

# High-Assurance Cryptography in Jasmin & Spectre Security

Vincent Laporte;  
and the Formosa Crypto team members

2022-09-26, Cambium, Inria, Paris

# Challenges for a (post-quantum) cryptography library

## Ambitious goals

- Execution speed
- Functional correctness
- Safety
- Security against:
  - quantum computers
  - side-channel attacks
  - speculative execution attacks
  - fault attacks

## Everything is broken



Illustration: `crypto/sha/asm/keccak1600-avx2.pl` (OpenSSL)

```
383     vmovdqu    %xmm0, 8+32*4-96($A_flat), $A41
384     vmovdqu    %xmm0, 8+32*5-96($A_flat), $A11
385
386     mov     $bsz, %rax
387
388     .Loop_squeeze_avx2:
389     mov     @A_jagged[$i]-96($A_flat), %r8
390     ____
391     for (my $i=0; $i<25; $i++) {
392     $code.=<<____;
393     sub     \ $8, $len
394     jc     .Ltail_squeeze_avx2
395     mov     %r8, ($out)
396     lea    8($out), $out
397     je     .Ldone_squeeze_avx2
398     dec     %eax
399     je     .Lextend_output_avx2
400     mov     @A_jagged[$i+1]-120($A_flat), %r8
401     ____
402     }
403     $code.=<<____;
404     .Lextend_output_avx2:
405     call    __KeccakF1600
406
407     vmovq    %xmm0, -96($A_flat)
408     vmovdqu  $A01, 8+32*0-96($A_flat)
409     vmovdqu  $A20, 8+32*1-96($A_flat)
410     vmovdqu  $A31, 8+32*2-96($A_flat)
411     vmovdqu  $A21, 8+32*3-96($A_flat)
412     vmovdqu  $A41, 8+32*4-96($A_flat)
413     vmovdqu  $A11, 8+32*5-96($A_flat)
```

## Jasmin

A programming language that enables both:

- crypto practitioners to write optimized implementations
- formal method enthusiasts to verify these implementations

## A tool-box

- Certified compiler: allows reasoning at source level
- Automatic checkers (safety, constant-time)
- EasyCrypt support for semi-automatic verification

## LibJade: work in progress

<https://github.com/formosa-crypto/libjade>

**Aim:** comprehensive library of (post-quantum) cryptography primitives

- efficient
- verified

# Jasmin Hello World

---

```
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 }
```

---

## Semantics judgment, defined in Coq

In program  $p$ ,  
calling function  $f$  with arguments  $\vec{a}$  from initial memory  $m$   
**terminates** in final memory  $m'$  and returns values  $\vec{r}$ :

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

## Automatic Checker, implemented in OCaml

Infers a sufficient precondition  $P$  (for a function  $f$  in program  $p$ ) such that:

$$\forall \vec{a} m, P(\vec{a}, m) \implies \exists \vec{r} m', f : (\vec{a}, m) \Downarrow_p (\vec{r}, m')$$

- polyhedra for numerical arguments
- range and alignment for pointer arguments

# Formal Verification of Jasmin Programs, using EasyCrypt

Jasmin programs are translated into pWhile programs

## For functional correctness

- Using (probabilistic) Hoare logic; or
- by proving program equivalence.

## For semantic security (e.g., IND $\$$ -CPA)

- This is where EasyCrypt shines

## Question to the audience

Which logic can express the correctness of the translation?

## 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”

# Consequences of Compiler Correctness

## Source-level reasoning is **correct**

- Functional properties carry down to the assembly code
- including semantic security

## Limits

- Non-determinism (caveat: next Jasmin release will have `#randombytes`)
- Changing representation of values
- Non-functional properties



# Verification of Constant-Time Security

## Instrumented Semantics

The adversary observes:

- control flow
- memory accesses:

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

## Security property ( $\varphi$ -CT)

Given a “low-equivalence” relation  $\varphi$  between initial states:

$$\forall \vec{a}_1 \ m_1 \ \vec{a}_2 \ m_2 \ \ell_1 \ \ell_2 \ \vec{r}_1 \ \vec{r}_2 \ m'_1 \ m'_2,$$

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

## Automatic checker

Programs can be annotated with security level annotations.

Able to deal with most cases

## Extraction to EasyCrypt

EC has relational program logics

Mostly automatic

## Leakage Transformers

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

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

- Stronger correctness theorem
- Proved for Jasmin v21 only
- Implies preservation of constant-time

# Summary

- Language & Tools & Theorems
- Case studies: Curve25519, ChaCha20/Poly1305, SHA3, Kyber, ...
- Ongoing work:
  - reduce the TCB
  - more target architectures
  - better programming environment
  - LibJade, a comprehensive (post-quantum) crypto library

# Speculative Execution: branch prediction and Spectre v1

- Do not wait
  - the end of an instruction before starting to execute the next one
- Speculate
  - what is the next instruction to execute

## Example

```
⋮  
mov    rax, 0  
cmp    rdi, 2  
jnb    ...  
mov    rax, qword ptr[rcx + rdi * 8]  
mov    rax, qword ptr[rdx + rax * 8]  
⋮
```

