

High-Assurance and High-Speed Cryptographic Implementations Using the Jasmin Language

J.B. Almeida, M. Barbosa, G. Barthe, B. Grégoire, **A. Koutsos**, V. Laporte, T. Oliveira, P-Y. Strub

Octobre 9th, 2019

Context

The screenshot shows a web browser displaying the Wikipedia Main Page. On the left, the Wikipedia logo and navigation menu are visible. The main content area is partially obscured by a 'Page Info' window. The 'Page Info' window has tabs for 'General', 'Media', 'Permissions', and 'Security', with 'Security' selected. The 'Security' tab shows the following information:

- Website Identity**
 - Website: en.wikipedia.org
 - Owner: This website does not supply ownership information.
 - Verified by: GlobalSign nv-sa [View Certificate](#)
 - Expires on: November 22, 2019
- Privacy & History**
 - Have I visited this website prior to today? Yes, 1,762 times
 - Is this website storing information on my computer? Yes, cookies [Clear Cookies and Site Data](#)
 - Have I saved any passwords for this website? No [View Saved Passwords](#)
- Technical Details**
 - Connection Encrypted (TLS_ECDHE_ECDSA_WITH_CHACHA20_POLY1305_SHA256, 256 bit keys, TLS 1.2)
 - The page you are viewing was encrypted before being transmitted over the Internet.
 - Encryption makes it difficult for unauthorized people to view information traveling between computers. It is therefore unlikely that anyone read this page as it traveled across the network.
 - [Help](#)

Cryptographic Libraries

Developing cryptographic libraries is hard, as the code must be:

- **efficient:** pervasive usage, on large amount of data.
- **functionally correct:** the specification must be respected.
- **protected against side-channel attacks:** constant-time implementation.

Side-Channel Attacks

Exploit **auxilliary information** to break a cryptographic primitive.

Side-Channel Attacks

Exploit **auxilliary information** to break a cryptographic primitive.

Constant-Time Programming

- Countermeasure against **timing** and **cache** attacks.
- **Control-flow** and **memory accesses** should not depend on **secret** data.
- Crypto implementations without this property are vulnerable.

Constraints

- Efficiency: **low-level** operations and **vectorized** instructions.
- Functional Correctness: **readable** code, with **high-level abstractions**.
- Side-Channel Attacks Protection: **control** over the executed code.

Gap Between Source and Assembly

Source

- High-level abstractions.
- Readable code.

Gap Between Source and Assembly

Source

- High-level abstractions.
- Readable code.

Source is not Security/Efficiency Friendly

- Trust compiler (GCC or Clang).
- Certified compilers are less efficient (CompCert).
- Optimizing can break side channel resistance.

Preservation of Constant-Timeness?

Before

```
int cmove(int x, int y, bool b) {  
    return x + (y-x) * b;  
}
```

Preservation of Constant-Timeness?

Before

```
int cmove(int x, int y, bool b) {  
    return x + (y-x) * b;  
}
```

After

```
int cmove(int x, int y, bool b) {  
    if (b) {  
        return y;  
    } else {  
        return x;  
    }  
}
```

Gap Between Source and Assembly

Assembly

- **Efficient** code.
- **Control** over the program execution.

Gap Between Source and Assembly

Assembly

- **Efficient** code.
- **Control** over the program execution.

Assembly is not Programmer/Verifier Friendly

- The code is obfuscated.
- More error prone.
- Harder to prove/analyze.

Fast and Formally Verified Assembly Code

- **Source language:** assembly in the head with formal semantics
⇒ programmer & verification friendly
- **Compiler:** predictable & formally verified (in Coq)
⇒ programmer has control and no compiler security bug
- **Verification tool-chain:**
 - Functional correctness.
 - Side-channel resistance (constant-time).
 - Safety.

Implementations in Jasmin

TLS 1.3 components : ChaCha20, Poly1305, Curve25519.

