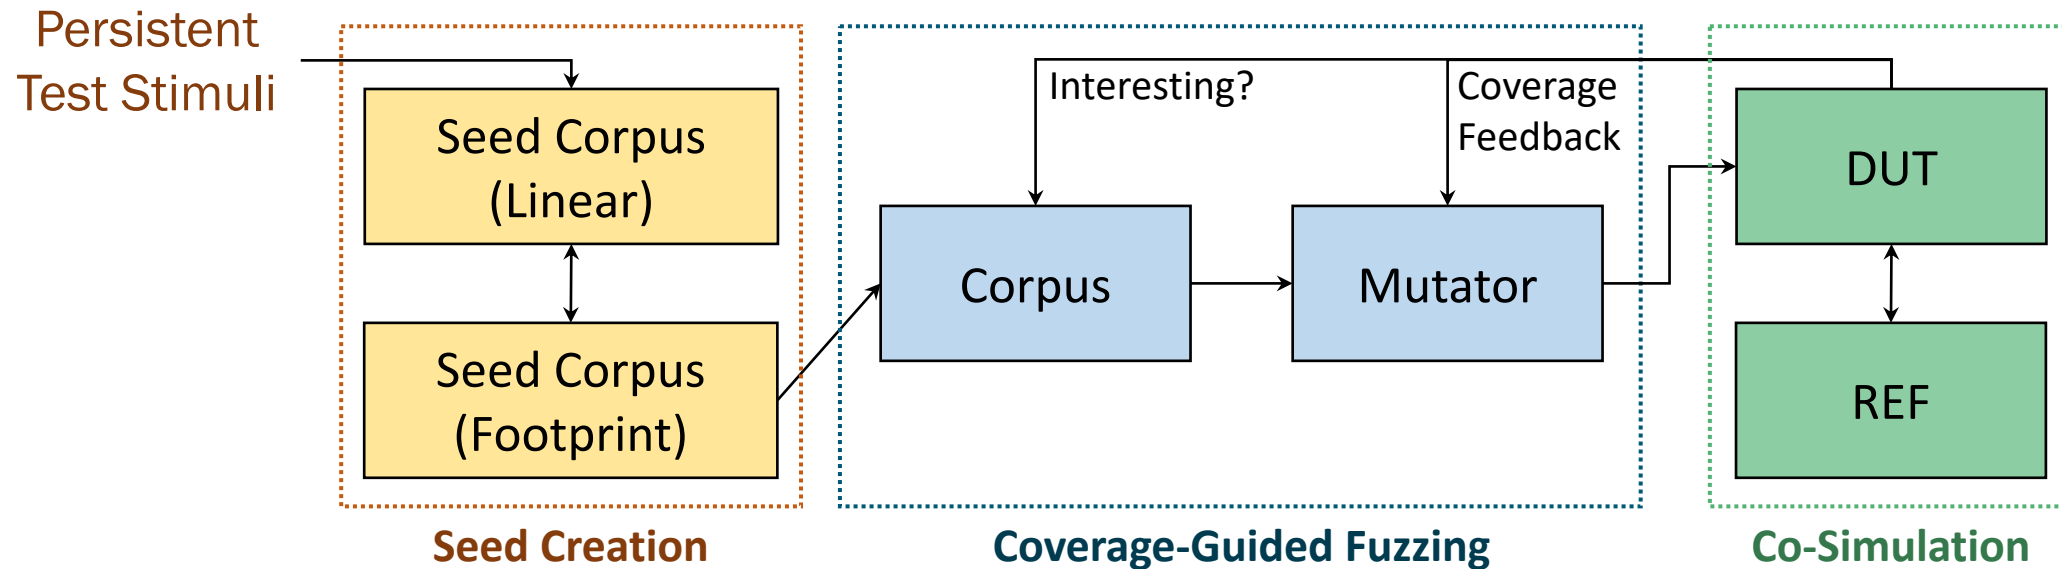
- *Mutations become more effective*
- *Seeds' size is significantly reduced*

# Footprint Memory: Capturing Execution Paths





# PathFuzz: Overview of the Workflow



Refer to our Paper Section 3.2 for more details in enhancing/adapting the three stages.

