

Side Channel Attack on OpenSSL ECDSA

Joop van de Pol

Nigel Smart

Yuval Yarom

Outline

- Background
 - ECDSA
 - wNAF scalar multiplication
 - The Hidden Number Problem
 - The Flush+Reload technique
- Exploiting the side channel
- Recovering the key
- Results

ECDSA

Signer has a private key $1 < \alpha < q-1$ and a public key $Q = [\alpha]G$

1. Compute $h = \text{Hash}(m)$
2. Randomly select an ephemeral key $1 < k < q$
3. Compute $(x, y) = [k]G$
4. Take $r = x \bmod q$; If $r = 0$ repeat from 2
5. Take $s = (h + r \cdot \alpha) \cdot k^{-1} \bmod q$; if $s = 0$ repeat from 2
6. (r, s) is the signature

Note that $k = \left| \left(r \cdot s^{-1} \right) \cdot \alpha + \left(h \cdot s^{-1} \right) \right|_q$

wNAF Form

To compute $[k]G$, first write k in wNAF form:

$$k = \sum_{i=0}^{n-1} d_i 2^i \text{ for } d_i \in \{0, \pm 1, \pm 3, \dots, \pm(2^w - 1)\}$$

Such that if $d_i \neq 0$ then $d_{i+1} = \dots = d_{i+w+1} = 0$.

Scalar Multiplication with wNAF form

Precompute $\{\pm G, \pm[3]G, \dots, \pm[2^w-1]G\}$

$x=0$

for $i=n-1$ **downto** 0

$x = \text{Double}(x)$

if $(d_i \neq 0)$ **then**

$x = \text{Add}(x, [d_i]G)$

end

end

return x

The Hidden Number Problem [BV96]

We know enough triples t_i , u_i and z_i such that

$$v_i = |\alpha t_i - u_i|_q < q / 2^{z_i}$$

for an unknown α .

We can find α by reducing HNP to a lattice problem.


•

HNP and ECDSA [HGS01,NS03]

Recall that $k = \left| (r \cdot s^{-1}) \cdot \alpha + (h \cdot s^{-1}) \right|_q$

We want $\left| \alpha t_i - u_i \right|_q < q / 2^{z_i}$

In terms of k :



Or in terms of t_i and u_i :

$$\left| (r \cdot s^{-1}) \cdot \alpha + (h \cdot s^{-1}) \right|_q = k$$

HNP and ECDSA [HGS01,NS03]

Recall that $k = \left| (r \cdot s^{-1}) \cdot \alpha + (h \cdot s^{-1}) \right|_q$

We want $\left| \alpha t_i - u_i \right|_q < q / 2^{z_i}$

In terms of k :



Or in terms of t_i and u_i :


$$\left| (r \cdot s^{-1}) \cdot \alpha - \left(- (h \cdot s^{-1}) \right) \right|_q < q$$

HNP and ECDSA [HGS01,NS03]

Recall that $k = \left| (r \cdot s^{-1}) \cdot \alpha + (h \cdot s^{-1}) \right|_q$

We want $\left| \alpha t_i - u_i \right|_q < q / 2^{z_i}$

In terms of k :



Or in terms of t_i and u_i :

$$\left| (r \cdot s^{-1}) \cdot \alpha - (a - (h \cdot s^{-1})) \right|_q < q$$

HNP and ECDSA [HGS01,NS03]

Recall that $k = \left| \left(r \cdot s^{-1} \right) \cdot \alpha + \left(h \cdot s^{-1} \right) \right|_q$

We want $\left| \alpha t_i - u_i \right|_q < q / 2^{z_i}$

In terms of k :

$$(k - a) / 2^{z_i}$$


$n - z_i$ 0

Or in terms of t_i and u_i :



$$\left| \left(\left(r \cdot s^{-1} \right) \cdot \alpha - \left(a - \left(h \cdot s^{-1} \right) \right) \right) / 2^{z_i} \right|_q < q / 2^{z_i}$$

HNP and ECDSA – State of the Art

- Useful information:
 - l bits for l known LSBs
 - Between $l-1$ and l bits for l known MSBs
 - $l/2$ bits for arbitrary l consecutive bits

FLUSH+RELOAD [YF14]

- A cache-based side-channel attack technique
- **FLUSH** memory line
- Wait a bit
- Measure time to **RELOAD** line
 - fast-> access
 - slow-> no access
- Repeat