The Jasmin Language

Initialization of ChaCha20 State

```
inline fn init(reg u64 key nonce, reg u32 counter) → stack u32[16] {
  inline int i;
  stack u32[16] st;
  reg u32[8] k;
  reg u32[3] n;

  st[0] = 0x61707865;
  st[1] = 0x3320646e;
  st[2] = 0x79622d32;
  st[3] = 0x6b206574;

  for i=0 to 8 {
    k[i] = (u32)[key + 4*i];
    st[4+i] = k[i];
  }

  st[12] = counter;

  for i=0 to 3 {
    n[i] = (u32)[nonce + 4*i];
    st[13+i] = n[i];
  }

  return st;
}
```

Zero-Cost Abstractions

- Variable names.
- Arrays.
- Loops.
- Inline functions.

User Control: Loop Unrolling

```
for i=0 to 15 {  
  k[i] = st[i];  
}
```

For Loops

- Fully unrolled.
- The value of the counter is propagated.
- The source code still readable and compact.

```
while(i < 15) {  
  k[i] = st[i]; i += 1;  
}
```

While Loops

- Untouched.

User Control: Register or Stack

- Jasmin has three kinds of variables:
 - register variables (`reg`).
 - stack variables (`stack`).
 - global variables (`global`).
- Arrays can be register arrays or stack arrays.
- Spilling is done manually (by the user).

```
inline fn sum_states(reg u32[16] k, stack u32 k15, stack u32[16] st) → reg u32[16], stack u32
{
  inline int i;
  stack u32 k14;

  for i=0 to 15 {
    k[i] += st[i];
  }

  k14 = k[14]; k[15] = k15; // Spilling
  k[15] += st[15];
  k15 = k[15]; k[14] = k14; // Spilling

  return k, k15;
}
```

User Control: Instruction-Set

- Direct memory access.

```
reg u64 output, plain;
```

```
for i=0 to 12 {  
    k[i] = (u32)[plain + 4*i];  
    (u32)[output + 4*i] = k[i]; }  
}
```

- The carry flag is an ordinary boolean variable.

```
reg u64[3] h;  
reg bool cf0 cf1;  
reg u64 h2rx4 h2r;
```

```
h2r    += h2rx4;  
cf0, h[0] += h2r;  
cf1, h[1] += 0 + cf0;  
_ , h[2] += 0 + cf1;
```

- Most assembly instructions are available.

```
of, cf, sf, pf, zf, z = x86_ADC(x, y, cf);
```

```
of, cf, x = x86_ROL_32(x, bits);
```

- Vectorized instructions (SIMD).

```
k[0] +8u32= k[1]; // vectorized addition of 8 32-bits words;
```

```
k[1] = x86_VPSHUFD_256(k[1], (4u2)[0,3,2,1]);
```

The Jasmin Compiler

Goals And Features

- Predictability and control of generated assembly.
- Preserves semantics (machine-checked in Coq).
- Preserves side-channel resistance

Passes and Optimizations

- For loop unrolling.
- Function inlining.
- Constant-propagation.
- Sharing of stack variables.
- Register array expansion.
- Lowering.
- Register allocation.
- Linearisation.
- Assembly generation.

Compilation Theorem (Coq)

$$\begin{aligned} &\forall p, p'. \text{compile}(p) = \text{ok}(p') \Rightarrow \\ &\forall v_a, m, v_r, m'. \text{enough-stack-space}(p', m) \Rightarrow \\ &v_a, m \Downarrow^p v_r, m' \Rightarrow v_a, m \Downarrow^{p'} v_r, m' \end{aligned}$$

Remarks

- The compiler uses validation.
- We may need some extra memory space for p' :
 $\text{enough-stack-space}(p', m)$
- If p is not safe, i.e. $v_a, m \Downarrow^p \perp$, then we have no guarantees.

Functional Correctness

Methodology

- We start from a **readable reference implementation**:
 - Using a mathematical specification (e.g. in $\mathbb{Z}/p\mathbb{Z}$).
 - Or a simple imperative specifications.
- We gradually transform the **reference implem.** into an **optimized implem.**:
 - We prove that each transformation **preserves functional correctness** by equivalence (game-hoping).
- We prove additional properties of the final implementation:
 - **Constant-time** by program equivalence.
 - **Safety** by static analysis.

