


Detecting Microarchitectural Vulnerabilities via Fuzz Testing of White-box CPUs

Bo (Brian) Fu
Simon Fraser University

Alaa R. Alameldeen
Simon Fraser University

Gururaj Saileshwar
NVIDIA Research



Existing Work – Revizor (ASPLOS '22)

Idea: Fuzz μ arch with generated test cases

- Randomly generated programs
 - Bounded memory accesses to a sandbox
- Randomly generated architectural state (“Input”)
 - Register values and memory contents

Speculation Contracts

- Specify expected μ arch side effects
- Augment ISA with “allowed” speculation clauses

Trace-generation Framework

- Contract Traces
 - Expected observations from given contract; via Unicorn emulator
- Hardware Traces
 - Actual observations from hardware; via Prime+Probe

	Observation Clause	Execution Clause
Load	expose: ADDRESS	None
Store	expose: ADDRESS	None
Cond. Jump	None	speculate: if(INVERTED_CONDITION){ IP = IP + TARGET}
Other	None	None

Example: MEM-COND

-Observe memory addresses, collect for both correct and mispredicted paths

$$\text{Contract}(\text{Prog}, \text{Data}) = \text{Contract}(\text{Prog}, \text{Data}')$$

$$\implies \text{Attack}(\text{Prog}, \text{Data}, \text{Ctx}) = \text{Attack}(\text{Prog}, \text{Data}', \text{Ctx})$$

For the same program, test with different inputs, and check if traces match.

If NOT: Contract violation, Side-Channel found!

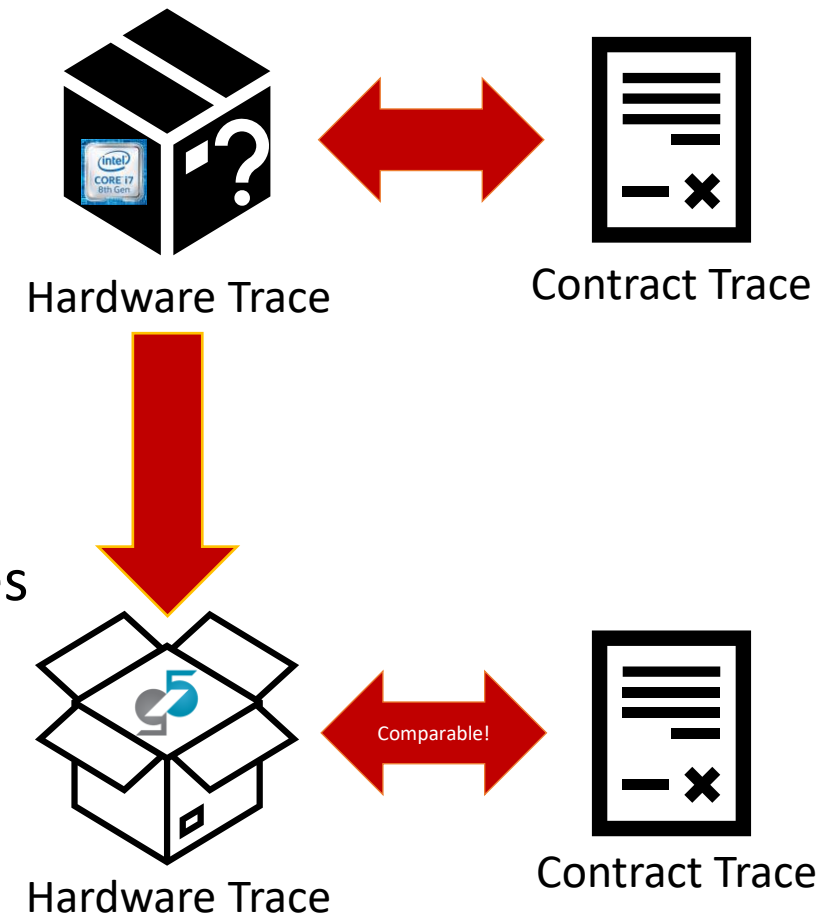
ISSUE:

