

[BEN GRAS/@BJG](mailto:ben.gras@bjg), KAVEH RAZAVI, CRISTIANO GIUFFRIDA, HERBERT BOS
VRIJE UNIVERSITEIT AMSTERDAM



NLUUG 2018



TLBLEED:
WHEN PROTECTING YOUR CPU CACHES IS NOT ENOUGH



TEASER


- We would like to protect against cache attacks generically
- You won't believe this one thing that people forget



OVERVIEW

- Side channels
- Cache attacks
- TLBleed
- Evaluation

CACHE SIDE CHANNELS



SIDE CHANNELS

- Leak secrets outside the regular interface




- The first combination safes in the 1950s

EXAMPLE: FLUSH+RELOAD

- Can attack AES implementation with T tables
- A table lookup happens $T_j [x_i = p_i \oplus k_i]$
- p_i is a plaintext byte, k_i a key byte
- We can detect lookups into the table using F+R

EXAMPLE: FLUSH+RELOAD


- Again: secrets are betrayed by memory accesses
- Known plaintext + accesses = key recovery



EXAMPLE: LIBCRYPT ECC





EXAMPLE: LIBCRYPT ETC




DEFENCE EXAMPLE: TSX

- Intel TSX: Transactional Synchronization Extensions
- Intended for hardware transactional memory
- But relies on unshared cache activity
- Transactions fit in cache, otherwise auto-abort
- We can use this as a defence - all solved now right?



TLBLEED




TLBLEED: TLB AS SHARED STATE

- Other structures than cache shared between threads?
- What about the TLB?
- Documented: TLB has L1iTLB, L1dTLB, and L2TLB
- Not documented: structure

TLBLEED: TLB AS SHARED STATE

- Let's experiment with performance counters
- Try linear structure first
- All combinations of ways (set size) and sets (stride)
- Smallest number of ways is it
- Smallest corresponding stride is number of sets



