GoFetch: Breaking Constant-Time Cryptographic Implementations Using Data Memory-Dependent Prefetchers

Boru Chen, Yingchen Wang, Pradyumna Shome,

Christopher Fletcher, David Kohlbrenner, Riccardo Paccagnella,

Daniel Genkin



Timing Attacks





// secret = 1 or 0
if (secret)
{
 trash = *addr_A
}







secret = 1





















This talk:

Show that these principles are insufficient.



This talk:

Show that these principles are insufficient.





Data Memory Dependent Prefetcher

This talk:

Show that these principles are insufficient.





Data Memory Dependent Prefetcher

secret	=	*non-sec-addr
secret	=	<pre>*non-sec-addr</pre>
secret	=	<pre>*non-sec-addr</pre>
secret	=	*non-sec-addr

program *without* secret-dependent load

This talk:

Show that these principles are insufficient.





Data Memory Dependent Prefetcher

This talk:

Show that these principles are insufficient.



ý

Data Memory Dependent Prefetcher

	<pre>secret = *non-sec-addr secret = *non-sec-addr</pre>	Cache Timing Attacks
r	<pre>secret = *non-sec-addr secret = *non-sec-addr</pre>	*secret
•	program <i>without</i> secret-dependent load	secret

Apple DMP could **treat loaded data as memory address** and perform access.

Apple DMP could **treat loaded data as memory address** and perform access.

program load



Apple DMP could **treat loaded data as memory address** and perform access.

program load





Apple DMP could **treat loaded data as memory address** and perform access.

program load





Apple DMP could **treat loaded data as memory address** and perform access.

program load





Apple DMP could **treat loaded data as memory address** and perform access.

program load





Apple DMP could **treat loaded data as memory address** and perform access.

program load





Apple DMP could **treat loaded data as memory address** and perform access.

program load





Apple DMP could **treat loaded data as memory address** and perform access.

program load





Augury¹

• Comprehensive reverse engineering of Apple DMPs.



- Comprehensive reverse engineering of Apple DMPs.
- Develop DMP-aided choseninput attack framework.



- Comprehensive reverse engineering of Apple DMPs.
- Develop DMP-aided choseninput attack framework.
- Undermine four cryptographic implementations in the wild or submitted to NIST PQC standardization.



How do classical prefetchers work?





Does memory access pattern even matter?

// Array-of-pointer pattern
for (i = 0; i < M; i++)
 trash += *arr[i];</pre>



Does memory access pattern even matter?




Key Observations of Apple DMP

Does memory access pattern even matter?







Where does the DMP scan for pointers?

// Single load
trash += arr[0];



Cache Line Aligned

Key Observations of Apple DMP

How does DMP determine pointers to dereference in each line?



Cache Line Aligned

Key Observations of Apple DMP

History filter: *how DMP avoids redundant dereference?* 4GByte region: *heuristic of predicting pointer value.*

Do-not-scan hint: how DMP avoids redundant scan? Top byte ignore: *how DMP synergizes with TBI?*

...

...

Check out the paper!

cstate: crypto state

sec: secret ci: chosen input



cstate: crypto state

sec: secret ci: chosen input

cstate = ci AND sec



cstate: crypto state

sec: secret ci: chosen input

```
cstate = ci AND sec
```



Choose ci as valid pointer! cstate = ptr AND sec if sec = 0xffffffffffffffff => cstate = ptr if sec = 0x000000000000000 => cstate = 0

cstate: crypto state

sec: secret ci: chosen input

```
cstate = ci AND sec
```



Choose ci as valid pointer! cstate = ptr AND sec if sec = 0xffffffffffffffff => cstate = ptr DMP *ptr if sec = 0x000000000000000 => cstate = 0 DMP ;

Cryptanalysis for DMP exploit

End-to-end key extraction PoCs

Cryptanalysis for DMP exploit End-to-end key extraction PoCs



Cryptanalysis for DMP exploit End-to-end key extraction PoCs



CryptanalysisEnd-to-end keyfor DMP exploitextraction PoCs







$$N = p \times q$$

RSA Decryption

$$m \equiv c^d \mod N$$
$$\downarrow$$
$$N = p \times q$$







Observations of crypto state *c* mod *p*:

- *c* is chosen input and *p* is secret.
- If c < p, then $c \mod p = c$, else $c \mod p = c lp$.

