

# No Linux, No Problem:

## Fast and Correct Windows Binary Fuzzing via Target-embedded Snapshotting

---

**Leo Stone**  
leo@vt.edu

**Rishi Ranjan**  
rishiranj@vt.edu

**Stefan Nagy**  
snagy@cs.utah.edu


**Matthew Hicks**  
mdhicks2@vt.edu




**COMPUTER SCIENCE**  
VIRGINIA TECH.

# Introduction

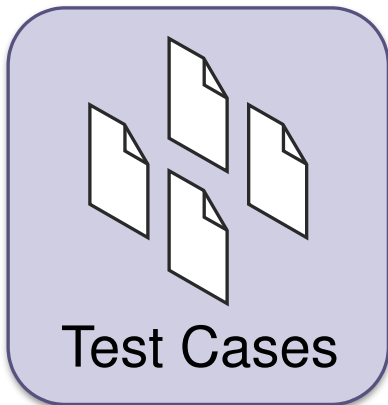
- **Fuzzing:** most successful method for automated software testing
- Attempt to break program using **many** randomized inputs



parse\_jpeg.exe



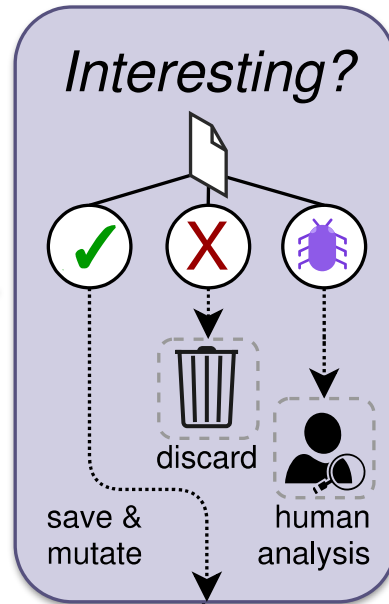
# 1 GENERATION



# 2 EXECUTION



# 3 TRIAGE




Limiting factor

Depends on speed: faster fuzzing = more effective fuzzing

# Motivation

- On Linux, execution is very fast → fuzzing is effective
  - `fork()`
  - Kernel modifications
- No equivalents on Windows...
  - Windows software ecosystem is larger, but fuzzing is orders of magnitude less effective than on Linux!
  
- **Challenge: without kernel support, how can we build an efficient Windows fuzzer?**

# One Solution: Persistent Mode




Fast, but **incorrect**

Execution changes state...

**Many binaries crash!**

# Our Approach



Copy state to a **snapshot**

Process restores own state

Fast like persistent mode, but correct!