# PathFuzz: Broadening Sources of Seed Corpus

- Modern CPU DV reaches a good coverage; let fuzzers take a step further
- The test cases we are currently using for the system-level DV of CPUs

## 1) hand-written directed tests

- riscv-software-src/riscv-tests
- riscv-non-isa/riscv-arch-test
- riscv-ovpsim/imperas-riscv-tests
- litmus-tests/litmus-tests-riscv
- josecm/riscv-hyp-tests

## 2) instruction-stream generators

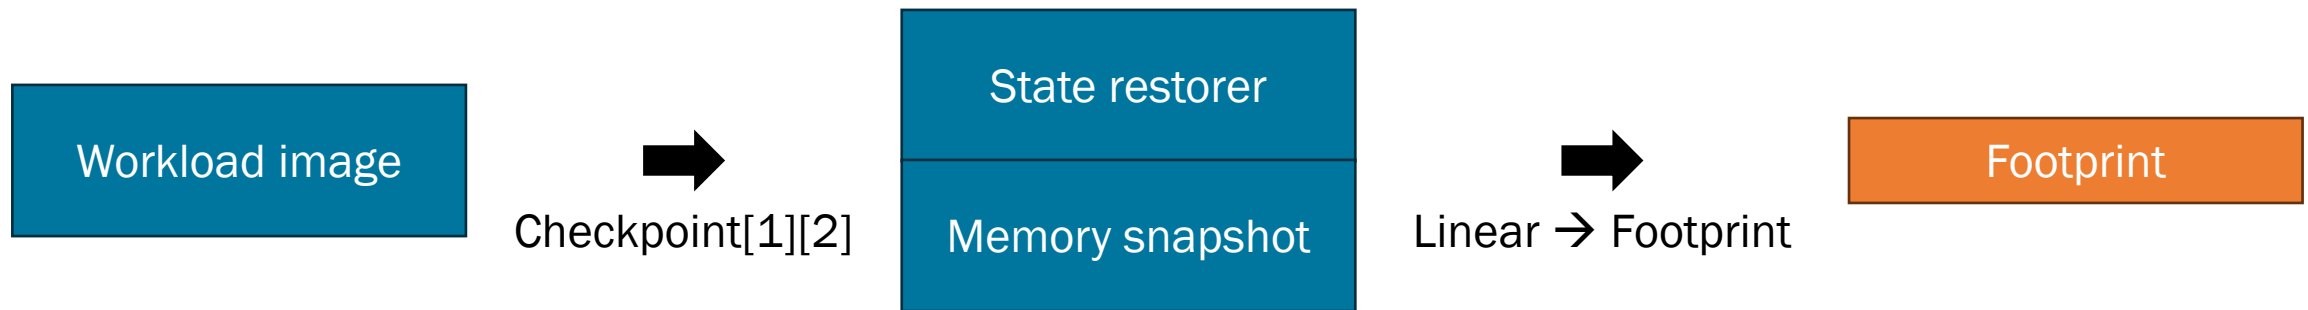
- chipsalliance/riscv-dv
- openhwgroup/force-riscv
- ksco/riscv-vector-tests
- sifive/riscv-vector-intrinsic-fuzzing
- chad-q/andes-vector-riscv-dv

## 3) real-world programs

- ucb-bar/riscv-benchmarks
- eembc/coremark
- SPEC CPU® 2017
- SPECjbb® 2015
- gcc,clang,rustc,verilator

# PathFuzz: Enhancing DV with Practical Fuzzing

- Incorporating existing, valuable CPU test cases as seeds for fuzzing
  - Extracting the footprints when CPU executes these test cases
  - Using the (short-running) footprints as fuzzing seeds
- Contribution: fuzzing with any start point at any program phase
- How: architectural checkpoints + footprint memory