Outline

- Background
 - ECDSA
 - wNAF scalar multiplication
 - The Hidden Number Problem
 - The Flush+Reload technique
- **Exploiting the side channel**
- Recovering the key
- Results

Attacking OpenSSL wNAF

- Use FLUSH+RELOAD to recover the double and add chain of the wNAF calculation
 - Divide time into slots of 1200 cycles (about $0.4\mu s$)
 - In each slot, probe a memory line in the code of the Double and Add functions.

```
 $x=0$   
for  $i=n-1$  downto 0  
   $x = \text{Double}(x)$   
  if ( $d_i \neq 0$ ) then  
     $x = \text{Add}(x, [d_i]G)$   
  end  
end  
return  $x$ 
```

Sample Trace

Raw:

```
D | | | D | D |
| D | | D | |
| | D | D |
| D | | D |
| | D | | D |
| | D | | ...

| | D | | D | | A | A | | D | | D | |
D | D | | D | | D | | D | | A | A | | D
| | D | | D | | D | | A | A | | | D | |
D | | D | | A | | D | | D | | D | D |
D | D | | D | | D | | D | | A | | D | D
```

Errors occur in approx. 1 symbol in 1000.

Processed:

```
DDDADDDDDADDDDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDD
DDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDAD
DDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDD
DDDADDDDDDDDDADDDDDADDDDDADDDDDDDADDDDDADDDDDDDADDDDDADDDDD
ADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDDADDDDD
DDADDDDDADDDDDDDADDDDDADDDDDADDDDDDDADDDDDDDADDDDDADDDDDADD
```

Sample Trace

We know how to use the revealed **LSBs**

Processed:

DDDDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
DDDDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
DDDDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
DDDDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
ADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
DDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD

Sample Trace

We know how to use the revealed **LSBs**

But these give an average of 2 bits per observed signature.

Can we use the information about the **MSBs**?


Processed:

DDDDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
DDDDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
DDDDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
DDDDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
ADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD
DDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADDADD

Using the MSBs

Assume $d_{m+l}, d_m \neq 0$

Before adding $[d_m]G$, x is:



$x=0$

for $i=n-1$ **downto** 0

$x = \text{Double}(x)$

if $(d_i \neq 0)$ **then**

$x = \text{Add}(x, [d_i]G)$

end


end

return x


Using the MSBs

Assume $d_{m+l}, d_m \neq 0$


Before adding $[d_m]G$, x is:



After adding $[d_m]G$, for $d_m > 0$ it is



And for $d_m < 0$




```

x=0
for i=n-1 downto 0
  x = Double(x)
  if (di≠0) then
    x = Add(x, [di]G)
  end
end
return x
    
```


Using the MSBs

Assume $d_{m+l}, d_m \neq 0$


Before adding $[d_m]G$, x is:




After adding $[d_m]G$, for $d_m > 0$ it is



And for $d_m < 0$



Either way,



$x=0$

for $i=n-1$ **downto** 0

$x = \text{Double}(x)$

if $(d_i \neq 0)$ **then**

$x = \text{Add}(x, [d_i]G)$

end

end

return x

Observation

For many “standard” curves, q , the group order is close to a power of two. That is, $q=2^n-\varepsilon$ such that $|\varepsilon|<2^p$ for $p\ll n$.

For example for **secp256k1**

q =FF
EBAaedce6af48a03bbfd25e8cd0364141


Adding or subtracting q to an n bit number is unlikely to change the MSBs

Using all the Information

For $m > p$


Recall that
 $q = 2^n - \varepsilon$
 or
 $-aq = a\varepsilon - a2^n$

$$k + 2^{m+w}$$



$$(k + 2^{m+w}) \cdot 2^{n-m-l-1}$$


$$(k + 2^{m+w}) \cdot 2^{n-m-l-1} - 2^{n-1}$$



$$(k + 2^{m+w}) \cdot 2^{n-m-l-1} - 2^{n-1} - aq$$



$$\left| (k + 2^{m+w}) \cdot 2^{n-m-l-1} - 2^{n-1} \right|_q$$



Summary

- FLUSH+RELOAD provides a nearly perfect side channel
 - Probability of error 1/1000 symbols
 - 99% of errors are noticed
 - 2 out of every 3 observed traces are perfect
- Curve choice allows using almost half of the information in each perfect trace

Results

Previous results:

- [LN13]: 160-bit key, 100 signatures with 2 known bits
- [BPSY14]: 256-bit key, expected 200 signatures

Our results:

# perfect traces	Time (s)	Success probability
10	2.25	0.07
11	4.66	0.25
12	7.68	0.38
13	11.3	0.54

With a very high probability, observing 25 signatures yields more than 13 perfect traces.