Target i-th bit of \boldsymbol{p} and

set input as $c = prefix \parallel 10000 \dots \parallel ptr$!

c prefix || 10000 ... || ptr
p prefix || xxxxxxxx ... x
i-th bit

Observation of *c mod p*:

- c is chosen input and p is victim's secret.
- if c < p, $c \mod p = c$
- if $c \ge p$,

 $c \mod p = c - lp$

Target i-th bit of p and set input as $c = prefix \parallel 10000 \dots \parallel ptr$!



• if
$$p = prefix \parallel 1_{XXXXXXX} \dots x \implies p > c$$

• then $c \mod p = c = \cdots \parallel ptr$

• if *p* = *prefix* ∥ 0*xxxxxxx* ... *x*. ⇒ *p* < *c* • then *c* mod *p* = *c* − *lp* = *unknown*

Observation of *c* mod *p*:

- c is chosen input and p is victim's secret.
- if c < p, $c \mod p = c$
- if $c \ge p$,
 - $c \mod p = c lp$

Target i-th bit of \boldsymbol{p} and

set input as $c = prefix \parallel 10000 \dots \parallel ptr$!



DMP recovers upper **560** bits

- + Coppersmith²
- = full extraction!

Observation of *c mod p*:

- c is chosen input and p is victim's secret.
- if c < p, $c \mod p = c$
- if $c \ge p$,
 - $c \mod p = c lp$

• CCA-KEM built on a CPA-PKE scheme.

• CCA-KEM built on a CPA-PKE scheme.

Key Mismatch Attack³

60

c = Enc(m, pk, e)

$$c = Enc(m, pk, e)$$

62

$$c = Enc(m, pk, e)$$

$$m \quad 1/0 \quad 1/0 \quad \cdots \quad 1/0$$

$$e \quad 0 \sim q \quad 0 \sim q \quad \cdots \quad 0 \sim q$$
Kyber-512: $q = 3329$

$$c = Enc(m, pk, e)$$



3 Qin et al., "A Systematic Approach and Analysis of Key Mismatch Attacks on Lattice-Based NIST Candidate KEMs", ASIACRYPT'21.



3 Qin et al., "A Systematic Approach and Analysis of Key Mismatch Attacks on Lattice-Based NIST Candidate KEMs", ASIACRYPT'21.



3 Qin et al., "A Systematic Approach and Analysis of Key Mismatch Attacks on Lattice-Based NIST Candidate KEMs", ASIACRYPT'21.

$$sk[i] < \frac{q}{4} - e[i] \qquad \Longrightarrow \qquad m'[i] = m[i]$$
$$sk[i] \ge \frac{q}{4} - e[i] \qquad \Longrightarrow \qquad m'[i] \neq m[i]$$

67

- CCA-KEM built on a CPA-PKE scheme.
- FO transformation wraps CPA to CCA.



Observations of decrypted message *m*':

 sk is secret, m and e are chosen inputs.

• if
$$(e + sk)[i] < \frac{q}{4}$$
,
 $m'[i] = m[i]$.
• if $(e + sk)[i] \ge \frac{q}{4}$,

 $m'[i] \neq m[i].$

- CCA-KEM built on a CPA-PKE scheme.
- FO transformation wraps CPA to CCA.



Observations of decrypted message *m*':

 sk is secret, m and e are chosen inputs.

• if
$$(e + sk)[i] < \frac{q}{4}$$
,
 $m'[i] = m[i]$.
• if $(e + sk)[i] \ge \frac{q}{4}$,
 $m'[i] \ne m[i]$.

