Theory and Practice of Finding Eviction Sets

Pepe Vila IMDEA Software Institute Boris Köpf Microsoft Research José F. Morales IMDEA Software Institute





CACHE

Find addresses that collide in cache: i.e. addresses mapped into the same cache set





CACHE

Find addresses that collide in cache: i.e. addresses mapped into the same cache set





CACHE

Find addresses that collide in cache: i.e. addresses mapped into the same cache set





CACHE

Find addresses that collide in cache: i.e. addresses mapped into the same cache set

Find associativity many colliding addresses: i.e. an **eviction set**





Efficient attacks require <u>small eviction sets</u>

Efficient attacks require <u>small eviction sets</u>

Prime+Probe



Efficient attacks require <u>small eviction sets</u>

Prime+Probe



Rowhammer



Efficient attacks require <u>small eviction sets</u>

Prime+Probe



Rowhammer



Spectre



Problem





Problem

In some scenarios, even unknown virtual address



Problem



Contributions

Systematic study of the problem of finding eviction sets

Contributions

Systematic study of the problem of finding eviction sets

Find eviction sets in O(n) compared to previous $O(n^2)$

Contributions

Systematic study of the problem of finding eviction sets

Find eviction sets in O(n) compared to previous $O(n^2)$

Reliability and performance evaluation of algorithms in real hardware

Finding minimal eviction sets

Find a large eviction set for an address V:

- Pick "enough" addresses at random

- Timing test:
$$a_v a_0 a_1 \dots a_n a_v^{\circ}$$

Finding minimal eviction sets

Find a large eviction set for an address V:

- Pick "enough" addresses at random

$$- \operatorname{Timing test:} a_v a_0 a_1 \dots a_n a_v^{\mathfrak{S}}$$



Reduce initial large eviction set into its minimal core

Finding minimal eviction sets

Find a large eviction set for an address V:

- Pick "enough" addresses at random

$$= \operatorname{Timing test:} a_v a_0 a_1 \dots a_n a_v^{\mathfrak{S}}$$



Reduce initial large eviction set into its minimal core



















Until we find an element C such that when removed the remaining set stops being an eviction set:

TEST(S\{C}) = False


























Until we have identified ASSOCIATIVITY many elements representing the eviction set's core!



O(N²) memory accesses



Threshold Group Testing

Group testing problem by Robert Dorfman (1943)



Blood samples

Threshold Group Testing

Group testing problem by Robert Dorfman (1943)



Blood samples

Threshold Group Testing

Group testing problem by Robert Dorfman (1943)

Generalization by Peter Damaschke (2006):

- Positive test only if at least "u" defectives
- Negative test only if at most "I" defectives
- Random otherwise



Blood samples



Y-right (lines): Average running time for eviction set reductionY-left (columns): Cost of finding an initial eviction set of certain sizeX: Eviction set size in number of addresses

Robustness Evaluation

Modern <u>replacement policies</u> break our test assumption and introduce errors.

- X: Cache set offset (each points aggregates all slices)
- Y: Average success rate for

Green: reduction rate w/o error correcting mechanisms.

Yellow: Test rate reliability

Find minimal eviction sets on Chrome with JS and Wasm

Conclusions

Finding minimal eviction sets is a threshold group-testing problem: new insight for research on principled countermeasures

Novel linear-time algorithm makes attacks faster and enables them in scenarios previously considered impractical

Thanks for your attention

Questions?

Noise on difference machines and cache sets

Demo verification scripts

```
run.sh:
google-chrome-beta --user-data-dir=/tmp/tmp.u9lo18kaTh
--js-flags='--allow-natives-syntax --experimental-wasm-bigint'
http://localhost:8000/ | ./verify_addr.sh
```

```
--allow-natives-syntax: used for printing found indices to stdout
--experimental-wasm-bigint: only for convenience, will have default support soon
```

```
verify addr.sh:
```

- find chrome's PID
- use pmap to find base virtual address for JS buffer
- read JS indices and add them to virtual address base
- execute ./virt_to_phys to translate virtual to physical addresses using /proc/pid/pagemap
- extract slice and cache index set from physical address (uses Intel's reverse engineered hash function)

Example of TEST() in Wasm

TurboFan x86 64 output

	0 55 1 4889e5 4 6a0a 6 56 7 4883ec10 b 488b9ea7000000 12 6666660f1f840000000000 1d 0f1f00 20 488b96c7000000 27 483922 2a 0f834000000 30 8bc0 32 49ba000000001000000 3c 4c3bd0 3f 7320 61 488b0403 65 4883f800 69 75b5 6b 488be5 6e 5d 6f c3	<pre>push rbp movq rbp,rsp push 0xa push rsi subq rsp,0x10 movq rbx,[rsi+0xa7] nop nop movq rdx,[rsi+0xc7] cmpq [rdx],rsp jnc <+0x70> mov1 rax,rax movq r10,0x100000000 cmpq r10,rax jnc <+0x61> movq rax,[rbx+rax*1] cmpq rax,0x0 jnz <+0x20> movq rsp,rbp pop rbp ret1</pre>
--	---	---

traverse.wat

```
(func $x (param $ptr i64)
 (loop $iter
    (set_local $ptr
        (i64.load
            (i32.wrap/i64 (get_local $ptr))))
    (br_if $iter
            (i32.eqz (i64.eqz (get_local $ptr)))))
)
```