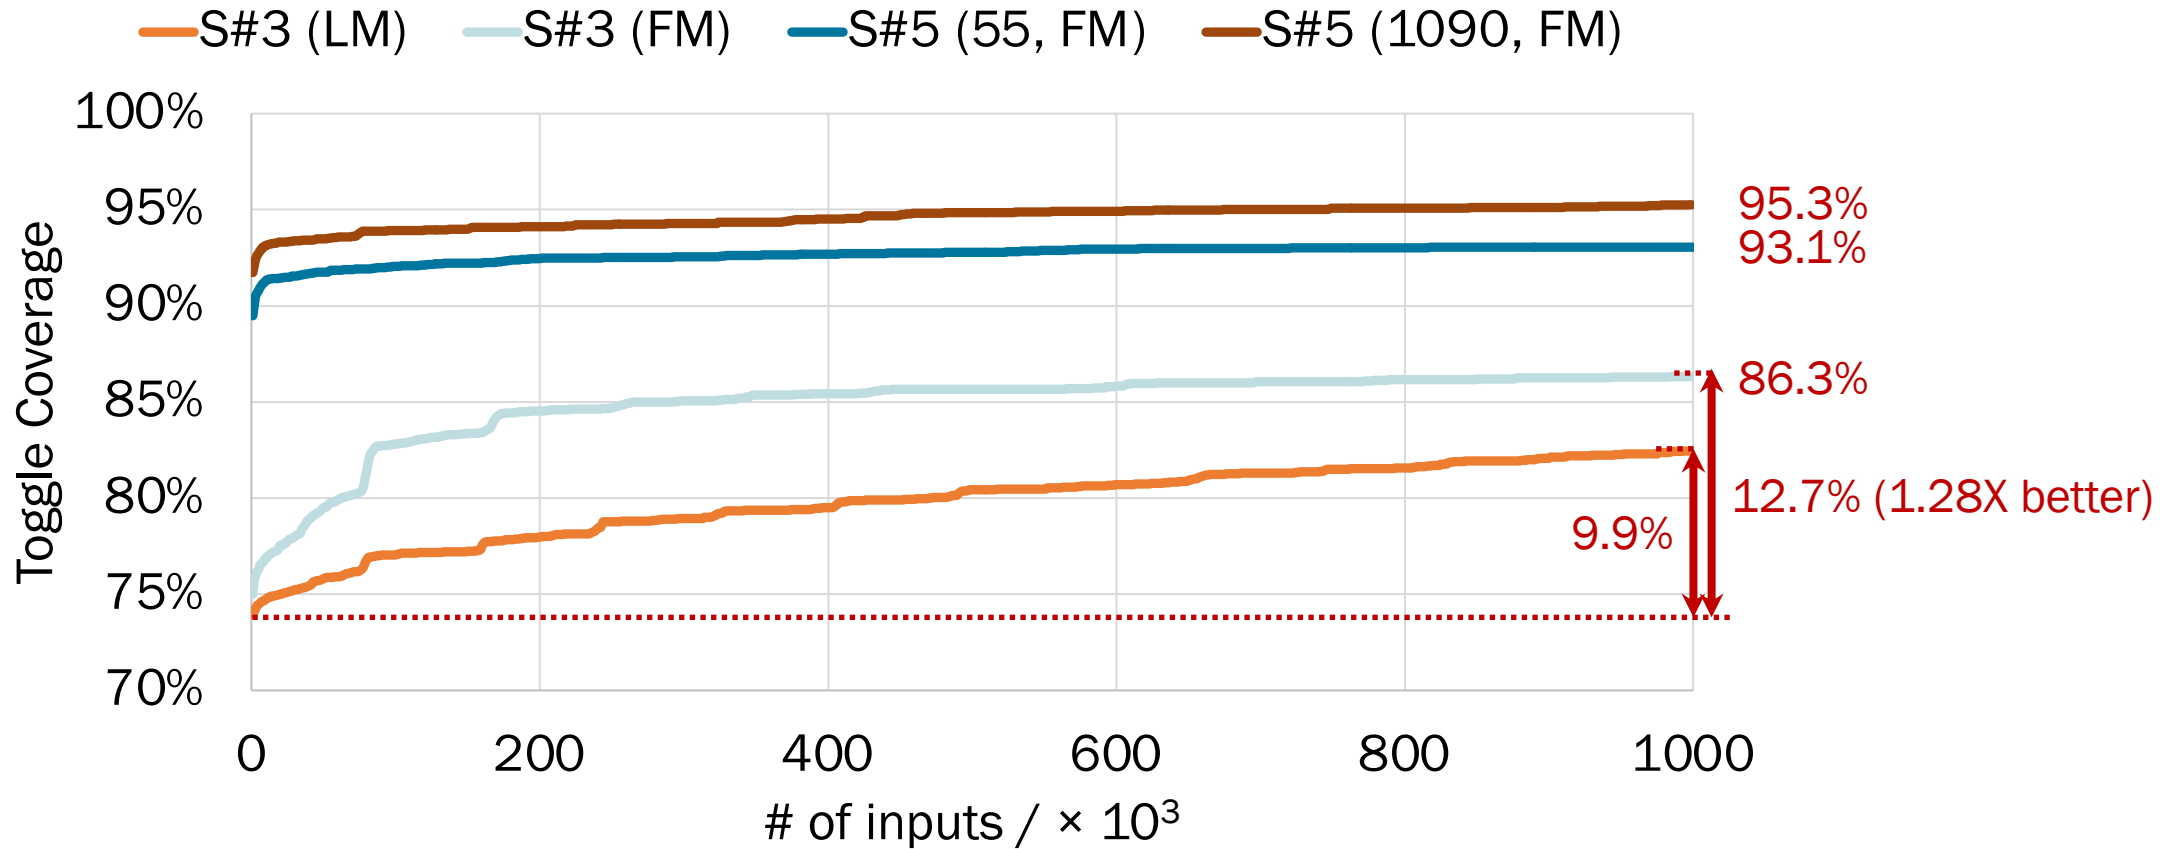
[1] Nursultan Kabylkas, et al., 2021. Effective Processor Verification with Logic Fuzzer Enhanced Co-simulation. *MICRO'21*.