Gradual Transformation

We perform functional correctness proofs by game hopping:

$$C_{\text{ref}} \sim C_1 \sim \dots \sim C_n \sim C_{\text{opt}}$$

EasyCrypt

- Jasmin programs are translated into EasyCrypt programs.
- EasyCrypt model for Jasmin (memory model + instructions).
- Equivalences are proved in EasyCrypt.

Relational Hoare Logic

A judgment $\{P\} c_1 \sim c_2 \{Q\}$ is valid if:

$$(m_1, m_2) \in P \Rightarrow m_1 \Downarrow^{c_1} m'_1 \Rightarrow m_2 \Downarrow^{c_2} m'_2 \Rightarrow (m'_1, m'_2) \in Q$$

Relational Hoare Logic is provided in EasyCrypt.

Example

- c_1 is the reference implementation (the specification)
- c_2 is the optimized implementation

$$\{\text{args}\langle m_1 \rangle = \text{args}\langle m_2 \rangle\} c_1 \sim c_2 \{\text{res}\langle m_1 \rangle = \text{res}\langle m_2 \rangle\}$$

Example: ChaCha20

Stream cipher that iterates a *body* on all the blocks of a message.

Reference

```
while (i < len) {  
  chacha_body;  
  i += 1;  
}
```

Loop tiling

```
while (i + 4 ≤ len) {  
  chacha_body;  
  chacha_body;  
  chacha_body;  
  chacha_body;  
  i += 4;  
}  
chacha_end
```

Scheduling

```
while (i + 4 ≤ len) {  
  chacha_body4_swapped;  
  i += 4;  
}  
chacha_end
```

Vectorization

```
while (i + 4 ≤ len) {  
  chacha_body4_vectorized;  
  i += 4;  
}  
chacha_end
```

Safety

Definition

A program p is safe under precondition ϕ if and only if:

$$\forall (v, m) \in \phi. v, m \not\Downarrow^p \perp$$

Why do we Need Safety?

- If p is safe, its execution never crashes.
- The compilation theorem gives no guarantees if p is not safe.
- Jasmin semantics in Easycrypt assumes that p is safe.


Properties to Check

- Division by zero.
- Variable and array initialization.
- Out-of-bound array access.
- Termination.
- Valid memory access.


Jasmin

Safety is checked automatically by **static analysis**.

Abstract Interpretation: Abstract Values




Abstract Interpretation: Abstract Values



Soundness

$X^\#$ over-approximates X if and only if $X \subseteq \gamma(X^\#)$

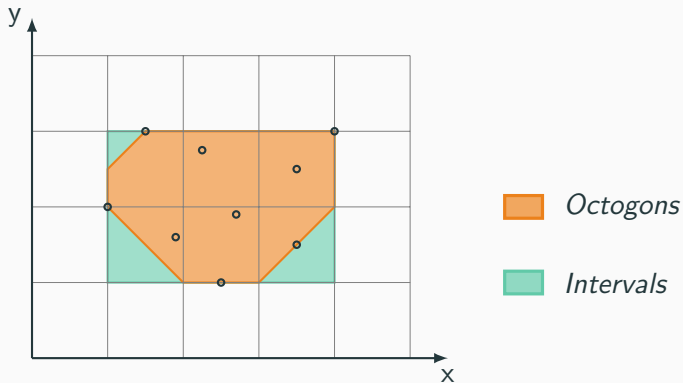
Abstract Interpretation: Abstract Values



