

Fine-Grained Constant-Time Policies with Jasmin & EasyCrypt

Vincent Laporte;
and the Formosa Crypto team members

2022-11-24, GLSec, CNAM, Paris

Remote Side-Channel Attacks

1996: Timing attacks on implementations. . . , Paul Kocher
2003: Remote timing attacks are practical, Brumley & Boney
2005: Cache missing for fun and profit; Cache-timing attacks against AES
2010: Efficient cache attacks on AES
2011: Cache Games

2013: Lucky Thirteen
2014: Flush+Reload
2015: Subnormal floating point; Lucky microseconds
2016: Cache bleed; Armageddon
2017: May the 4th be with you
2018: Spectre
2022: Hertz bleed

Constant-Time Programming

Two (basic) simple rules

- no branching on secret data
- no memory access at secret addresses

Pros

- Simple
- Source-level counter-measure

Cons

- Unusual coding patterns
- May negatively impact performances

Constant-Time programming is a **source-level** countermeasure. . .

Compilers are too eager at optimizing

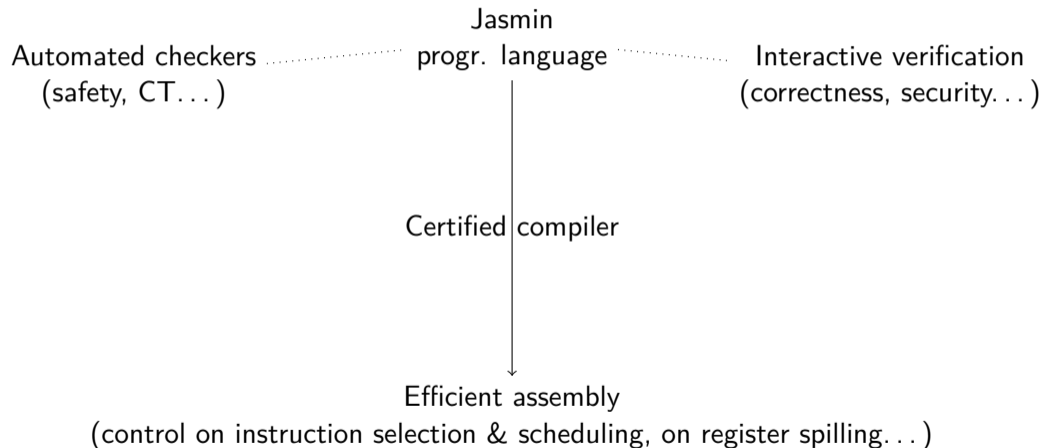
- $b \times p + (1 - b) \times q$ becomes if b then p else q
- run-time libraries have fast path for small operands

Verifying CT Security

For each branching condition and each dereferenced address prove that it is public.

- Simple **taint** analysis
- Needs prior knowledge about which inputs are public

Main difficulty: pointer analysis (aka alias analysis)



Jasmin Example

```
1 export
2 fn lehmer(reg u64 state) → reg u64 {
3   reg u64[2] s, m;
4   stack u64[2] t;
5   inline int i;
6   reg u64 j, result;
7   for i = 0 to 2 {
8     s[i] = [state + i * 8];
9   }
10  m[0] = 0x261fd0407a968add;
11  m[1] = 0x45a31efc5a35d971;
12  t = mul128(s, m);
13  result = t[1];
14  j = 0;
15  while (j < 2) {
16    [state + j * 8] = t[(int) j];
17    j += 1;
18  }
19  return result;
20 }
```

```
1 inline
2 fn mul128(reg u64[2] x, y) → stack u64[2] {
3   reg u64 xhi, ylo, lo, hi, tmp;
4   stack u64[2] r;
5   xhi = x[1];
6   ylo = y[0];
7   hi, lo = #MULX(ylo, x[0]);
8   tmp = xhi * y[0];
9   hi += tmp;
10  y[1] *= x[0];
11  y[1] += hi;
12  r[0] = lo;
13  r[1] = y[1];
14  return r;
15 }
```

Constant-Time, Formally

Instrumented Semantics

The adversary observes:

- control flow
- memory accesses:

$$f : (\vec{a}, m) \Downarrow_p^{\ell} (\vec{r}, m')$$

Syntax of Structured Leakage

$l_e ::= \bullet$	empty	$l ::= l_e := l_e$	assignment
v	value	$\text{if}_b(l_e, l)$	conditional
(l_e, \dots, l_e)	sub-leakage	$\text{while}_t(l_e, l, l)$	iteration
		$\text{while}_f(l_e)$	loop end
		$\{l; \dots; l\}$	sequence

Security property (φ -CT)

