

The Jasmin Workbench for High-Assurance & High-Speed Cryptography

Benjamin GRÉGOIRE



2023-06-28

This Talk

- 1 Why & What is Jasmin
- 2 A few Case Studies
- 3 Side channels
- 4 We are not done yet!

Why & What is Jasmin

Challenges for a (post-quantum) cryptography library

Ambitious goals

- Security proof (algorithm, quantum adversaries, EasyCrypt)
- Execution speed
- Functional correctness
- Safety
- Security against:
 - side-channel attacks
 - speculative execution attacks
- All those guaranties should be provided at the assembly level

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

```
382         vmovdqu    %xmm0, -96(%r11),%xmm1
383         vmovdqu    8+32*4-96(%r11),%xmm1
384         vmovdqu    8+32*5-96(%r11),%xmm1
385
386         mov     $bsz,%rax
387
388     .Loop_squeeze_avx2:
389         mov     @A_jagged[$i]-96(%r11),%r8
390     ___
391     for (my $i=0; $i<25; $i++) {
392     $code.=<<___;
393         sub     \S8,$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(%r11),%r8
401     ___
402     }
403     $code.=<<___;
404     .Lextend_output_avx2:
405         call   __KeccakF1600
406
407         vmovq   %xmm0, -96(%r11)
408         vmovdqu %xmm1, 8+32*0-96(%r11)
409         vmovdqu %xmm1, 8+32*1-96(%r11)
410         vmovdqu %xmm1, 8+32*2-96(%r11)
411         vmovdqu %xmm1, 8+32*3-96(%r11)
412         vmovdqu %xmm1, 8+32*4-96(%r11)
413         vmovdqu %xmm1, 8+32*5-96(%r11)
```

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

An illustrative Jasmin program

```
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

Compiler Correctness (Coq)

Semantics Preservation (forward simulation)

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

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

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-functional properties

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

For implementation security (e.g., constant-time)

- Using relational Hoare logic
- on an instrumented program with explicit information leakage.

A few Case Studies

Fast Implementations

Two example of implementations in Jasmin, as fast as the fastest available implementations:

Curve25519

[CCS17]

- Scalar multiplication on a standard elliptic curve
- Verified for safety and constant-time security

Chacha20/Poly1305

[SP20]

- Authenticated encryption scheme
- Verified for safety and constant-time security
- With formal proofs of functional correctness

Secure High-Assurance Implementations of SHA-3

[CCS19]

- Fast (optimized for AVX2)
- Secure (constant-time)
- Correct (wrt. a reference implementation)

Indifferentiability proof of the Sponge construction

- Main theorem about security of SHA-3
- Bounds the probability for an adversary to break it:
 - in particular to find collisions, preimages, or second preimages
- Theorem applies to the optimized implementation!

Post Quantum implementations

Kyber

[CHES23]

- KEM
- Reference + avx2 implementations
- Verified for safety, functional correctness (hard part)

Dilithium

[CRYPTO23]

- Signature Scheme
- Reference implementation
- Formal security proof in EasyCrypt

Side channels

Implementation Security

Adversaries may observe the machine running a victim program.

Is any sensitive information leaked into these observations?

Constant-Time

- A popular mitigation against timing (cache-based) side-channel attacks
- Two rules
 - No branching on secret data
 - No memory access at secret addresses
- Can be checked using a taint analysis
 - Propagate from entry-points “security labels” (low/high)

Preservation of Constant-Time (Swarn Priya Thesis) [CSF18, POPL19, CCS21]

- Source is CT implies target is CT

Fine-Grained Leakage Models for Constant-Time

Variation of the model

[CCS22]

The base-line constant-time model is too coarse in practice:

- some arithmetic operations leak a function of their arguments (DIV/MOD)
- leaking only the cache line
- Bug found in OpenSSL, patch accepted.

Problem with rejection sampling

- Kyber, Dilithium and Falcon are all based on rejection sampling
- This is not constant time, but it can be (approximate) probabilistic constant time

A serious vulnerability

- Leakage of sensitive information due to speculative execution (branch prediction)
- Not a hardware bug



Efficiently mitigated

[SP23]

- Manually protect (against V1) the whole LibJade library (using SSLH)
- Automatically prove the result secure (type system)
- Experimentally assess that the protection cost is *low*
- Don't know how to prove preservation !!!

We are not done yet!

- Still hard to program in Jasmin
 - Lacking documentation
 - Register allocation error message are not helpful
 - Detailed knowledge of the target architecture is often needed
- Lack of modularity
 - No separate compilation
 - Existing libraries are difficult to reuse
- Checking safety of Kyber decapsulation takes 16 CPU-hours
- Functional correctness hard to establish

Target architectures

- Currently: x86_64
- Soon: ARMv7
- Later: ARMv8, RISC-V, Open-titan?

Open question

How to build & maintain a comprehensive, multi-platform, verified library ?

- More than Spectre V1, zeroing the stack and register ...
- Correctness proofs too hard (link with Cryptoline?)
- Post-quantum security proof (EasyPQC, [CCS21])
- Falcon needs floating point (dealing with error) and new notion (Rényi divergence)

Questions?

Thanks