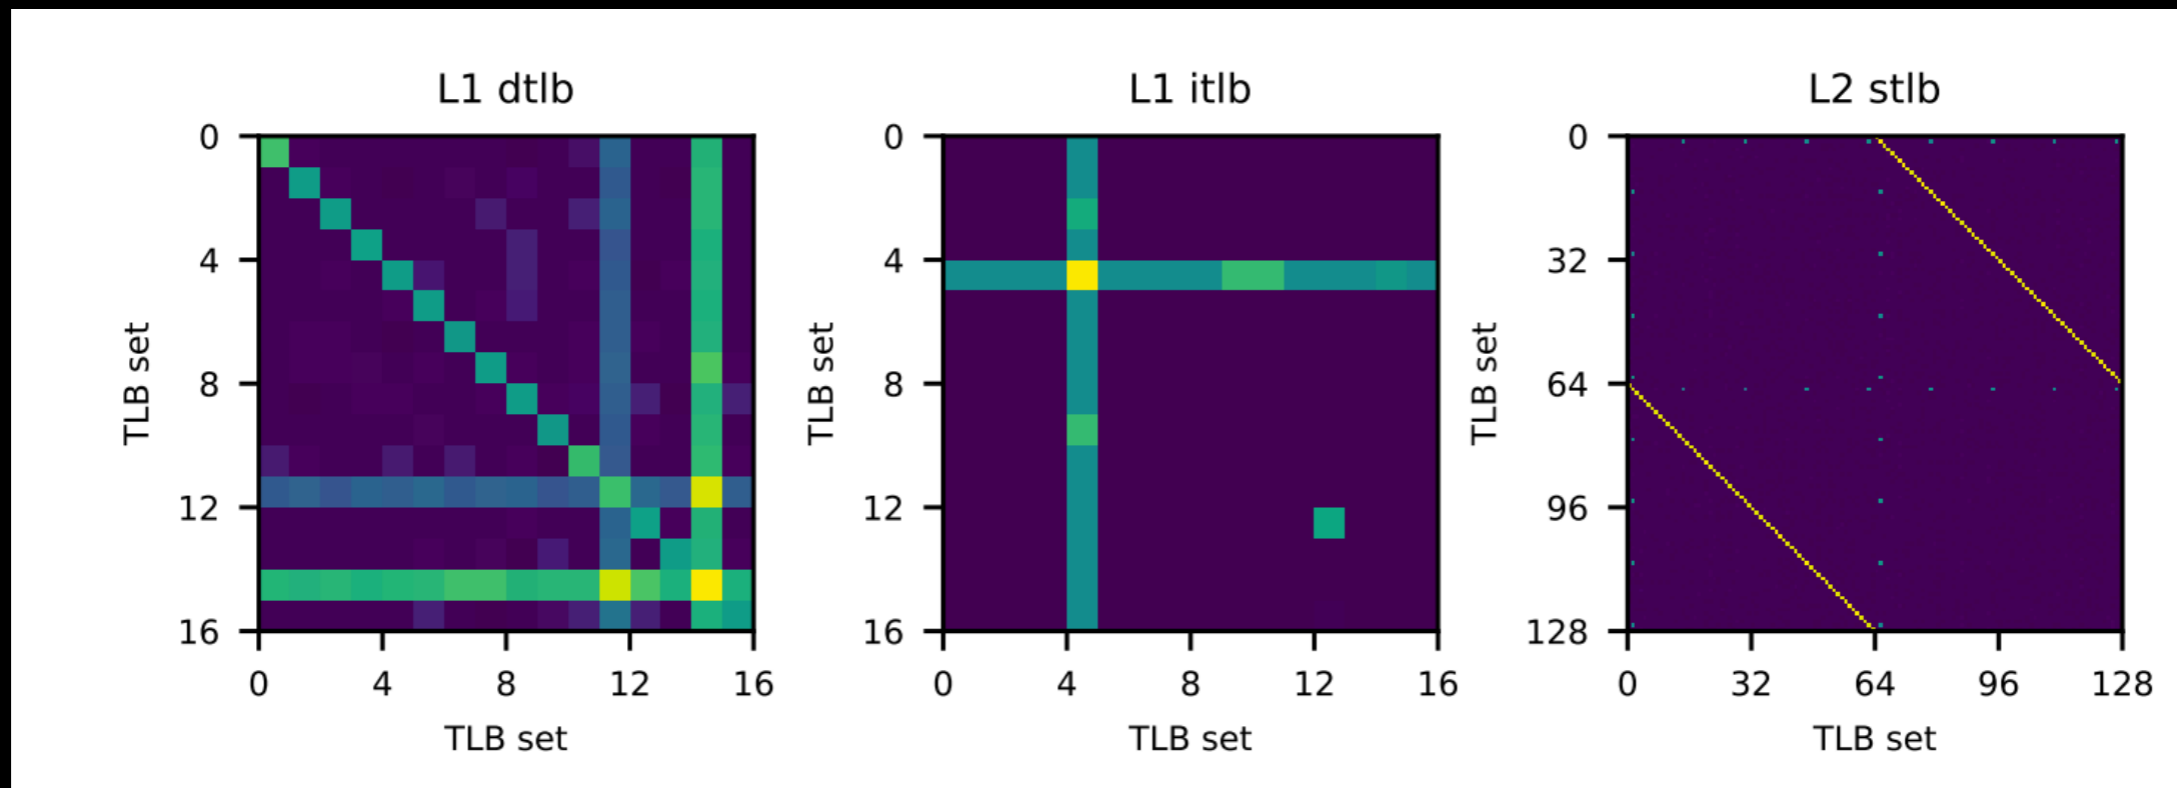
TLBLEED: TLB AS SHARED STATE

- For L2TLB:
We reverse engineered a more complex hash function
- Skylake XORs 14 bits, Broadwell XORs 16 bits
- Represented by this matrix, using modulo 2 arithmetic

$$H = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

TLBLEED: TLB AS SHARED STATE

- Let's experiment with performance counters
- Now we know the structure..
Are TLB's shared between hyperthreads?
- Let's experiment with misses when accessing the same set




TLBLEED: TLB AS SHARED STATE

Name	year	L1 dTLB					L1 iTLB					L2 sTLB				
		set	w	pn	hsh	shr	set	w	pn	hsh	shr	set	w	pn	hsh	shr
Sandybridge	2011	16	4	7.0	lin	✓	16	4	50.0	lin	✗	128	4	16.3	lin	✓
Ivybridge	2012	16	4	7.1	lin	✓	16	4	49.4	lin	✗	128	4	18.0	lin	✓
Haswell	2013	16	4	8.0	lin	✓	8	8	27.4	lin	✗	128	8	17.1	lin	✓
HaswellXeon	2014	16	4	7.9	lin	✓	8	8	28.5	lin	✗	128	8	16.8	lin	✓
Skylake	2015	16	4	9.0	lin	✓	8	8	2.0	lin	✗	128	12	212.0	XOR-7	✓
BroadwellXeon	2016	16	4	8.0	lin	✓	8	8	18.2	lin	✗	256	6	272.4	XOR-8	✓
Coffeelake	2017	16	4	9.1	lin	✓	8	8	26.3	lin	✗	128	12	230.3	XOR-7	✓