Set m as $ptr \parallel 0 \dots 0$ and tweak e[0]!

- if $(e + sk)[\mathbf{0}] < \frac{q}{4}$ • then m' = ptr||0 ... 0 ptr
- if $(e + sk)[\mathbf{0}] \ge \frac{q}{4}$
 - then m' = ptr' ||0 ... 0 No ptr

Observations of decrypted message *m*':

sk is secret, *m* and *e* are chosen inputs.

• if
$$(e + sk)[i] < \frac{q}{4}$$
,
 $m'[i] = m[i]$.

• if
$$(e + sk)[i] \ge \frac{q}{4}$$
,
 $m'[i] \ne m[i]$.

5

Set m as $ptr \parallel 0 \dots 0$ and tweak e[0]!

- if $(e + sk)[\mathbf{0}] < \frac{q}{4}$ • then m' = ptr||0 ... 0 ptr
- if $(e + sk)[\mathbf{0}] \ge \frac{q}{4}$
 - then m' = ptr' ||0 ... 0 No ptr

Kyber-512:

DMP recovers 392 indexes

- + Lattice Reduction³
- = **full** extraction!

3 May et al., "Too Many Hints – When LLL Breaks LWE", ASIACRYPT'23.

Observations of decrypted message *m*':

sk is secret, *m* and
 e are chosen
 inputs.

• if
$$(e + sk)[i] < \frac{q}{4}$$
,

$$m'[i] = m[i].$$

• if
$$(e + sk)[i] \ge \frac{q}{4}$$
,
 $m'[i] \ne m[i]$.



```
bool set_DIT_enabled(void) {
    bool was_DIT_enabled = get_DIT_enabled();
    __asm__ __volatile__("msr dit, #1");
    return was_DIT_enabled;
}
```

Enable DIT for constant-time cryptographic operations
Impact



Apple: Disable DMP with DIT=1Only works on M3.



Go: Propose an opt-in DIT mode in Go binary.



proposal: runtime: implement a DIT/DOIT mode #66450



rolandshoemaker opened this issue on Mar 21 · 20 comments

Impact



Apple: Disable DMP with DIT=1

Only works on M3.



Go: Propose an opt-in DIT mode in Go binary.



Asahi Linux: Found chicken bit to disable DMP on M1/M2.

Asahi Linux



Hector Martin @marcan@treehouse.systems

Found the DMP disable chicken bit. it's HID11_EL1<30> (at least on M2).

So yeah, as I predicted, GoFetch is entirely patchable. I'll write up a patch for Linux to hook it up as a CPU security bug workaround.

(HID4_EL1<4> also works, but we have a name for that and it looks like a big hammer: HID4_FORCE_CPU_OLDEST_IN_ORDER)

Code here: github.com/AsahiLinux/m1n1/blo... (Thanks to @dkohlbre for the userspace C version this is based off of!)

One interesting finding is that the DMP is *already disabled* in EL2 (and presumably EL1), it only works in EL0. So it looks like the CPU designers already had *some* idea that it is a security liability, and chose to hard-disable it in kernel mode. This means kernel-mode crypto on Linux is already intrinsically safe.

Apr 08, 2024, 20:17 · Edited Apr 08, 20:24 🗸 · 🕲 · Web · 🛱 97 · ★ 215

	99	☆	

Impact



Apple: Disable DMP with DIT=1Only works on M3.



Go: Propose an opt-in DIT mode in Go binary.





Asahi Linux: Found chicken bit to disable DMP on M1/M2.

Asahi Linux



Pwine Awards: Best Cryptographic Attack winner.

Conclusion

- Data memory-dependent prefetchers (DMPs) performs secret-dependent memory access to leak data.
- Exploiting DMPs to perform key extraction attacks to constant-time cryptography is feasible.



Check our Website: gofetch.fail

boruc2@illinois.edu