Given a “low-equivalence” relation φ between initial states:

$$\forall \vec{a}_1 \ m_1 \ \vec{a}_2 \ m_2 \ l_1 \ l_2 \ \vec{r}_1 \ \vec{r}_2 \ m'_1 \ m'_2,$$

$$(\vec{a}_1, m_1) \varphi (\vec{a}_2, m_2) \implies \left\{ \begin{array}{l} f : (\vec{a}_1, m_1) \Downarrow_p^{\ell_1} (\vec{r}_1, m'_1) \\ f : (\vec{a}_2, m_2) \Downarrow_p^{\ell_2} (\vec{r}_2, m'_2) \end{array} \right. \implies l_1 = l_2.$$

Example: Explicit Leakage

```
1 fn max(reg u64[2] a) → reg u64 {  
2   reg u64 x, y;  
3   x = a[0];  
4   y = a[1];  
5   if x < y {  
6     x = y;  
7   }  
8   return x;  
9 }
```

Applied to { 19; 23 }, it yields:

$$\{ \bullet := (\bullet, 0); \\ \bullet := (\bullet, 1); \\ \text{if}_{\#}(((\bullet, \bullet), \bullet), \{\bullet := \bullet\}) \}$$

Quizz

Is this function constant-time?

A Type-System for Constant-Time

- Two-level security lattice: $L \leq H$
- A (flow-sensitive) typing environment mapping variables to security level
- Usual type-system for non-interference with public branching only (no implicit flows)
- Selected rules:

$$\frac{\Gamma \vdash e : \tau}{\Gamma \vdash x := e : \Gamma\{x \leftarrow \tau\}} \text{ASSIGN}$$

$$\frac{\Gamma \vdash i : \Gamma_i \quad \Gamma_i \vdash c : \Gamma'}{\Gamma \vdash i; c : \Gamma'} \text{SEQ}$$

$$\frac{\Gamma \vdash b : L \quad \Gamma \vdash c_1 : \Gamma_1 \quad \Gamma \vdash c_2 : \Gamma_2}{\Gamma \vdash \text{if } b \text{ then } c_1 \text{ else } c_2 : \Gamma_1 \sqcup \Gamma_2} \text{COND}$$

$$\frac{\Gamma \vdash p : L \quad \Gamma \vdash e : \tau}{\Gamma \vdash *p := e : \Gamma} \text{STORE}$$

$$\frac{\Gamma \vdash n : L \quad \Gamma \vdash e : \tau}{\Gamma \vdash a[n] := e : \Gamma\{a \leftarrow \tau \sqcup \Gamma(a)\}} \text{ASSIGN-ARRAY}$$

A Wide Range of Adversary Models

The *base-line* adversary model does not fit all scenarios:
sometimes too coarse, sometimes too precise.

Time-Variable instructions

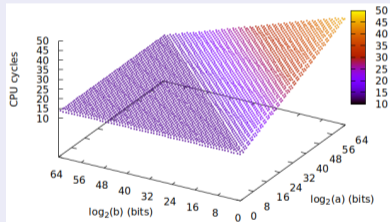


Figure 1: Timing behavior of the 64-bit div instruction on a x86 microprocessor

- Need more precise leakage model

Caches operate at a *line* level

Can observers *really* distinguish addresses within a single cache-line? Defenses under such threat model are expensive.

- Need less precise leakage model

Leakage models form a **lattice**

(A. Shivakumar, Barthe, Grégoire, Laporte, and Priya 2022)

Interactive Proofs with EasyCrypt

- A program translation $[\cdot]$ to EasyCrypt makes leakage explicit
- Security (program c is φ -CT) is reduced to a **Relational Hoare Logic** statement:

$$[c] \sim [c] : \varphi \wedge =_{\{\text{leak}\}} \implies =_{\{\text{leak}\}}$$

- Reasoning is vastly simplified through the use of automated tactics (wp, smt...)
- Two parameters model:
 - leakage of operators
 - leakage of memory accesses

```
clone import ALeakageModel as LeakageModel.  
  
module M = {  
  var leak : leakages_t  
  
  proc max (a: W64.t Array2.t) : W64.t = {  
    var aux: W64.t;  
  
    var x : W64.t;  
    var y : W64.t;  
  
    leak <- LeakAddr([0]) :: leak;  
    aux <- a.[0];  
    x <- aux;  
    leak <- LeakAddr([1]) :: leak;  
    aux <- a.[1];  
    y <- aux;  
    leak <- LeakCond(x \ult y) :: LeakAddr([]) :: leak;  
    if (x \ult y) {  
      aux <- y;  
      x <- aux;  
    }  
    return (x);  
  }  
}.
```