TLBLEED: TLB AS SHARED STATE

- Can we use only latency?
- Map many virtual addresses to same physical page




TLBLEED: TLB AS SHARED STATE

- Let's observe EdDSA ECC key multiplication
- Scalar is secret and ADD only happens if there's a 1
- Like RSA square-and-multiply
- But: we can not use code information! Only data..!


```
void _gcry_mpi_ec_mul_point (mpi_point_t result,  
    gcry_mpi_t scalar, mpi_point_t point,  
    mpi_ec_t ctx)  
{  
    ...  
    for (j=nbits-1; j >= 0; j--) {  
        _gcry_mpi_ec_dup_point (result, result, ctx);  
        if (mpi_test_bit (scalar, j))  
            _gcry_mpi_ec_add_points(result,result,point,ctx);  
    }  
    ...  
}
```

TLBLEED: TLB AS SHARED STATE




TLBLEED: TLB AS SHARED STATE


- Monitor a single TLB set and use temporal information
- Use machine learning (SVM classifier) to tell the difference



TLBLEED: TLB AS SHARED STATE

- Monitor
- Use






EVALUATION

TLBLEED RELIABILITY: ECC

Microarchitecture	Trials	Success	Median BF
Skylake	500	0.998	$2^{1.6}$
Broadwell	500	0.982	$2^{3.0}$
Coffeelake	500	0.998	$2^{2.6}$
Total	1500	0.993	




- Single trace capture: 1ms
- Median end-to-end time: 17s

Microarchitecture	Trials	Success	Median BF
Broadwell (CAT)	500	0.960	$2^{2.6}$
Broadwell	500	0.982	$2^{3.0}$

TLBLEED RELIABILITY: RSA

- 1024-bit RSA square-and-multiply in libgcrypt
- Old version
- F+R hardened: conditional pointer swap



RECEPTION

- Intel: dismissed TLBleed
- OpenBSD disabled Intel HT
- Widespread media coverage, logo thanks to TheRegister
- Wikipedia

```
CVS: cvs.openbsd.org: src
Mark Kettenis | Tue, 19 Jun 2018 12:30:19 -0700
CVSROOT: /cvs
Module name: src
Changes by: kette...@cvs.openbsd.org 2018/06/19 13:29:52

Modified files:
sys/arch/amd64/amd64: cpu.c
sys/arch/amd64/include: cpu.h
sys/kern : kern_sched.c kern_sysctl.c
sys/sys : sched.h sysctl.h

Log message:
SMT (Simultaneous Multi Threading) implementations typically share
TLBs and L1 caches between threads. This can make cache timing
attacks a lot easier and we strongly suspect that this will make
```



TLBleed

From Wikipedia, the free encyclopedia

TLBleed is a cryptographic *side-channel attack* that uses *machine learning simultaneous multithreading*.^{[1][2]} As of June 2018, the attack has only vulnerable to a variant of the attack, but no proof of concept has been

The attack led to the **OpenBSD** project disabling simultaneous multithreading, which theoretically be prevented by preventing tasks with different security co

References [edit]

- [↑] Williams, Chris (2018-06-22). "Meet TLBleed: A crypto-key-leaking C
- [↑] *abc* Varghese, Sam (25 June 2018). "OpenBSD chief de Raadt says

CREDIT

- Work also by Kaveh Razavi, Cristiano Giuffrida, Herbert Bos
- Some diagrams in these slides were taken from other work:
FLUSH+RELOAD, Cloak
- Yuval Yarom, Katrina Falkner, Peter Peßl, Daniel Gruss



CONCLUSION

- Practical, reliable, high resolution side channels exist outside the cache
- They bypass defenses
- @bjg @kavehrazavi
- @vu5ec
- vusec.net/projects/tlbleed/
- Thank you for listening