Soundness

$X^\#$ over-approximates X if and only if $X \subseteq \gamma(X^\#)$


Abstract Interpretation: Abstract Values



Soundness

$X^\#$ over-approximates X if and only if $X \subseteq \gamma(X^\#)$


Abstract Interpretation: Abstract Values



Soundness

$X^\#$ over-approximates X if and only if $X \subseteq \gamma(X^\#)$

Abstract Interpretation: Abstract Transformers




Soundness

$f^\#$ over-approximates f if and only if:

$$\forall X^\#. f \circ \gamma(X^\#) \subseteq \gamma \circ f^\#(X^\#)$$

Abstract Interpretation: Abstract Transformers




Soundness

$f^\#$ over-approximates f if and only if:

$$\forall X^\#. f \circ \gamma(X^\#) \subseteq \gamma \circ f^\#(X^\#)$$

Abstract Interpretation: Abstract Transformers



Soundness

$f^\#$ over-approximates f if and only if:

$$\forall X^\#. f \circ \gamma(X^\#) \subseteq \gamma \circ f^\#(X^\#)$$

Features of the Language

Jasmin is a simple language for static analysis:

- No recursion.
- Arrays size are statically known.
- No dynamic memory allocation.

Example

```
fn load(reg u64 in, reg u64 len) {  
  inline int i;  
  reg u8 tmp;  
  
  tmp = 0;  
  while (len >= 16) {  
    for i = 0 to 16 {  
      tmp = (u8)[in + i]; }  
    in += 16;  
    len -= 16; }  
  
  for i = 0 to 16 {  
    if i < len {  
      tmp = (u8)[in + i]; }}  
  
  return tmp;  
}
```

Example

```
fn load(reg u64 in, reg u64 len) {  
  inline int i;  
  reg u8 tmp;  
  
  tmp = 0;  
  while (len >= 16) {  
    for i = 0 to 16 {  
      tmp = (u8)[in + i]; }  
    in += 16;  
    len -= 16; }  
  
  for i = 0 to 16 {  
    if i < len {  
      tmp = (u8)[in + i]; }}  
  
  return tmp;  
}
```

Memory Calling Contract

$$\text{valid-mem}_{\text{load}}(\text{in}_0, \text{len}_0) =$$
$$[\text{in}_0; \text{in}_0 + \text{len}_0]$$

Variables in the Abstract Domain

Let \mathcal{P} be a set of pointers. To a variable $x \in \mathcal{V}$, we associate:

- $x \in \mathcal{V}^\#$: its abstract value.
- $x_0 \in \mathcal{V}^\#$: its abstract initial value.
- $\text{pt}_x \subseteq \mathcal{P}$: points-to information.
- $\text{offset}_x \in \mathcal{V}^\#$: its abstract offset.

Variables in the Abstract Domain

Let \mathcal{P} be a set of pointers. To a variable $x \in \mathcal{V}$, we associate:

- $x \in \mathcal{V}^\#$: its abstract value.
- $x_0 \in \mathcal{V}^\#$: its abstract initial value.
- $\text{pt}_x \subseteq \mathcal{P}$: points-to information.
- $\text{offset}_x \in \mathcal{V}^\#$: its abstract offset.

Moreover, for every $p \in \mathcal{P}$, we have:

- $\text{mem}_p \in \mathcal{V}^\#$: memory accesses at p (plus an offset).

Concretization Function

We decompose x into a base pointer p and an offset offset_x :

$$\gamma(\text{pt}_x = \{p\} \wedge \text{offset}_x = \mathcal{S}^\#) = x \mapsto \{p + o \mid o \in \gamma(\mathcal{S}^\#)\}$$

Concretization Function

We decompose x into a base pointer p and an offset offset_x :

$$\gamma(\text{pt}_x = \{p\} \wedge \text{offset}_x = \mathcal{S}^\#) = x \mapsto \{p + o \mid o \in \gamma(\mathcal{S}^\#)\}$$

Example

- $\gamma(\text{pt}_x = \{p\} \wedge \text{offset}_x = [32; 63]) = x \mapsto [p + 32; p + 63]$

Concretization Function

We decompose x into a base pointer p and an offset offset_x :

$$\gamma(\text{pt}_x = \{p\} \wedge \text{offset}_x = \mathcal{S}^\#) = x \mapsto \{p + o \mid o \in \gamma(\mathcal{S}^\#)\}$$