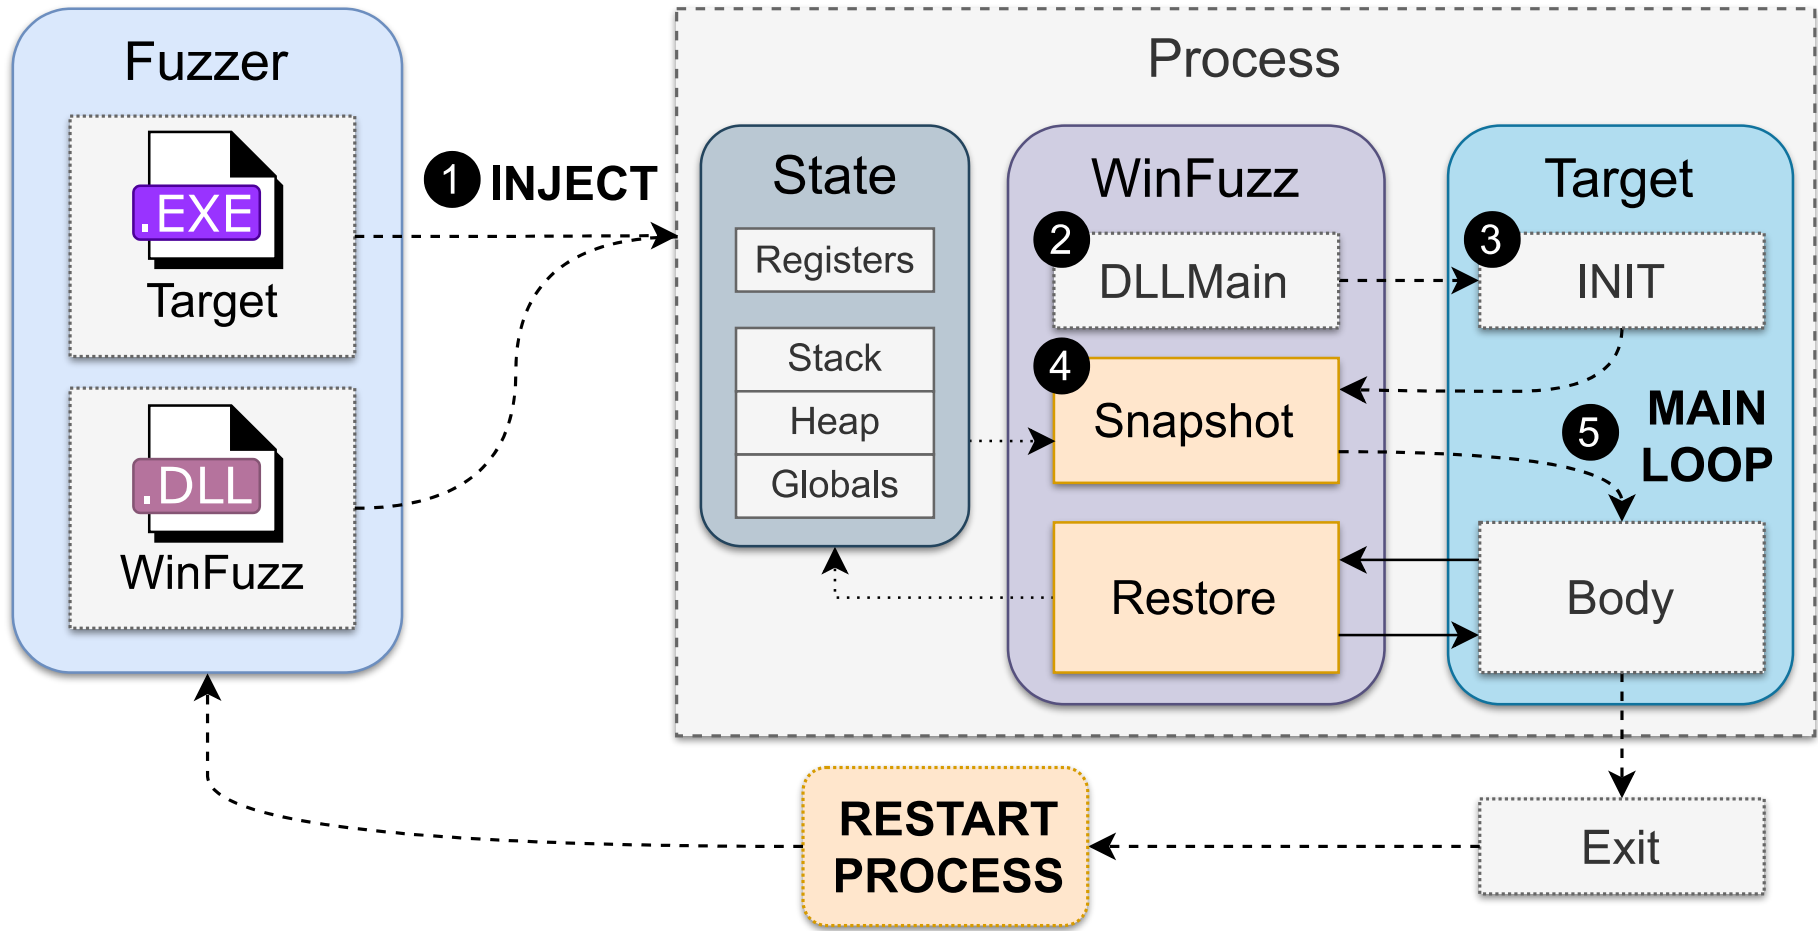
## Key Insight

All relevant state is controllable with language-level constructs

No kernel support necessary

# Implementation - WinFuzz

- Based on Winnie, an existing Windows fuzzer
- Main steps:
  1. DLL injection
  2. `DLLMain()` - hook setup
  3. Initializing target
  4. Taking state snapshot
  5. Main fuzzing loop: run and restore state