[2] Yinan Xu, et al., 2023. Towards Developing High Performance RISC-V Processors Using Agile Methodology. *MICRO'22*.

# Evaluation

- Setup: famous, widely-adopted, open-source projects
  - Fuzzer: LibAFL v0.10.1 (unmodified QueueScheduler, StdMapObserver, StdFuzzer)
  - CPU design under test: rocket-chip
  - CPU reference/golden model: Spike (riscv-isa-sim)
  - Various seeds, seed count (linear/footprint formats)
    - S#1: riscv-tests, 140 (LM, FM)
    - S#2: riscv-arch-test, 257 (LM, FM)
    - S#3: riscv-dv, 1150 (LM, FM)
    - S#4: force-riscv, 969 (FM)
    - S#5: SPEC CPU2006, 1090 (FM)
- To show the coverage increase, coverage reach, discovered bugs

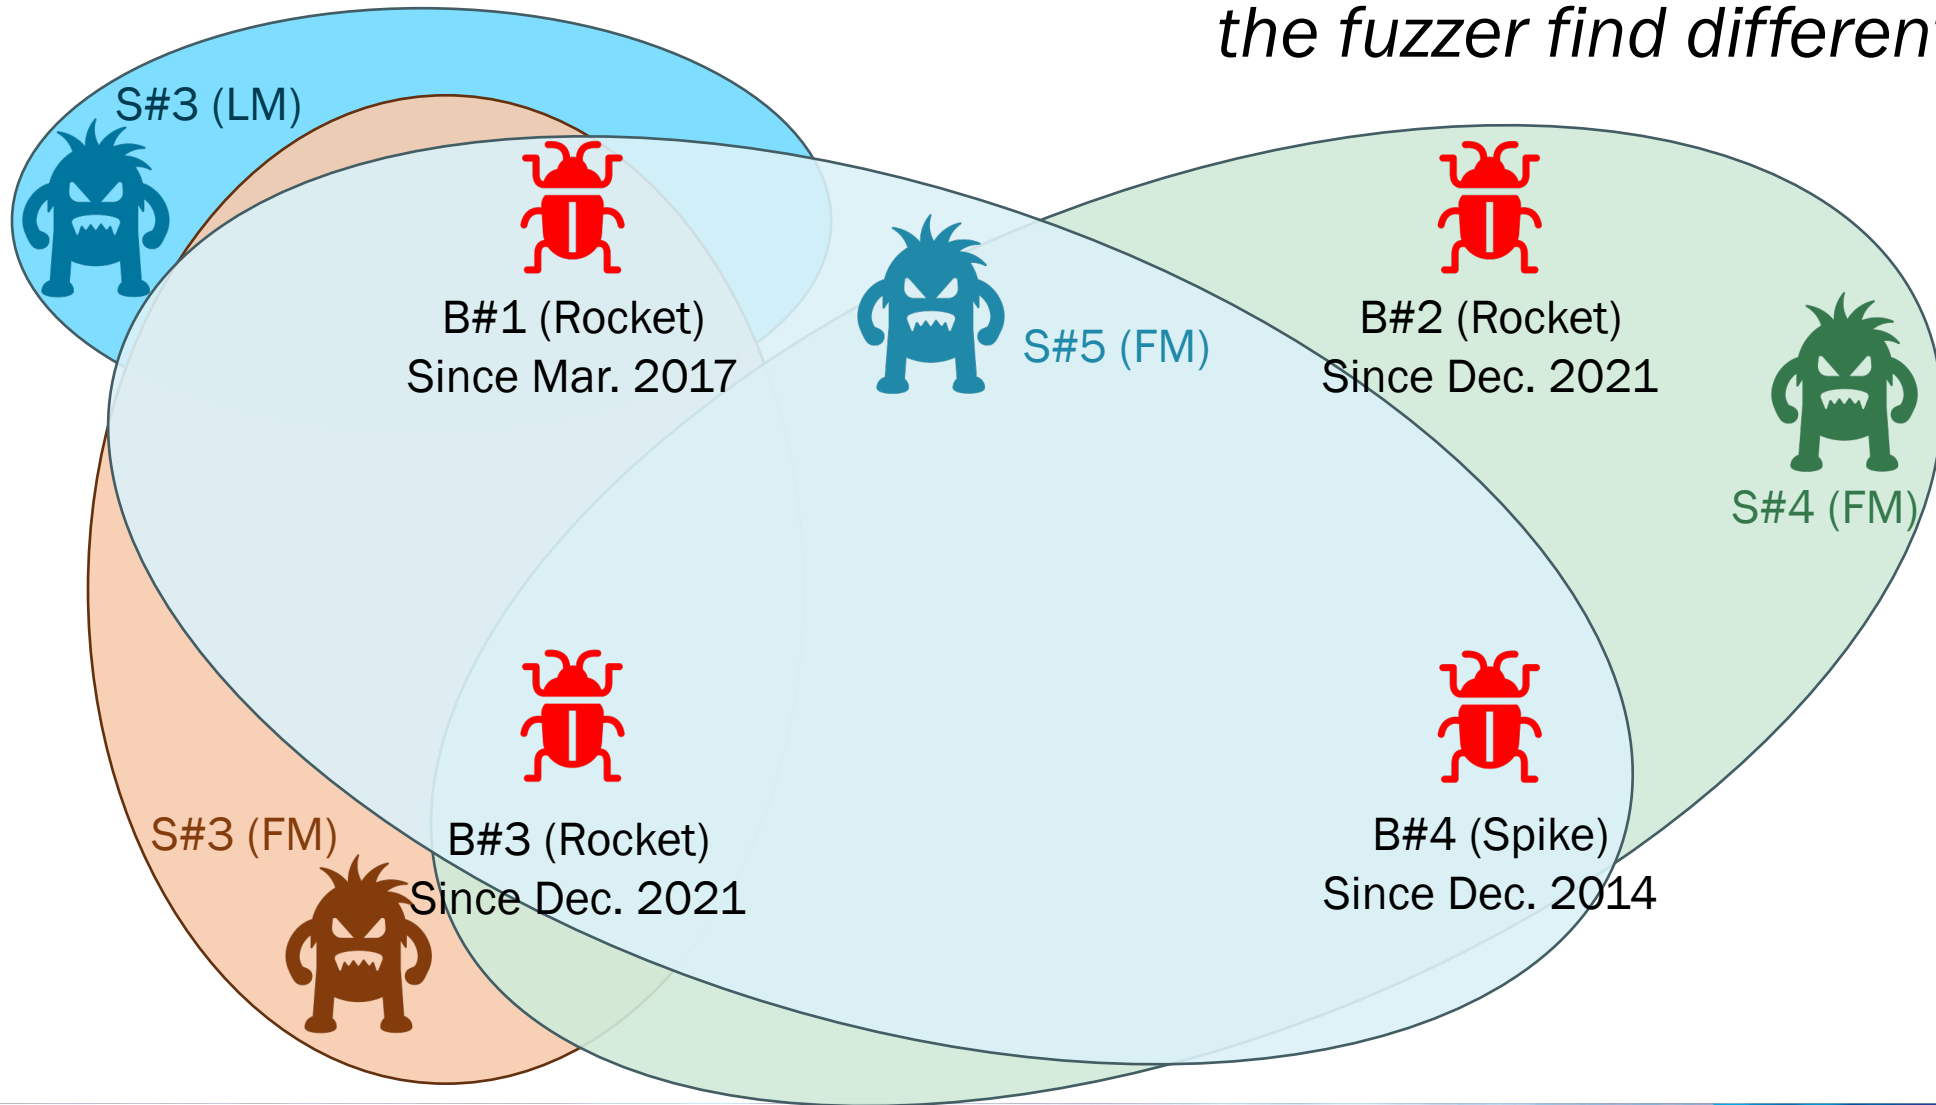
# Evaluation: Coverage



\* 95% criteria: formal-assisted HyPFuzz takes 72 hours to achieve 94.9% on CVA6; we take ~10 hours to achieve 95.3% on rocket-chip

# Evaluation: Bugs

Insight<sup>4</sup>: *Different seeds help the fuzzer find different bugs!*