Example

- $\gamma(\text{pt}_x = \{p\} \wedge \text{offset}_x = [32; 63]) = x \mapsto [p + 32; p + 63]$
- Abstract transformer:
 - $\mathcal{S}^\# : \text{pt}_x = \{p\} \wedge \text{offset}_x = [32; 63]$
 $y \leftarrow x + 16$
 - $\mathcal{S}'^\# : \text{pt}_y = \{p\} \wedge \text{offset}_y = [48; 79]$

Remark

- In $y \leftarrow x + z$, we can either use x or z as a base pointer.
- In practice, it is never a problem (assembly coding style).

Memory Calling Contract

Let f be a procedure with pointers \mathcal{P} . If:

$$\llbracket f \rrbracket^\#(\mathcal{S}_{\text{init}}^\#) \doteq \bigwedge_{p \in \mathcal{P}} \text{mem}_p = \mathcal{S}_p^\# \wedge \dots$$

Then for every $\mathcal{S}_{\text{init}} \subseteq \gamma(\mathcal{S}_{\text{init}}^\#)$:

$$\text{valid-mem}_f(\mathcal{S}_{\text{init}}) \subseteq \bigcup_{p \in \mathcal{P}} \gamma(\mathcal{S}_p^\#)$$

Example

- $S^\#$: $pt_x = \{p\} \wedge mem_p = [0; 127] \wedge offset_x = [128; 128 + 16]$
 $tmp \leftarrow (u8)[x + 16]$
- $S'^\#$: $mem_p = [0; 127] \cup^\# [128; 128 + 32] = [0; 160]$

Example

```
fn load(reg u64 in, reg u64 len) {  
  inline int i;  
  reg u8 tmp;  
  
  tmp = 0;  
  while (len >= 16) {  
    for i = 0 to 16 {  
      tmp = (u8)[in + i]; }  
    in += 16;  
    len -= 16; }  
  
  for i = 0 to 16 {  
    if i < len {  
      tmp = (u8)[in + i]; }}  
  
  return tmp;  
}
```

After the While Loop

$$\begin{aligned} & 0 \leq \text{offset}_{\text{in}}, \text{len}, \text{len}_0, \text{mem}_{\text{in}} \\ & \wedge \text{offset}_{\text{in}} + \text{len} = \text{len}_0 \\ & \wedge \text{len}_0 - 15 \leq \text{offset}_{\text{in}} \leq \text{len}_0 \\ & \wedge \text{mem}_{\text{in}} \leq \text{offset}_{\text{in}} \end{aligned}$$

At the End

$$0 \leq \text{mem}_{\text{in}} \leq \text{len}_0$$

The Analyzer

- Intervals + Relational domain (polyhedra).
- Basic syntactic pre-analysis.
- Disjunctive domain (using the control flow).
- Simple non-relational boolean abstractions (for bools and initialization).
- Brutal handling of function calls.

Result

For Poly1305, with signature:

```
export fn poly1305_avx2(reg u64 out, reg u64 in, reg u64 len, reg u64 k)
```

We infer the ranges:

$\text{mem}_{\text{out}} : \text{out} + [0; 16[$

$\text{mem}_{\text{len}} : \emptyset$

$\text{mem}_k : k + [0; 32[$

$\text{mem}_{\text{in}} : \text{in} + [0; \text{len}[$

Caveat

We manually provide some information to the analyser:

- pointers (input) variables: k, in and out in Poly1305.
- relational (input) variables: len in Poly1305.

Conclusion

Contributions

A framework to build high-speed certified implementations of cryptographic primitives.

- Code is manually optimized.
- Functional correctness is obtained by game hopping.
- Safety and security against timing attacks are proved automatically.
- Efficient implementation of Poly1305, ChaCha20 and Gimli.

Future Works

- More TLS 1.3 primitives.
- More architectures, more general purpose language.
 - procedure calls.
 - register allocation/spilling.
- Certification for safety proofs.