Many processors feature branch prediction:



Figure 1: IBM Stretch (by Rama, CC BY-SA)

# Spectre Vulnerabilities (2018) & Speculative Load Hardening (SLH)

---

```
1 u64[2] offsets = { 1, 0 };
2 u64[2] data = { 0xcafe, 0xbeef };
3
4 export fn archetype(#public reg u64 n) → reg u64 {
5     reg u64 p r;
6     reg bool c;
7
8
9     r = 0;
10    c = n <u 2;
11    if c { // Branch trained with in-bounds n
12        // Speculatively, out-of-bounds array access
13        p = offsets[(int) n];
14        // Speculatively, load from secret address → Cache attack!
15        r = data[(int) p];
16    }
17    return r;
18 }
```

---



# Spectre Vulnerabilities (2018) & Speculative Load Hardening (SLH)

---

```
1 u64[2] offsets = { 1, 0 };
2 u64[2] data = { 0xcafe, 0xbeef };
3
4 export fn archetype(#public reg u64 n) → reg u64 {
5     reg u64 p r;
6     reg bool c;
7
8
9     r = 0;
10    c = n <u 2;
11    if c { // Branch trained with in-bounds n
12
13        p = offsets[(int) n];
14        ? #LFENCE; // Secure but inefficient
15        r = data[(int) p];
16    }
17    return r;
18 }
```

---



# Spectre Vulnerabilities (2018) & Speculative Load Hardening (SLH)

---

```
1 u64[2] offsets = { 1, 0 };
2 u64[2] data = { 0xcafe, 0xbeef };
3
4 export fn archetype(#public reg u64 n) → reg u64 {
5     reg u64 p r;
6     reg bool c;
7
8
9     r = 0;
10    c = n <u 2;
11    if c {
12
13        p = offsets[(int) n];
14        p = #protect(p); // Wish: only prevent insecure flows
15        r = data[(int) p];
16    }
17    return r;
18 }
```

---



# Spectre Vulnerabilities (2018) & Speculative Load Hardening (SLH)

---

```
1 u64[2] offsets = { 1, 0 };
2 u64[2] data = { 0xcafe, 0xbeef };
3
4 export fn archetype(#transient reg u64 n) → reg u64 {
5     reg u64 p r;
6     reg bool c;
7     #msf reg u64 mask;    // Is the hardware misspeculating?
8     mask = #init_msf();  // Initial fence
9     r = 0;
10    c = n <u 2;
11    if c {
12        mask = #set_msf(c, mask); // Detect misspeculation
13        p = offsets[(int) n];
14        p = #protect(p, mask);    // Mask loaded value
15        r = data[(int) p];
16    }
17    return r;
18 }
```

---





## Primitives for SLH

Primitive	Semantics	Implementation
$m = \text{init-msf}()$	$m = 0$	fence; $m = 0$
$m = \text{set-msf}(c, m)$	$\text{assert}(c)$	$m = c ? m : -1$
$v = \text{protect}(v, m)$	$\text{assert}(m == 0)$	$v \mid= m$

# Speculative semantics for Jasmin

- The adversary has full control over the speculation through **directives**:

$$d \in \text{Dir} ::= \text{step} \mid \text{force} \mid \text{load } a, i \mid \text{store } a, i$$

- Side-channel leakage is modeled through **observations**:

$$o \in \text{Obs} ::= \bullet \mid \text{read } a, v \mid \text{write } a, v \mid \text{branch } b$$

- Small steps relate states with an explicit mispeculation bit:

$$\frac{}{x := e; c, \langle \rho, \mu, b \rangle \xrightarrow[\text{step}]{\bullet} c, \langle \rho\{x \leftarrow \llbracket e \rrbracket_{\rho}\}, \mu, b \rangle} \text{ASSIGN}$$

# Small Step Speculative Semantics

$$\frac{e = \text{if } (d = \text{force}) \text{ then } \neg \llbracket t \rrbracket_\rho \text{ else } \llbracket t \rrbracket_\rho}{\text{if } t \text{ then } c_\top \text{ else } c_\perp; c, \langle \rho, \mu, b \rangle \xrightarrow[d]{\text{branch}(e)} c_e; c, \langle \rho, \mu, b \vee d = \text{force} \rangle} \text{COND}$$

$$\frac{\llbracket e \rrbracket_\rho = n \in |a| \quad \mu(a, n) = v}{x := a[e]; c, \langle \rho, \mu, b \rangle \xrightarrow[\text{step}]{\text{addr}(n)} c, \langle \rho\{x \leftarrow v\}, \mu, b \rangle} \text{LOAD}$$

$$\frac{\llbracket e \rrbracket_\rho = n \notin |a| \quad m \in |\alpha|, \mu(\alpha, m) = v}{x := a[e]; c, \langle \rho, \mu, \top \rangle \xrightarrow[\text{load}(\alpha, m)]{\text{addr}(n)} c, \langle \rho\{x \leftarrow v\}, \mu, \top \rangle} \text{U-LOAD}$$