Cache storage attacks

CT-RSA 2015

B. B. Brumley

Department of Pervasive Computing
Tampere University of Technology, Finland
`billy.brumley AT tut.fi`

21 Apr 2015



TAMPEREEN TEKNILLINEN YLIOPISTO
TAMPERE UNIVERSITY OF TECHNOLOGY

Covert channels

Rainbow Series: Light Pink Book



A communication channel is covert if it is neither designed nor intended to transfer information at all.

Covert Channel Analysis of Trusted Systems (NCSC-TG-030)

Covert timing channels

A potential covert channel is a timing channel if its scenario of use involves a process that “signals information to another by modulating its own use of system resources (e.g., CPU time) in such a way that this manipulation affects the real response time observed by the second process.”

Covert storage channels

A potential covert channel is a storage channel if its scenario of use “involves the direct or indirect writing of a storage location by one process and the direct or indirect reading of the storage location by another process.”

Crypto and side-channel attacks

High level: one party is legit, the other isn't

Timing attacks

- ▶ D. Page (DES, 2002)
- ▶ D. J. Bernstein (AES, 2005)
- ▶ C. Percival (RSA, 2005)
- ▶ Osvik et al. (AES, 2005)
- ▶ Neve and Seifert (AES, 2006)
- ▶ B. B. Brumley and R. Hakala (ECDSA, 2009)
- ▶ Aciicmez et al. (DSA, 2010)
- ▶ B. B. Brumley and N. Taveri (ECDSA, 2011)
- ▶ Yarom et al. (ECDSA, 2014)
- ▶ ...

Storage attacks



Data caching

idx	tag	data
0	de30ec	????
1	096324	????
.	.	.
.	.	.
F	61eff8	????

Cache-timing attacks

Time access to own data and infers cache hits / misses to determine victim state that caused the eviction.

Cache debugging and HW privilege separation

Invasive cache debug

- ▶ Direct reads to cache lines and metadata
- ▶ Cache footprint profiling
- ▶ HW errata or coherency

Access control

- ▶ Cannot allow unchecked direct access to cache lines (exception)
- ▶ x86 protection mode: Ring 0/3
- ▶ ARM ft. TrustZone: NS-bit

Data caching with HW privilege separation

Note priv 0 can evict priv 1 and vice versa.

idx	priv	tag	data
0	1	de30ec	????
1	0	096324	????
.	.	.	.
.	.	.	.
F	1	61eff8	????

The covert storage channel (1)

Alice (priv 0) wants to send one nibble (0x2) to Bob (priv 1).

Bob pollutes:

idx	priv	tag	data
0	1	de30ec	????
1	1	096324	????
2	1	ce5f84	????
.	.	.	.
.	.	.	.
F	1	61eff8	????

The covert storage channel (2)

Alice does one read to evict one line (0x2):

idx	priv	tag	data
0	1	de30ec	????
1	1	096324	????
2	0	b7d710	????
.	.	.	.
.	.	.	.
F	1	61eff8	????

The covert storage channel (3)

Bob then:

1. Reads from Line 0. No exception.
2. Reads from Line 1. No exception.
3. Reads from Line 2. Exception – receives 0x2.

idx	priv	tag	data
0	1	de30ec	????
1	1	096324	????
2	0	b7d710	????
.	.	.	.
.	.	.	.
F	1	61eff8	????

From cache storage to clean cache-timing traces

1. Read directly from a cache line. A processor exception indicates M, otherwise H.
2. If M go back to the first step. This requires another query because the processor exception most like wipes the cache state and/or triggers a reset.
3. If H continue with the next line.

