



The Last Yard: Foundational End-to-End Verification of High-Speed Cryptography

Philipp G. Haselwarter*

Aarhus University
Denmark
philipp@haselwarter.org

Théo Winterhalter

Inria
France
theo.winterhalter@inria.fr

Benjamin Salling Hvass*

Aarhus University
Denmark
bsh@cs.au.dk

Cătălin Hrițcu

MPI-SP
Germany
catalin.hritcu@mpi-sp.org

Lasse Letager Hansen*

Aarhus University
Denmark
letager@cs.au.dk

Bas Spitters

Aarhus University
Denmark
spitters@cs.au.dk

Abstract

The field of high-assurance cryptography is quickly maturing, yet a unified foundational framework for end-to-end formal verification of efficient cryptographic implementations is still missing. To address this gap, we use the Coq proof assistant to formally connect three existing tools: (1) the Hacspect emergent cryptographic specification language; (2) the Jasmin language for efficient, high-assurance cryptographic implementations; and (3) the SSProve foundational verification framework for modular cryptographic proofs. We first connect Hacspect with SSProve by devising a new translation from Hacspect specifications to imperative SSProve code. We validate this translation by considering a second, more standard translation from Hacspect to purely functional Coq code and generate a proof of the equivalence between the code produced by the two translations. We further define a translation from Jasmin to SSProve, which allows us to formally reason in SSProve about efficient cryptographic implementations in Jasmin. We prove this translation correct in Coq with respect to Jasmin’s operational semantics. Finally, we demonstrate the usefulness of our approach by giving a foundational end-to-end Coq proof of an efficient AES implementation. For this case study, we start from an existing Jasmin implementation of AES that makes use of hardware acceleration and prove that it conforms to a specification of the AES standard written in Hacspect. We use SSProve to formalize the security of the encryption scheme based on the Jasmin implementation of AES.

*Equal Contribution.



This work is licensed under a Creative Commons Attribution 4.0 International License.

CPP ’24, January 15–16, 2024, London, UK

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0488-8/24/01

<https://doi.org/10.1145/3636501.3636961>

CCS Concepts: • Theory of computation → Program verification; Program specifications; • Security and privacy → Symmetric cryptography and hash functions; Logic and verification;

Keywords: high-assurance cryptography, formal verification, computer-aided cryptography, AES, Coq

ACM Reference Format:

Philipp G. Haselwarter, Benjamin Salling Hvass, Lasse Letager Hansen, Théo Winterhalter, Cătălin Hrițcu, and Bas Spitters. 2024. The Last Yard: Foundational End-to-End Verification of High-Speed Cryptography. In *Proceedings of the 13th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP ’24)*, January 15–16, 2024, London, UK. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3636501.3636961>

1 Introduction

Research on high-assurance cryptography recently led to significant practical success, with formally verified cryptographic code making its way into mainstream libraries and software products [7, 14, 16, 19, 21, 34, 37, 41, 42]. Since in this area missing any bugs can have a serious security impact, some additionally try to reduce the trusted computing base of their verification tools for cryptographic code and construct foundational proofs [5, 14, 21, 25, 29, 32]. Such foundational proofs rely on strong logical foundations—usually by working in a proof assistant like Coq or Isabelle/HOL—and only on standard, clearly stated assumptions. Yet despite good progress in this direction, a couple of important gaps remain for foundational end-to-end cryptographic verification.

First, there is a specification gap. Currently, cryptographic primitives and protocols are specified only using informal pseudo-code in the standards (e.g., in IETF RFCs). The Hacspect language [15, 31] aims to improve this, by making the code of these cryptographic specifications executable, which allows them to also serve as reference implementations that can be used as oracles for testing more efficient implementations. Hacspect is a simple subset of the Rust programming language, which aims to be understandable for both ordinary developers and cryptographers. Hacspect can be translated




Figure 1. Proposed workflow for foundational end-to-end verification of high-speed cryptography

to the typed, purely functional language of proof assistants such as Coq, EasyCrypt, or F^{*}, which allows sharing cryptographic specifications across these proof assistants.