# Applying Fuzzing to Open-Source XiangShan

欢迎 ChinaSys 社区研究者多多关注开源硬件验证（测试）领域

Version	#Error / #All (seed corpus)	Potential Bug Count
20230905	*** / 50000	5
20230907	**** / 300000	6+ (***/**** analyzed)
20230915	**/1838 (riscv-tests, LM)	Not analyzed yet
	**/3772 (riscv-arch-test, LM)	
	*** / 2181 (riscv-dv, LM)	
	*** / 25087 (riscv-tests, FM)	
	**/4132 (riscv-arch-test, FM)	
	*** / 2532 (riscv-dv, FM)	
	*** / 2196 (force-riscv, FM)	
	*** / 3751 (SPECCPU2006, FM)	

\* Preliminary testing results on unstable versions of XiangShan; do not necessarily reflect the final design verification quality.

# Conclusion

- **Motivation:** broadening the fuzzing horizons on CPUs
  - More effective mutations, richer seed sources for better exploration capabilities
- **PathFuzz:** a coverage-guided CPU fuzzing workflow
  - Input format: both linear and footprint memory
  - Incorporate large-scale programs as fuzzing seeds
- **Evaluation**
  - Achieve better coverage increase/reach
  - Detect 4 long-standing bugs in well-known projects
- **Open-sourced** at GitHub with open-source components\*
  - Contribute to the reproducible, reusable research community

\* <https://github.com/OpenXiangShan/xfuzz>. Thank LibAFL, rfuzz, DifuzzRTL, SIC, and DiffTest.



# Conclusion; Questions?

- **Motivation:** broadening the fuzzing horizons on CPUs
  - More effective mutations, richer seed sources for better exploration capabilities
- **PathFuzz:** a coverage-guided CPU fuzzing workflow
  - Input format: both linear and footprint memory
  - Incorporate large-scale programs as fuzzing seeds
- **Evaluation**
  - Achieve better coverage increase/reach
  - Detect 4 long-standing bugs in well-known projects
- **Open-sourced** at GitHub with open-source components\*
  - Contribute to the reproducible, reusable research community

\* <https://github.com/OpenXiangShan/xfuzz>. Thank LibAFL, rfuzz, DifuzzRTL, SIC, and DiffTest.