# Elements of Program State

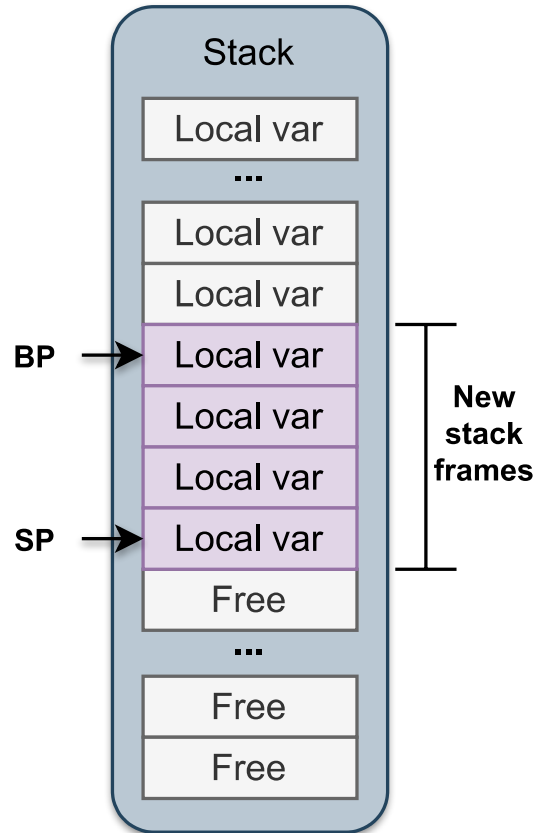
<b>Stack</b>
<b>Registers</b>
<b>Heap</b>
<b>Global variables</b>