Such translations from Hacspec to a proof assistant produce a functional specification that can be used for verifying cryptographic code. In such a verification one often starts by proving the equivalence of the functional specification with an imperative specification, which is closer to the code of an implementation to be verified [5]. We automate this step by devising a new translation from Hacspec to imperative programs in SSProve, which is a recent foundational verification framework for modular cryptographic proofs in Coq [1, 25]. Moreover, we provide translation validation infrastructure for automatically proving the equivalence of the code produced by these two translations.

Second, there is an implementation gap. Implementing cryptography in C has pitfalls: (1) unverified C compilers cannot be trusted to be always correct and secure [39], and (2) the CompCert verified C compiler does not perform aggressive optimizations and generates code with efficiency comparable only to GCC at optimization level 1 [2, 28]. Moreover, even aggressively optimized C programs are sometimes not fast enough since they cannot make use of special instructions providing hardware acceleration for cryptographic primitives (e.g., Intel AES-NI [23]). So cryptographic primitives are often implemented directly in assembly, at the cost of loss of abstraction, clarity, and convenience. The Jasmin language [4] was proposed as a solution to this problem. It is a language for implementing cryptographic primitives combining structured control flow with assembly instructions, which allows one to produce efficient code for x86 and ARM. Moreover, the Jasmin compiler comes with Coq proofs that

it preserves the semantics of the source code [4, 5] and that it does not introduce timing side-channel attacks [6].

In the fundamental ‘Last Mile’ paper [5], Jasmin programs are given semantics in Coq and compiled with a compiler verified in Coq, but reasoning about the security and correctness of Jasmin programs is done only after an *unverified* translation to EasyCrypt. In this paper, we close this gap by providing a *verified* translation from Jasmin to SSProve. Staying in Coq not only allows us to reduce the trusted computing base, but it also facilitates reusing existing mathematical Coq libraries [3, 30] to verify Jasmin implementations.

Contributions. We formally connect three existing tools, Hacspec, Jasmin, and SSProve, into a unified foundational Coq framework for the end-to-end verification of high-speed cryptography (Figure 1). This includes the following novel contributions:

- We devise a new translation from Hacspec specifications to imperative SSProve code. In contrast to the existing functional translations, it allows us to reason about the *stateful* behavior of Hacspec.
- We provide a translation validation infrastructure, which automatically produces Coq proofs of program equivalence between the results of this imperative translation and those of a more standard functional translation. We do this by performing a *compositional* symbolic evaluation, relating imperative code to its mathematical model.
- We connect the Jasmin language and verified compiler to SSProve, by providing a translation of Jasmin source code to SSProve. We overcome the challenge created by the fact that SSProve only supports global state while Jasmin programs can use local state.

- We give a mechanized proof in Coq that this translation from Jasmin to SSProve preserves Jasmin’s operational semantics.
- We demonstrate the usefulness of our approach on a case study by producing a foundational end-to-end Coq proof of an efficient AES implementation. We start from an existing Jasmin implementation of AES using the Intel AES-NI instructions for hardware acceleration [23] and prove (in ~2500 lines of Coq code) that it conforms to a Hacspect specification of the AES standard [20]. Finally, we instantiate a PRF-based symmetric encryption scheme with the implementation of AES, and use SSProve to prove IND-CPA security of this scheme under the (standard) assumption that AES is a pseudo-random function (PRF).

Outline. We start by giving an overview of our methodology and illustrating it on a very simple one-time pad example (Section 2). We then discuss necessary background (Section 3), before diving in the two formal connections we establish: the one between Hacspect and SSProve (Section 4), the other between Jasmin and SSProve (Section 5). We finally present the AES case study (Section 6), before discussing related (Section 7) and future work (Section 8).

2 Foundational End-to-End Verification, from Specification to Efficient Implementation

In this section, we first give an overview of our methodology following Figure 1 and then demonstrate its workings on the very simple example of a one-time pad. At a high level, we provide a foundational framework for proving the equivalence between a specification in Hacspect and an efficient, low-level implementation in Jasmin, by translating both to imperative SSProve programs. Once translated, we relate the programs and prove properties about them in Coq using SSProve’s probabilistic relational Hoare logic.

2.1 Workflow

The workflow is illustrated in Figure 1. Starting from an informal description, such as an official *standard* (e.g., published by NIST or IETF) for a cryptographic primitive or protocol, one uses a subset of Rust with a simple, well-defined semantics to develop a *Hacspect specification*.¹ We then automatically translate this specification in two ways:

- once to the purely functional language of Coq; this translation produces a *functional specification*; and
- once to the imperative language of SSProve; this translation produces an *SSProve specification*.