Example

HMHHHMHHHHHHMHHH

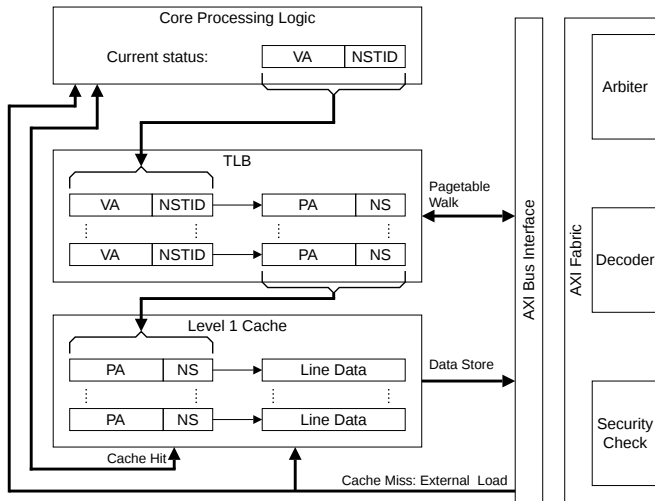
A practical architecture (1)

SRC: ARM Security Technology documentation

The content of the caches, with regard to the security state, is dynamic. Any non-locked down cache line can be evicted to make space for new data, regardless of its security state. It is possible for a Secure line load to evict a Non-secure line, and for a Non-secure line load to evict a Secure line.

A practical architecture (2)

SRC: ARM Security Technology documentation



A practical architecture (3)

SRC: ARM1156T2-S Technical Reference Manual

The purpose of the data cache Tag RAM operation is to:

- ▶ read the data cache Tag RAM contents and write into the Data Cache Debug Register.
- ▶ write into the Data Cache Debug Register and write into the Data Tag RAM.

To read the Data Tag RAM, write CP15 with:

```
; Data Tag RAM read operation  
MCR p15, 3, <Rd>, c15, c2, 0
```

Transfer data to the Data Cache Debug Register to the core:

```
; Read Data Cache Debug Register  
MRC p15, 3, <Rd>, c15, c0, 0
```

AES and side-channels

SRC: Jeff Moser

Plaintext in 4x4 grid

0	4	8	C
1	5	9	D
2	6	A	E
3	7	B	F

Initial Round

AES Crib Sheet

(Handy for memorizing)

0123
3210 Rows Row Shift

General Math

$11B = \text{AES Polynomial} = M(x)$

$x^8 + x^4 + x^3 + x + 1$ Fast Multiply

$x \cdot a(x) = (a \ll 1) \oplus (a_7 = 1) ? 1B : 00$

$\log(x \cdot y) = \log(x) + \log(y)$

Use $(x+1) = 03$ for log base

Intermediate Rounds

#	Key
9	128
11	192
13	256

S-Box (SRD)

$SRD[a] = f(g(a))$

$g(a) = a^{-1} \text{ mod } m(x)$

Sea Think $53 \oplus 63$

5 is and 3 0's $[0110\ 0011]^T$

01111000	a_7	010001
11111100	a_6	010001
00111110	a_5	010001
00011111	a_4	010001
10001111	a_3	010001
11000111	a_2	010001
11000111	a_1	010001
11100011	a_0	010001

Key Expansion: Round Constants

First Column: 01 02 04 08...

S			
0	1	B	K
M	Z	I	E
E	B	T	Y

Round Key

Other Columns:

T	E1	E1
Z	Z1	10
B	B1	B4
	F2	CA

Mix Columns:

2	1	1	3
1	3	1	1
3	1	1	1
1	1	3	1

Inverse Mix

E	B	D	9
9	E	B	D
D	9	E	B
B	D	9	E

Key Expansion: Round Constants

K	B3	01	B2
E	6E	00	6E
Y	CB	00	CB
	B7	00	B7

Mix Columns:

2	1	1	3
1	3	1	1
3	1	1	1
1	1	3	1

Inverse Mix

E	B	D	9
9	E	B	D
D	9	E	B
B	D	9	E

Prev Col \oplus Col from Previous round key

S	B2	E1
O	6E	Z1
M	CB	86
E	B7	F2


Final Round

Ciphertext

?	?	?	?
?	?	?	?
?	?	?	?
?	?	?	?

AES cache storage attack simulation

Borrowed cache-timing attack by Neve and Seifert (2006)



Conclusion

- ▶ Cache debugging mechanisms and hardware-enforced privilege separation can create covert channels
- ▶ Can potentially put crypto keys at risk
- ▶ Compared to cache-timing attacks:
 - ▶ Each “cache miss” costs an additional query
 - ▶ But the traces themselves are much more accurate
- ▶ For ARM ft. TrustZone, attack would be from Normal World kernel space to Secure World
- ▶ Concrete mitigation: implementation defined instructions
- ▶ Thanks! Questions?