### 1 Snapshot

savedBP = BP  
savedSP = SP

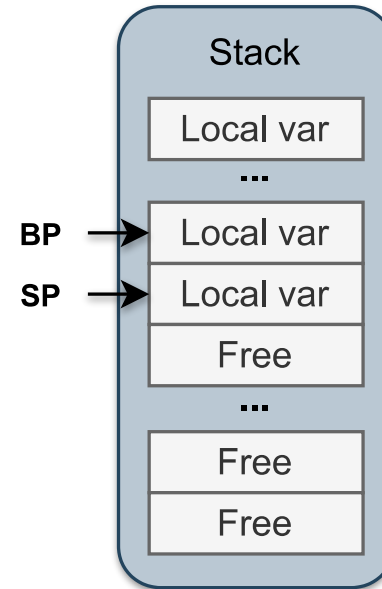


### 2 Target runs



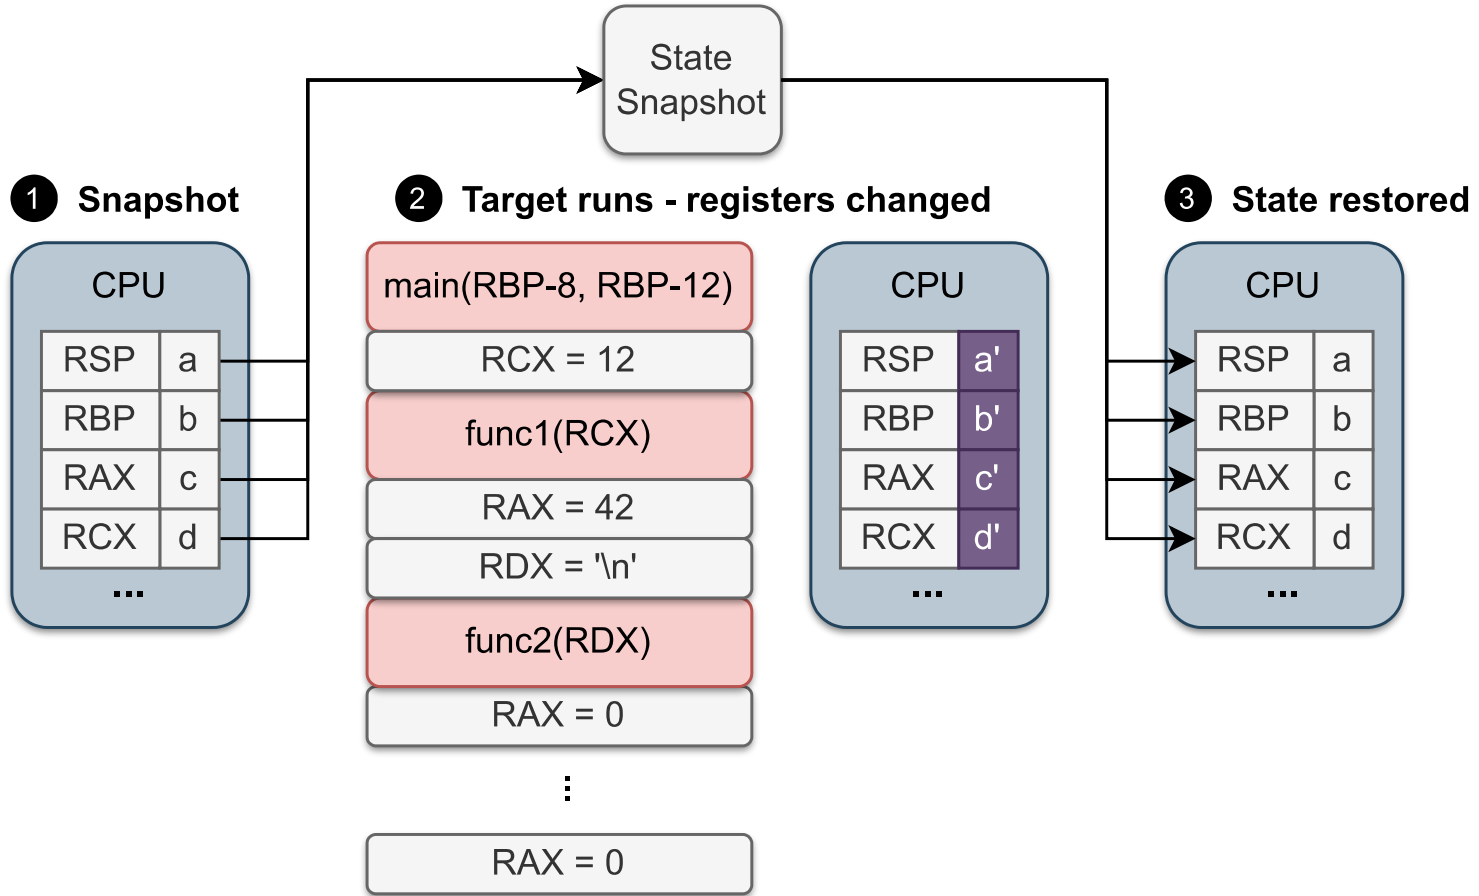
### 3 Stack restored

BP = savedBP  
SP = savedSP



# Resetting Registers

- 2 types of parameters:
  - Simple values
  - Pointers
- We reset each register to its original value
  
- We avoid returning from the target function and destroying any stack frames that could hold target parameters



# Resetting the Heap

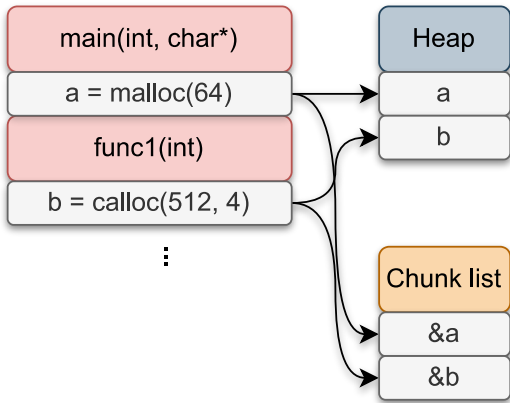
- Memory leaks are expected - we're trying to break the program!
- Small memory leak can cause a fuzzer crash
- We use heap API hooks to prevent memory leaks