The functional translation [35] targets Coq’s mathematical language, and is similar to the usual functional semantics

of Hacspect in F^{*} and EasyCrypt. The imperative translation serves as a stepping stone towards a Jasmin implementation, which is inherently imperative.

We then perform translation validation [33] to automatically construct an *equivalence proof* in Coq, which formally shows that the functional and imperative Hacspect translations produce equivalent SSProve code from a given Hacspect program. More specifically, we prove that in a clean state, the program produced by the imperative translation will return the same value as the one produced by the functional translation. The proof is conducted in SSProve’s relational Hoare logic (see Section 3.3.4).

The second part of our framework concerns efficient cryptographic implementations written in Jasmin. We implemented a translation from Jasmin to the imperative language of SSProve and proved that it preserves semantics. This proof is entirely mechanized in Coq, which is possible because both SSProve and Jasmin already have formal semantics in Coq [4–6, 25]. So from the same *Jasmin implementation* (1) we can produce an *assembly implementation* using the existing Jasmin compiler, which was proved in Coq to preserve the source language semantics [4, 25]; and (2) we can obtain the *Jasmin Coq AST* of the Jasmin implementation, which we then translate to an *SSProve implementation* in a way that we proved to preserve semantics.

We are now in a position to reason about the SSProve implementation using the relational probabilistic Hoare logic of SSProve. On the one hand, we can conduct an *equivalence proof* between the SSProve implementation obtained from Jasmin and the SSProve specification obtained from Hacspect. On the other hand, we can connect the translated Hacspect specification with *security proofs* done in the SSProve framework. These proofs use the standard security games from the cryptographic literature [13, 24, 36, 38].

Formal Guarantees. By combining the correctness theorems of the Jasmin compiler and our translation to SSProve, we get the following corollary: for any function in a Jasmin program with well-defined semantics, there exists a corresponding compiled assembly function and translated SSProve function with the same semantics, i.e., which maps equal arguments to equal results and which modifies memory in an equivalent manner. In particular, the semantics of the SSProve and assembly functions agree and we can prove the properties of the assembly program by analyzing the corresponding SSProve program; probabilistic properties cannot however be carried to the assembly level, since the semantics there are deterministic. Note that we inherit some assumptions from the compiler proof (e.g., assuming sufficient stack-space) and introduce some in the translation proof (e.g., functions cannot use while-loops), see also Section 5.

¹In fact, Hacspect is directly used in the upcoming hash-to-curve IETF standard [22] for writing a [reference implementation](#).

2.2 One-time Pad Example

We now illustrate this methodology using a very simple example: We construct a one-time pad (OTP) from exclusive or (XOR). This toy example should convey intuition on the methodology. A more interesting case study for the AES encryption scheme is presented in Section 6. This section is to get an idea of the workflow and the ideas, however, we will introduce the background theory in more detail in Section 3.

2.2.1 Specification. The Hacspecc specification for `xor` takes two 64-bit words as input, puts them into mutable variables, and computes their XOR ($\wedge$ in Hacspecc). The result is stored in a mutable variable², which is then returned.

```
fn xor(w1 : u64, w2 : u64) -> u64 {
  let mut x : u64 = w1;
  let mut y : u64 = w2;
  let mut r : u64 = x ^ y;
  r
}
```

Our framework produces an automatic translation of this code to the following Coq function of type `both`.

```
Definition hacspecc_xor (w1 : int64) (w2 : int64) :=
  letbm x_0 : int64 loc( x_0_loc ) := w1 in
  letbm y_1 : int64 loc( y_1_loc ) := w2 in
  letbm r_2 : int64 loc( r_2_loc ) := x_0 .^ y_1 in
  r_2.
```

Here `letbm` stands for “let bind mutable”. The type `both` can be projected both to pure Coq and to SSProve code (see Section 4.3), resulting in the following two functions:

```
Definition hacspecc_xor_pure x y := x .^ y.
```

```
Definition hacspecc_xor_state (x y : int64) :=
  put x_loc := x ;;
  temp_x ← get x_loc ;;
  put y_loc := y ;;
  temp_y ← get y_loc ;;
  put r_loc := int_xor temp_x temp_y ;;
  temp_r ← get r_loc ;;
  ret temp_r.
```

