Revisiting Microarchitectural Side-Channels by Miro Haller

Responsible: Prof. Serge Vaudenay EPFL / LASEC Supervisor: Muhammed Fatih Balli EPFL/LASEC



Motivation

Cache Architecture

Prime+Probe

Applications I – OpenSSL AES-CBC

Prime+Probe on Physically Indexed Caches

Applications II – Argon2d

CacheSC

References



Motivation



3

4

5

6

7

8

Cache Side-Channel History

- 1996: Kocher publishes first side-channel attacks against PKC systems [1]
- 2002: Page presents theoretical cache attacks [2]
- 2005: Osvik et al. introduce Prime+Probe and implement practical cache attacks [3, 4]
- 2013: Yarom and Falkner apply Flush+Reload to LLC leading to cross-core and cross-VM attacks [5]
- 2016: Gruss et al. introduce the stealthy Flush+Flush cache attack [6]

Increased Complexity



EPFL



- Processors evolved significantly
- Side-Channels likely to persist
- Empirical approach
- Applicable to recently introduced algorithms?
 - Argon2d

11.06.2020



Cache Architecture



4

5

6

7

8

Memory Subsystem



concrete numbers for Ivy Bridge with Intel $\ensuremath{\mathbb{R}}$ CoreTM i7-3520M CPU @ 2.90GHz

Miro Haller







Prime+Probe





Prime+Probe in Theory



11.06.2020

11



Prime+Probe in Practice I

- Precise time measurements
 - Out-of-order execution
 - Fencing
 - Dynamic frequency scaling
 - CPU governor
 - BIOS
 - Busy loop heuristic

1	cpuid
2	rdtsc
3	/* store timestamp */
4	/* measured code */
5	rdtscp
6	/* store timestamp */
7	cpuid

Code benchmarking following the recommendation of Intel's benchmarking white paper [7]



Prime+Probe in Practice II







11.06.2020

S



Prime+Probe in Practice IV

- Further issues
 - Compiler
 - Function call overhead
 - Approximate LRU
 - Write buffers and LFBs
 - Branch prediction



Lessons Learned

- Simplicity is key
 - Reduce noise
 - Synchronous measurement
 - Increase complexity step-by-step
- The data structure is essential
 - Minimize cache pollution
 - Avoid hardware prefetching
 - Reduce overhead





Applications I – OpenSSL AES-CBC



AES-CBC with T-Tables

- T-tables Te₀, Te₁, Te₂, Te₃ combine the SubBytes, ShiftRows, and MixColumns transformations into lookup tables
- Set of bytes $\mathbb{B}=\{0,1\}^8,$ round $0\leq r\leq 10$
- Plaintext block $P = p_0 || p_1 || ... || p_{15}$, key $K = k_0 || k_1 || ... || k_{15}$
- Key schedule produces key matrices $K^{(r)} \in \mathbb{B}^{4 \times 4}$ for each round
- Let $A^{(r)} \in \mathbb{B}^{4 \times 4}$ be the input to the T-tables in round r
- $A_{i,j}^{(0)} = p_l \bigoplus k_l$ for $l = i + 4j, 0 \le i, j < 4$
- $A_{:,j}^{(r+1)} = Te_0 \left[A_{0,j}^{(r)} \right] \bigoplus Te_1 \left[A_{1,j+1}^{(r)} \right] \bigoplus Te_2 \left[A_{2,j+2}^{(r)} \right] \bigoplus Te_3 \left[A_{3,j+3}^{(r)} \right] \bigoplus K_{:,j}^{(r)}$ for $0 \le r < 9$
- Last round r = 9 uses another table Te_4



3

5

6

7

8

CPA attack on r = 0 Table Lookups

- First round table lookups depend on key and plaintext: $Te_t[p_l \bigoplus k_l]$ for $l = t + 4s, 0 \le t, s < 4$
- Lookup table entries are 4 bytes
- Cache set accesses leak 4 bits of k_l:

$$Te_t[a] \Longleftrightarrow *(Te_t+a) \Longleftrightarrow egin{array}{c} ptr_{Te_t} \ + a \cdot 4B \end{array} \leftrightarrow egin{array}{c} + b_{63}b_{62}\dots b_{12} \end{array} egin{array}{c} b_{11}b_{10}b_9b_8b_7b_6 \ a_7a_6a_5a_4 \end{array} egin{array}{c} b_5b_4b_3b_2b_1b_0 \ a_7a_6a_5a_4 \end{array} egin{array}{c} a_3a_2a_1a_000 \end{array} egin{array}{c} offset \end{array}$$

• Extendable to recover half of the entire key







5

6

7

8

Conclusion

- CPA on AES-CBC of OpenSSL-0.9.8 is reproducible on contemporary hardware
- OpenSSL-1.1.1 still includes this code
- But unlikely to be executed today
 - AES-NI



Prime+Probe on Physically Indexed Caches



Why Target Physically Indexed Caches?