### 1 Initialization

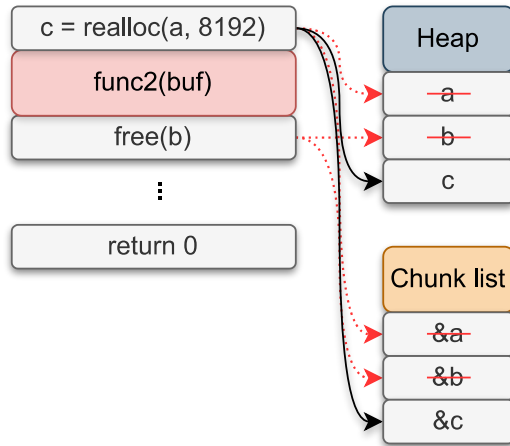


### 2 Target runs


a. Live pointers tracked



b. Pointers removed on free

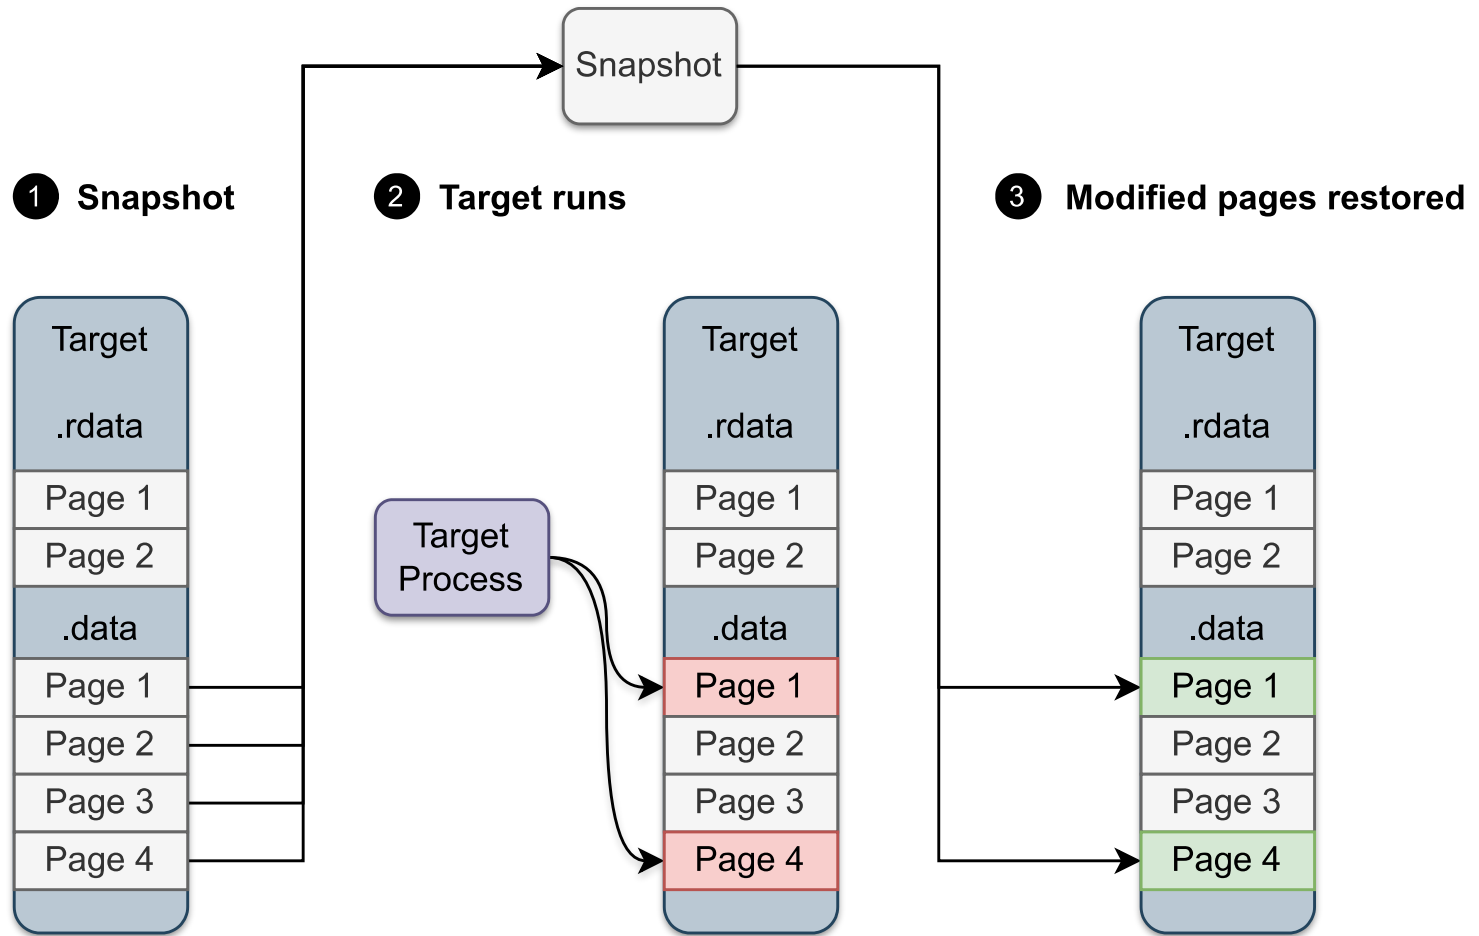


### 3 Remaining chunks freed



# Resetting Global Memory

- Copy correct starting state of all mutable global sections
- Use guard pages to track modifications
- Only restore *modified* pages