For achieving translation validation, the `both` type also carries an equivalence proof between these two functions:

```
∀ x y, ⊢ { λ '(h0, h1), ⊤ }
  hacspecc_xor_state x y ≈
  ret (hacspecc_xor_pure x y)
  { λ '(v0, h0) '(v1, h1), v0 = v1 }.
```

2.2.2 Jasmin Implementation. A Jasmin implementation of `xor` could look as follows.

```
export fn xor(reg u64 x, reg u64 y) -> reg u64 {
  reg u64 r;
  r = x;
  r ^= y;
  return r;
}
```

It takes two register-allocated arguments `x` and `y` (as indicated by the `reg` keyword) and writes the XOR of `x` and `y` into the return register `r`.

2.2.3 SSProve Implementation. The next step is to translate the Jasmin code to the following SSProve function.

```
Definition JXOR id0 w1 w2 :=
  put x := w1 ;;
  put y := w2 ;;
  put r := w1 ⊕ w2 ;;
  r1 ← get r ;;
  ret r1.
```

While this readable code is not the literal output of the translation, it is the result of some careful (but semi-automated and verified) unfolding and simplification. The produced code also takes an “identifier”, `id0`, as input: this determines which locations on the heap it will use for its local memory. This technical detail will be explained in Section 5 and can safely be ignored for now.

2.2.4 Equivalence of Implementation and Specification. Now that we have both translations to SSProve, we can prove that they are equivalent in our program logic.

```
Theorem xor_equiv : ∀ id0 w1 w2,
  ⊢ { λ '(h0, h1), ⊤ }
    JXOR id0 w1 w2 ≈ hacspecc_xor_state w1 w2
  { λ '(v0, h0) '(v1, h1), v0 = v1 }.
```

The precondition is a predicate over the two initial heap states and the postcondition is a predicate over the two final heaps and values. The notion of equivalence we use here to relate the two functions only requires the return values `v0`, `v1` of the two programs to be equal, provided we run them both on the same inputs. In particular, we do not make assumptions or restrict how the two programs use the heaps `h0`, `h1`. The programs are thus allowed to use different locations to store their intermediate values. This theorem is proved using the rules of the relational program logic of SSProve [1].

2.2.5 Security Proof for the OTP Implementation. We now prove perfect cryptographic security of the Jasmin implementation of OTP using XOR. To this end, we first need to define some terminology. In SSProve a *package* is a finite set of procedures that might contain calls to external procedures. The set it implements is called its *export interface* and the set on which it depends its *import interface*. A *game* is a package with no imports and a *game pair* is a pair of games that export the same procedures. These can be used to model cryptographic games, e.g., a game pair might consist of a

²This use of mutability is for illustrative purposes only.

real encryption scheme and an oracle: these have the same interfaces but different implementations.

For OTP we define the game pair consisting of an implementation of OTP using the Jasmin code and an implementation which is obviously secure. The Jasmin game is the package `JOTP_real` exporting the single procedure:

```
Definition JOTP id0 m :
  k_val ← sample_uniform('word n) ;;
  JXOR id0 m k_val.
```

We already have a security proof for the package `OTP_real` exporting the single procedure:

```
Definition OTP m :
  k_val ← sample_uniform('word n) ;;
  ret m ⊕ k_val.
```

This game is already proven to be indistinguishable under chosen plaintext attack from an implementation where the message is chosen at random. The statement and proof are in the SSProve library. This is done by proving that, when the input is disregarded and a random message is encrypted, the advantage of an attacker in distinguishing between `OTP_real` and a game `OTP_ideal` is zero.

If we can prove that `JOTP_real` is perfectly indistinguishable from `OTP_real`, then we can combine the two results using the triangle inequality for advantages of games (Lemma 1 in the SSProve paper [1]) and prove that an adversary also cannot distinguish between `JOTP_real` and `OTP_ideal`, i.e., the Jasmin implementation is IND-CPA. That is, we only need to prove the following theorem.

Lemma `JOTP_OTP_perf_ind id`: `JOTP_real id` $\approx_0$ `OTP_real`.