Application: Efficient Constant-Time Modulo

```
1 export fn mod_TV(reg u64 a, reg u64 b) → reg u64 {
2   reg u64 flag one zero dividend modulo result ...;
3   reg bool lzbz lzbz cf;
4   flag = 0x1234; one = 1; zero = 0;
5   lzbz, lzb = lzcnt(b);
6   flag = zero if lzbz;
7   lzaz, _ = lzcnt(a);
8   flag = one if lzaz;
9   lzb_m1 = #LEA(lzb - 1);
10  b_lzb_m1 = b << lzb_m1;
11  dividend = #LEA(b_lzb_m1 + a);
12  b_lzb = b_lzb_m1 << 1;
13  cf, temp2 = b_lzb + a;
14  dividend = temp2 if ! cf;
15  dividend = a if flag == 1;
16  temp2 = 0xFFFFFFFFFFFFFFFF;
17  dividend = temp2 if flag == 0;
18  modulo = dividend % b ;
19  result = modulo;
20  result = a if flag == 0;
21  return result;
22 }
```

- Efficiently uses hardware division
- Timing no longer depends on the value of the first argument (a):

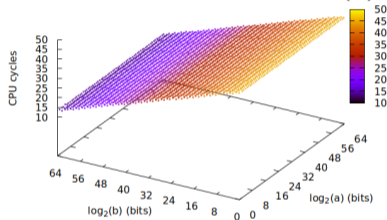


Figure 2: Timing behavior of `mod_TV`

- Proved **correct** and **secure** in EasyCrypt

Semantics Preservation (forward simulation)

If the compilation of program p produces a program p' , then its safe behaviors are preserved:

$$\forall \vec{a} \ m \ \vec{r} \ m', \quad f : (\vec{a}, m) \Downarrow_p (\vec{r}, m') \quad \Longrightarrow \quad f : (\vec{a}, m) \Downarrow_{p'} (\vec{r}, m').$$

Hidden Details

- Source and target languages are different
- Initial states are not the same (but tightly related)
- The target stack must be large enough
 - i.e., the compiler does not enforce the absence of “stack overflow”

Compilers do not preserve leakage

- common subexpression elimination removes observable events
- loop invariant code motion reorders events
- allocation of local variables introduces events
- instruction selection transforms events
- ...

Is this constant-time preserving?

Instrumented Correctness

Leakage Transformers

If the compilation of program p produces a program p' ,
there exists a leakage transformer F such that:

$$\forall \vec{a} \ m \ \ell \ \vec{r} \ m', \quad f : (\vec{a}, m) \Downarrow_p^\ell (\vec{r}, m') \quad \Longrightarrow \quad f : (\vec{a}, m) \Downarrow_{p'}^{F(\ell)} (\vec{r}, m').$$

Corollary

Constant-Time security is preserved

Beyond CT-Preservation

Leakage transformers allow to reason about **target instrumentation** at the source level
(e.g., execution time, resource usage...)

See recent work, focusing on execution cost (Barthe et al. 2021)

Branch Prediction and Spectre

```
1 u64[2] offsets = { 1, 0 };
2 u64[2] data = { 0xcafe, 0xbeef };
3
4 export
5 fn archetype(#public reg u64 n) → reg u64 {
6   reg u64 p r;
7   r = 0;
8   if n < 2 { // Branch trained with in-bounds n
9     // Speculatively, out-of-bounds
10    p = offsets[(int) n];
11    // Speculatively, load from secret address
12    // → Cache attack!
13    r = data[(int) p];
14  }
15  return r;
16 }
```

Security goal: speculative constant-time

No choice of speculation can make the leakage depend on sensitive information

Spoiler

Counter-measures can be efficiently implemented and automatically verified (A. Shivakumar, Barthe, Grégoire, Laporte, Oliveira, et al. 2022)

Open Questions

- Precisely check SCT at assembly-level?
- Justify that compilation preserves SCT?

Thanks

Get in touch at <https://formosa-crypto.org>

References

- A. Shivakumar, Basavesh, Gilles Barthe, Benjamin Grégoire, Vincent Laporte, Tiago Oliveira, Swarn Priya, Peter Schwabe, and Lucas Tabary-Maujean. 2022. “Typing High-Speed Cryptography Against Spectre V1.” Cryptology ePrint Archive, Paper 2022/1270. <https://eprint.iacr.org/2022/1270>.
- A. Shivakumar, Basavesh, Gilles Barthe, Benjamin Grégoire, Vincent Laporte, and Swarn Priya. 2022. “Enforcing Fine-grained Constant-time Policies.” In *CCS '22: 2022 ACM SIGSAC Conference on Computer and Communications Security*, 83–96. Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (CCS '22). Los Angeles CA, United States: ACM. <https://doi.org/10.1145/3548606.3560689>.
- Barthe, Gilles, Benjamin Grégoire, Vincent Laporte, and Swarn Priya. 2021. “Structured Leakage and Applications to Cryptographic Constant-Time and Cost.” In *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security*, 462–76. CCS '21. New York, NY, USA: ACM. <https://doi.org/10.1145/3460120.3484761>.