- Common architecture:
 - L1 virtually indexed
 - L2, L3 physically indexed
- Prime+Probe on L1 or Flush+Reload on L3
- Advantages of Prime+Probe on L2:
 - Larger than L1
 - No shared addresses needed
 - Higher cache miss penalty than for L1
- Challenge:
 - Building cache data structure







Building the Data Structure



11.06.2020



5

6

7

8

Limitations

- Limitations:
 - Privileged version requires CAP_SYS_ADMIN capability
 - Unprivileged version is slow
 - Precise L1 sets, permutation of L2 sets





Applications II – Argon2d



5

6

7

8

Argon2 [11]

- Memory hard password hash function
- Two variants:
 - Argon2d:
 - Data-dependent memory accesses
 - Faster
 - Suitable for cryptocurrencies
 - Argon2i:
 - Data-independent memory accesses
 - Preferred for password hashing



Argon2 [11]



Figure 2 from the Argon2 paper by Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich [11].



4

5

6

7

8

Motivation

- Password-dependent cache access pattern
- Efficient password cracking through early abortion
- Depends on cache observation accuracy

target access pattern





Argon2d Observation Granularity I



11.06.2020



Argon2d Observation Granularity II

• Empirical results:

Ratio of Target Sets	$\mu^{\mathcal{A}}_{\mathcal{C}_{P+P}}$	$\sigma^{\mathcal{A}}_{\mathcal{C}_{P+P}}$	$\mu_{C_{blocks}}^{\mathcal{V}}$	$\sigma_{{\cal C}_{blocks}}^{{\cal V}}$
512/512	34.4	2.3	6517.7	1460.9
32/512	595.6	41.9	6520.9	1445.6

- Results:
 - Independent of the attacker
 - Victim accesses approx. 25× the number of cache lines in L2
- Further research:
 - Required granularity to distinguish access patterns
 - Slow down victim



7 CacheSC



Conclusion

- Prime+Probe on L1
 - Classic OpenSSL AES-CBC CPA on contemporary hardware
- Prime+Probe on L2
 - Cache observation granularity of Argon2d
- Insights into cache attack data structure design
- CacheSC

GitHub

"Library for Prime+Probe cache side-channel attacks on L1 and L2"



References



References

- I. Paul C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In Neal Koblitz, editor, Advances in Cryptology — CRYPTO '96, pages 104–113, Berlin, Heidelberg, 1996. Springer Berlin Heidelberg.
- 2. Dan Page. Theoretical use of cache memory as a cryptanalytic side-channel. IACR Cryptology ePrint Archive, 2002:169, 01 2002.
- 3. Dag Arne Osvik, Adi Shamir, and Eran Tromer. Cache Attacks and Countermeasures: The Case of AES. January 2005.
- 4. Daniel J. Bernstein. Cache-timing attacks on AES. 2005.
- 5. Yuval Yarom and Katrina Falkner. FLUSH+RELOAD: A High Resolution, Low Noise, L3 Cache Side-Channel Attack. In 23rd USENIX Security Symposium (USENIX Security 14), pages 719–732, San Diego, CA, August 2014. USENIX Association.
- 6. Daniel Gruss, Clémentine Maurice, Klaus Wagner, and Stefan Mangard. Flush+ Flush: a fast and stealthy cache attack. In International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment, pages 279–299. Springer, 2016.
- 7. Gabriele Paoloni. How to Benchmark Code Execution Times on Intel® IA-32 and IA-64 Instruction Set Architectures. <u>https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/ia-32-ia-64-benchmark-code-execution-paper.pdf</u>, 2010-09.



6

8

References

8. Alex Biryukov, Daniel Dinu, and Dmitry Khovratovich. Argon2: new generation of memory-hard functions for password hashing and other applications. In 2016 IEEE European Symposium on Security and Privacy, pages 292–302. IEEE, 2016.



Backup Slides

EPFL

Single Eviction



Measure all cache lines in the same set simultaneously

Measure all cache lines individually and sum those of the same cache set





Normalization



OpenSSL AES-CBC library call without normalization

OpenSSL AES-CBC library call with normalization

Trimming

EPFL



without trimming

Prime+Probe (without any evictions) with trimming

EPFL OpenSSL CPA Accuracy