- No relation between contract and hardware traces directly
 - No insight into how hardware trace is generated
- ⇒ Can only construct an existence proof of a vulnerability

SOLUTION - White-box system:

- Can understand exactly how hardware trace is generated;
 - E.g. Which cache block miss caused traces to differ
- Can correlate contract (expected) and hardware (actual) traces
- Much easier to understand what is wrong

Importantly: Provides a method to validate (or disprove) existing speculation defenses e.g. Invisi-Spec, CleanupSpec



Example – Spectre v1 (Spec. Store Bypass):

```

401- 1343009500: system.cpu.icache.tags: MISS: PC 0x407f, Vaddr 0x4080, Paddr 0x4080
402- 1343011500: system.l2.tags: MISS: PC 0x407f, Vaddr 0x4080, Paddr 0x4080
403- 1343542000: system.cpu.dcache.tags: MISS: PC 0x4049, Vaddr 0x2c100, Paddr 0x1af100
404- 1343543000: system.l2.tags: MISS: PC 0x4049, Vaddr 0x2c100, Paddr 0x1af100
405- 1343545000: system.cpu.icache.tags: MISS: PC 0x40bc, Vaddr 0x40c0, Paddr 0x40c0
406- 1343546000: system.l2.tags: MISS: PC 0x40bc, Vaddr 0x40c0, Paddr 0x40c0

401+ 1343061000: system.l2.tags: MISS: PC 0x403f, Vaddr 0x4040, Paddr 0x4040
402+ 1343134500: system.cpu.icache.tags: MISS: PC 0x407f, Vaddr 0x4080, Paddr 0x4080
403+ 1343136500: system.l2.tags: MISS: PC 0x407f, Vaddr 0x4080, Paddr 0x4080
404+ 1343667000: system.cpu.dcache.tags: MISS: PC 0x4049, Vaddr 0x2c280, Paddr 0x1af280
405+ 1343668000: system.l2.tags: MISS: PC 0x4049, Vaddr 0x2c280, Paddr 0x1af280
406+ 1343670000: system.cpu.icache.tags: MISS: PC 0x40bc, Vaddr 0x40c0, Paddr 0x40c0

```

For 2 different inputs: Can see which memory addresses were accessed by the transient miss; 0x2c100 for the RHS, 0x2c280 for the LHS (VAddr)

```

: 0x00004047: JZ_I : rdip t1, t1
: 0x00004047: JZ_I : limm t2, 0x6
: 0x00004047: JZ_I : wrip t1, t2
: 0x00004049: MOV_R_M : ld rax, DS: ld
0x00004049: MOV_R_M : ld rax, DS: ld rax, DS x
: 0x0000404d: JMP_I : limm t2, 0x4
: 0x0000404d: JMP_I : wrip t1, t2
: 0x0000404f: MOV_R_M : ld rax, DS: ld
: 0x00004053: MFENCE
: 0x00004056: PUSH_R : st rax, SS: st r

```

Can be sure a contract violation occurs & results are not false positive.

Great speedup vs. original Revizor, which must repeatedly check to ascertain findings are real.

Can also measure other µarch structures i.e. LSQ Unit, TLB, BTB, etc. & look for possible side channels there! (via Speculation Contract Violation)

Goals & Future of our work:

- Implement our changes in existing gem5 implementations of speculative defenses (Invisi-Spec, CleanupSpec, etc.) and validate them, finding existing known vulnerabilities (and maybe unknown ones!)
- Smart mutation of generated programs
 - Don't test previously tested μ arch states/contexts
 - Can directly set initial μ arch context as another input
- Detect vulnerabilities in future proposed defenses
 - Create an extensible fuzzing framework that any gem5-based defense may be plugged into, after running a script to make our required modifications (within gem5)