# Evaluation

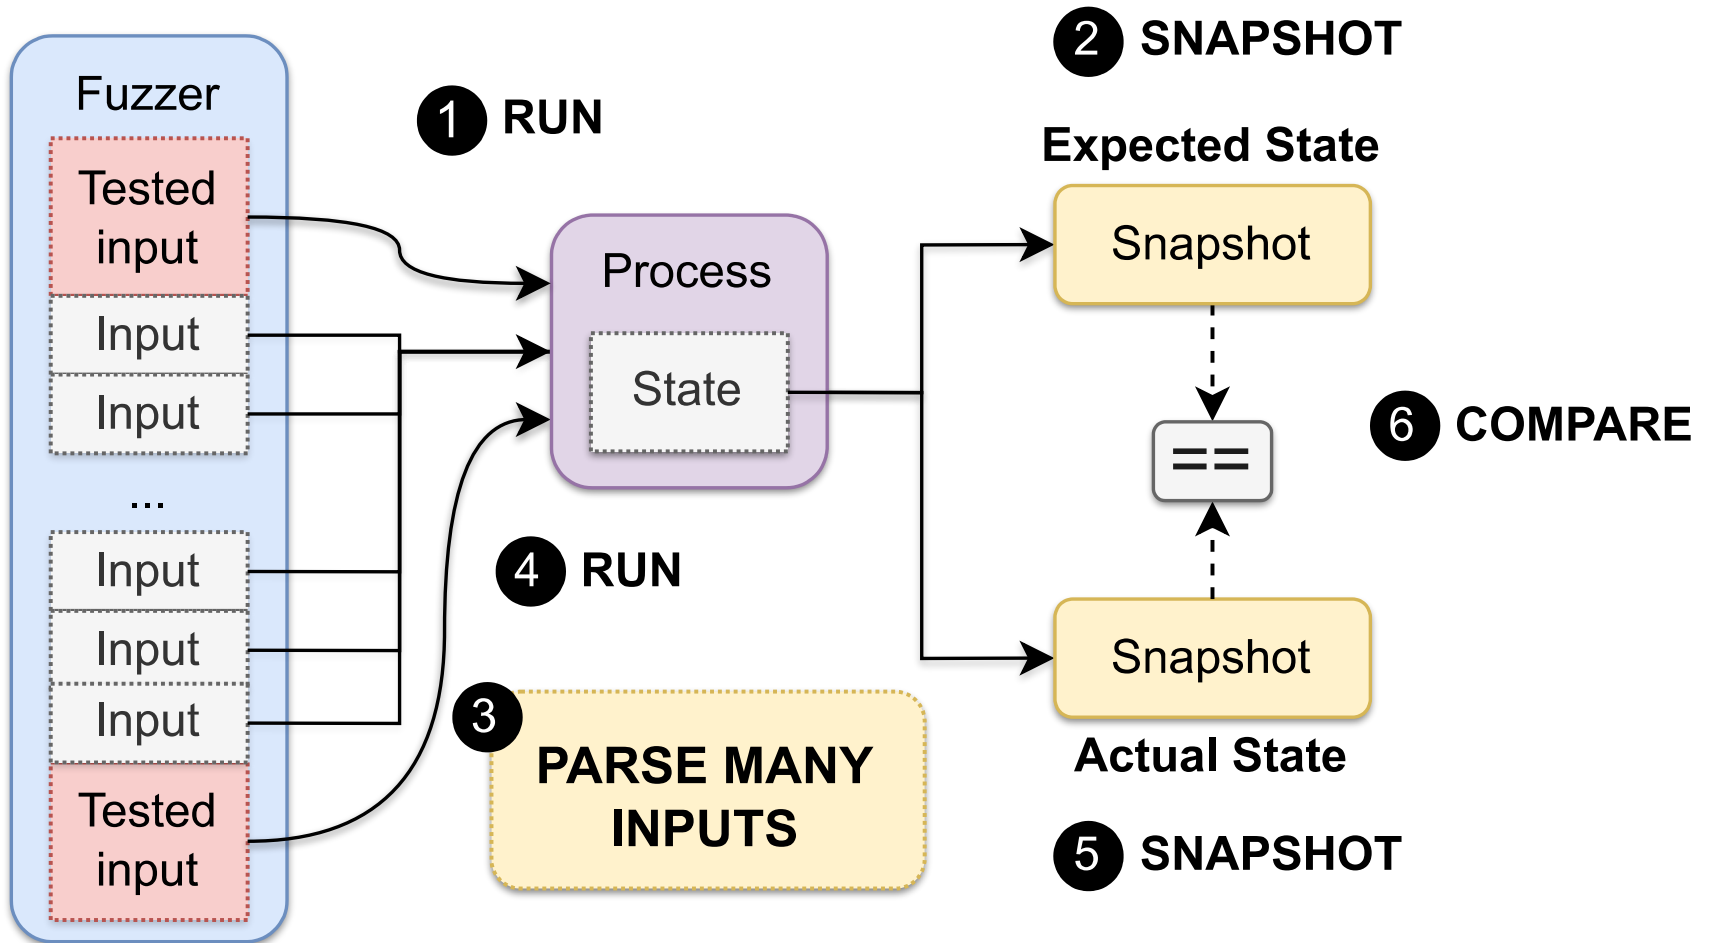
- Criteria: correctness, performance, bug discovery
  - Versus state-of-the art: Winnie (custom forkserver) and WinAFL (process creation/persistent mode)
- Setup: Azure instances running Windows 10 Pro with single-core 2.1 GHz Intel Xeon CPUs, 3.5 GB RAM
- Each fuzzing trial ran at least 5 times to collect statistically significant results
  - Mann-Whitney u-tests used to determine significance