## Speculative Constant-Time

Given an equivalence relation  $\varphi$  on states, a program  $c$  is  $\varphi$ -SCT when:

$$\forall D \ s_1 \ s_2 \ O_1 \ O_2 \ \dots, s_1 \ \varphi \ s_2 \ \Longrightarrow \\ c, \langle s_1, \perp \rangle \xrightarrow[D]{O_1} c', \langle s'_1, b'_1 \rangle \Longrightarrow c, \langle s_2, \perp \rangle \xrightarrow[D]{O_2} c', \langle s'_2, b'_2 \rangle \Longrightarrow \\ O_1 = O_2.$$

# Type System with Constraints

## Security Types

- Two security levels:  $L \leq H$
- A security type is a pair  $\tau = (\tau_n, \tau_s)$  of security levels
  - $\tau_n$ : security level under normal executions
  - $\tau_s$ : security level under all executions, including misspeculated ones

## MSF-tracking

Flow-sensitive typing state  $\Sigma$ :

- unknown: *top* state
- ms: register ms accurately models the (semantic) misspeculation bit
- $ms|_e$ : register ms and expression e model the MSF

## Theorem

If a program  $c$  is well-typed:  $\Sigma, \Gamma \vdash c : \Sigma', \Gamma' \mid C$   
then for all valuation  $\theta$  such that  $\theta \models C$ ,  $c$  is  $=_{\theta, \Gamma}^{\Sigma}$ -SCT.

## Selected Rules

$$\frac{\Gamma \vdash e : \tau \mid C}{\Sigma, \Gamma \vdash x = e : \Sigma^x, \Gamma\{x \leftarrow \tau\} \mid C} \text{ASSIGN}$$

$$\frac{\Gamma \vdash b : \sigma \mid C_b \quad \Sigma_{|b}, \Gamma \vdash c_1 : \Sigma_1, \Gamma_1 \mid C_1 \quad \Sigma_{|\neg b}, \Gamma \vdash c_2 : \Sigma_2, \Gamma_2 \mid C_2 \quad \Gamma' \text{fresh}}{\Sigma, \Gamma \vdash \text{if } b \text{ then } c_1 \text{ else } c_2 : \Sigma_1 \cap \Sigma_2, \Gamma' \mid C_b \cup C_1 \cup C_2 \cup \{\sigma \leq L; \Gamma_1 \leq \Gamma'; \Gamma_2 \leq \Gamma'\}} \text{COND}$$

$$\frac{\Gamma \vdash i : \sigma \mid C_i \quad \tau \text{ fresh}}{\Sigma, \Gamma \vdash x = a[i] : \Sigma^x, \Gamma\{x \leftarrow \tau\} \mid C_i \cup \{\sigma \leq L, \Gamma_n(a) \leq \tau_n, H \leq \tau_s\}} \text{LOAD}$$

$$\frac{\Gamma \vdash i : \sigma \mid C_i \quad \tau \text{ fresh}}{\Sigma, \Gamma \vdash \text{ssafe } x = a[i] : \Sigma^x, \Gamma\{x \leftarrow \tau\} \mid C_i \cup \{\sigma \leq L, \Gamma(a) \leq \tau\}} \text{SAFE LOAD}$$

## Selective SLH

A fence at the beginning of each export function

Not all loads need to be protected:

- secret loads are fine
- speculatively safe loads too

## A few implementation tricks

Reordering can help: load (public data) early, store (secret data) late

Can use `MMX` for public spills

When no protect are needed, tracking the `MSF` is not necessary

## Example: ChaCha20

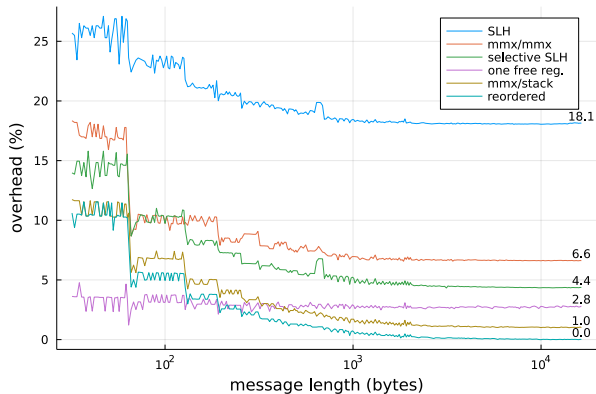


Figure 2: Run-time overhead for various implementations of ChaCha20



## Selective SLH

- One can use the type-checker to insert protect instructions systematically
- Security against Spectre v1 can be achieved at a very little cost and library scale

## Formal Verification

- Reasoning about speculative executions is possible
- Static analysis of optimized low-level programs is feasible

# Conclusion

## Formosa: High-speed, high-assurance cryptography

Fast, safe, correct, secure, ...

## Jasmin: a great environment for research

- Secure & correct compilation
- Language-based security
- ...

## Open questions

- SCT preservation
- Adoption (by crypto practitioners)
- Your questions...