# Benchmarks

Program	WINFUZZ	Winnie	WinAFL	Source	File Format	Size (KB)	Basic Blocks
<b>tar</b>	✓	✓	✗	Proprietary	.tar	606	30758
<b>nconvert</b>	✓	✓	✗	Proprietary	.png	2458	91550
<b>freetype</b>	✓	✓	✓	Open Source	.ttf	482	20891
<b>audiofile</b>	✓	✓	✓	Open Source	.wav	45	1504
<b>flac</b>	✓	✓	✓	Open Source	.flac	686	19292
<b>nanosvg</b>	✓	✓	✓	Open Source	.svg	47	1966
<b>sqlite</b>	✓	✓	✓	Open Source	.db	802	46758
<b>smpdf</b>	✓	No cov	✗	Proprietary	.pdf	3379	39816
<b>irfanview</b>	✓	No cov	✓	Proprietary	.png	1946	55187
<b>jq</b>	✓	No cov	✗	Open Source	.json	2662	13965


# Correctness Test

- New fuzzer mode that checks for state corruption by comparing program states
- Used to test all benchmarks in corpus
- Available as part of our open source implementation
  - New users can test their own targets and saved inputs

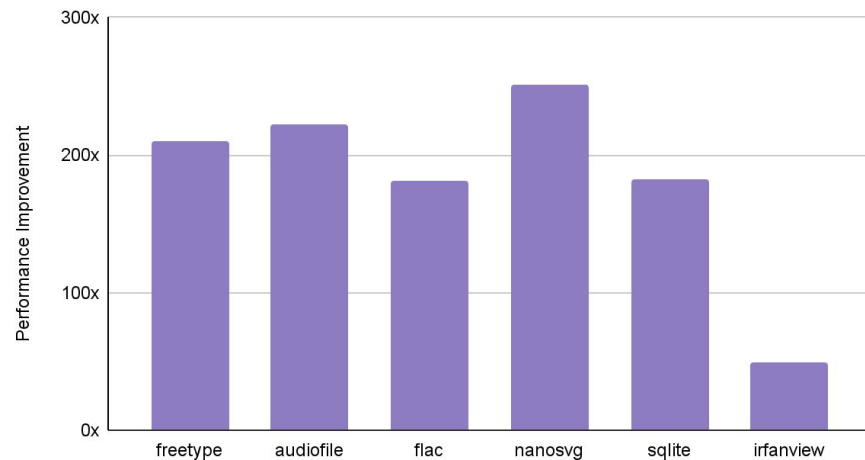


# Results: Throughput

Versus Winnie



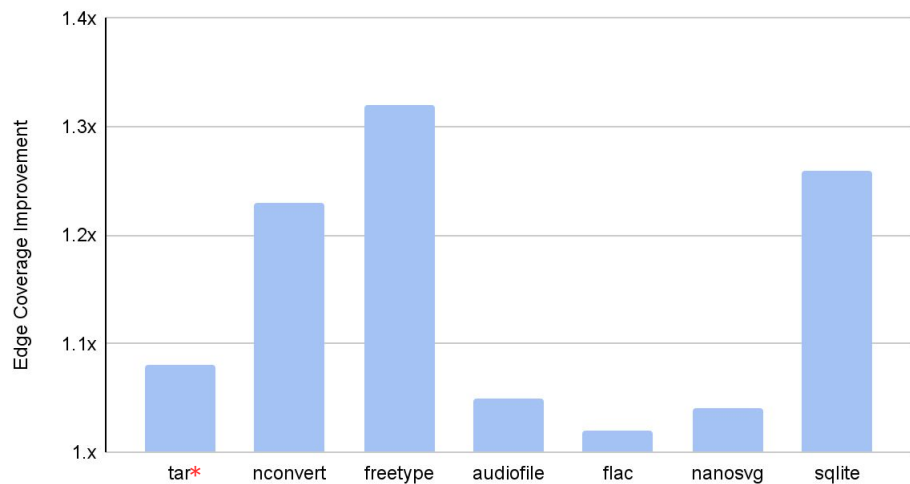
Versus WinAFL



Average improvement: 7x vs. Winnie and 182x vs. WinAFL

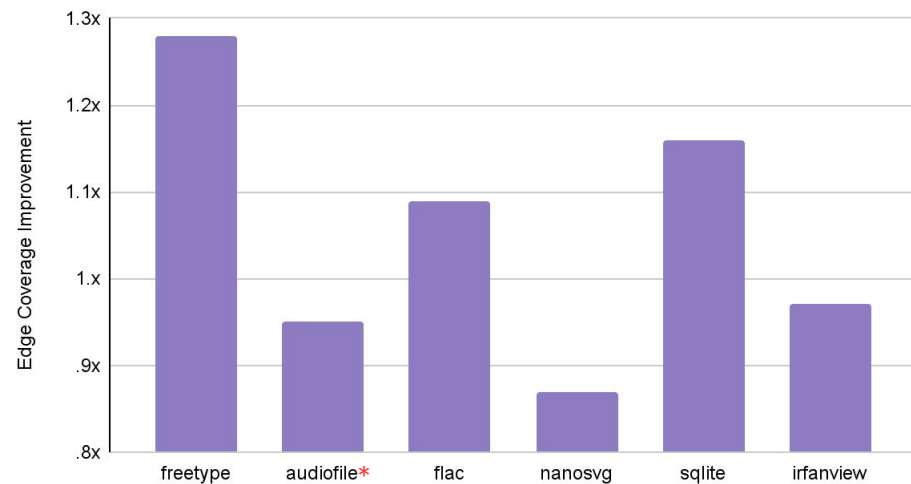
# Results: Edge Coverage

Versus Winnie



\* $p=0.148$

Versus WinAFL

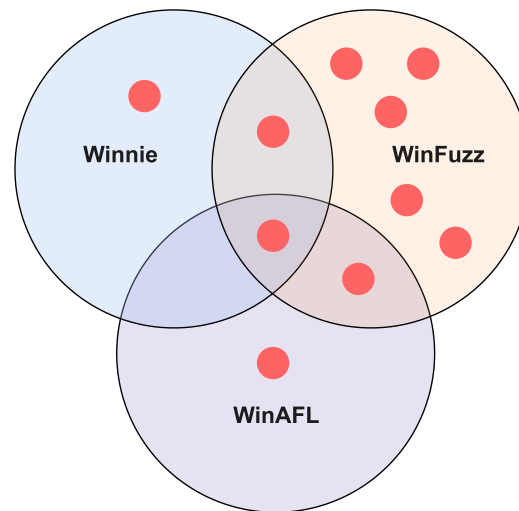


\* $p=0.216$

Average improvement: 15% vs. Winnie and 5% vs. WinAFL

# Bug Discovery Time

- We compared the average time taken to find specific bugs
- Some benchmarks (flac, nanosvg, audiofile) used older versions to increase bug count
- 10 unique bugs found across all fuzzers/trials
  - Winnie: 3/10
  - WinAFL: 3/10
  - WinFuzz: 8/10



Binary	Category	WINFUZZ	Winnie	WinAFL
<b>flac</b>	Null ptr deref	12.25 s	15.6 s	243.8 s
<b>nconvert</b>	Illegal address	2.1 hrs	X	X
<b>nconvert</b>	Invalid free	3.6 hrs	X	X
<b>nconvert</b>	Invalid ptr deref	7.2 hrs	X	X
<b>nconvert</b>	Heap overflow	8.5 hrs	15.8 hrs	X
<b>nconvert</b>	Illegal address	X	4.5 hrs	X
<b>nanosvg</b>	Stack overflow	1.4 min	X	X
<b>nanosvg</b>	Null ptr deref	1.8 min	X	X
<b>nanosvg</b>	Null ptr deref	X	X	21.9 hrs
<b>audiofile</b>	Illegal Address	12 min	X	5.2 hrs
	WINFUZZ's speedup		1.56x	23x



# Undiscovered Bugs

- We ran additional experiments with WinFuzz to find 0-day bugs
- All bugs were reported to authors

<b>Binary</b>	<b>Description</b>
<b>nconvert</b>	2 invalid ptr reads, 3 invalid ptr writes
<b>audiofile</b>	Infinite loop
<b>jhead</b>	Invalid ptr read
<b>flvmeta</b>	2 invalid ptr reads
<b>gpmf-parser</b>	1 invalid ptr read, 1 invalid ptr write
<b>gpmf-parser</b>	Invalid ptr write
<b>pdf2json</b>	Stack buffer overrun (ntdll.dll)
<b>pdf2json</b>	Stack buffer overrun (pdf2json.exe)
<b>pdf2json</b>	Stack overflow
<hr/>	
Total 0-day bugs:	9

# Thank you!

## Q&A

(Open source release pending)

[github.com/ForTE-Research](https://github.com/ForTE-Research)