Here $\approx_0$ means that the advantage of an adversary trying to distinguish between the two games is zero. To prove this lemma we use Theorem 1 from the SSProve paper [1], which allows us to conclude if we can prove the following code equivalence for all m and some *stable invariant* `inv`:

```
⊢ { λ '(s0, s1), inv (s0, s1) }
  JOTP id0 m ≈ OTP m
{ λ '(b0, s0) '(b1, s1), b0 = b1 ∧ inv (s0, s1) }.
```

For the precise definition of stable invariant see Section 4.2 of the SSProve paper [1]. In our case, we can use the invariant `heap_ignore`, which asserts that both heaps are preserved during execution if the locations used by `JXOR` are ignored.

Combining this result with the already established security of `OTP_real` we get security of `JOTP_real`.

```
Theorem unconditional_secretcy_jas : ∀ LA A,
  fdisjoint LA xor_locs → ValidPackage LA
  [ interface #val #[i1] : 'word → 'word ]
  A_export A →
  Advantage IND_CPA_jasmin A = 0.
```

That is, for all valid adversaries A with a matching interface, and all regions of adversarial memory LA , if the adversary

cannot use the same locations as `JXOR` then their advantage in distinguishing between `JOTP_real` and `OTP_ideal` is zero.

3 Background & Technical Preliminaries

3.1 Hacspect

Hacspect is a High Assurance Cryptography SPECification language [15, 27, 31] aiming to provide a common language to programmers, cryptographers and proof engineers. It proposes to make future internet standards, such as those published by IETF and NIST, machine-readable. Hacspect is a subset of Rust which makes it executable and accessible to cryptographic engineers.

The Hacspect language was carefully crafted to have a functional semantics, in which assignments are translated to let-expressions. The Hacspect tool comes with functional translations to the purely functional languages of several proof assistants, currently F^* , Coq, and EasyCrypt. As such it is a convenient tool to share specifications across proof assistants.³ Hacspect also comes with an operational semantics [31], but since the semantics is not formalized in the backends, the functional translation cannot be verified against it. Instead, this translation constitutes the authoritative semantics. This motivates our choice to relate our imperative translation via translation-validation⁴.

Currently, all Hacspect backends use a functional semantics. However, both in EasyCrypt and in Coq/SSProve, one could also choose to use a translation to an embedded imperative language. This can be seen as one of the benefits of Hacspect, as anyone familiar with either functional or imperative coding paradigms will understand the Hacspect specification. We will explain how to do so in Section 4.

3.2 Jasmin

Jasmin [4] is a low-level language designed for implementing high-speed cryptography, with a verified compiler backend supporting the x86 and ARM architectures. The language has a formal big-step operational semantics in Coq. The Jasmin compiler is also implemented and verified in Coq, in the sense that it preserves the semantics of the Jasmin source [4, 5] and also that it does not introduce timing side-channel attacks [6]. We give a condensed overview of Jasmin, focusing on the aspects that are interesting for the sake of our discussion, and limiting the explanation to a few representative examples. For more details please see the Jasmin paper [4].

³This also allows one to combine code generated from different proof assistants. For example, one could combine a hash function from F^* and an elliptic curve implementation from Coq, both of which would be specified in Hacspect, verified, and then extracted to C (or Rust, or ASM). This is the methodology proposed in the libcrux library [27].

⁴The operational semantics of Hacspect would be a good target for future formalization.

3.2.1 The Language. Jasmin is an imperative language with structured control flow in the form of loops, conditionals, and procedure calls. Jasmin has types for booleans, integers, bit-words of various sizes, and arrays. Despite these high-level features, the Jasmin compiler produces predictable assembly code, which enables efficient and secure cryptographic implementations. For instance, the programmer can use architecture-specific assembly instructions and can specify whether procedure-local variables should be stored in registers (using the `reg` keyword) or on the stack (using the `stack` keyword). Jasmin’s operational semantics was carefully crafted to hide low-level details such as the distinction between the storage types `reg` and `stack`. Our correctness theorem for the translation from Jasmin to SSProve, like Jasmin’s compiler correctness theorem, is proven with respect to this operational semantics, and we can thus safely ignore such distinctions.

A Jasmin program P consists of a list of non-recursive function definitions, associating to each function name f a list of variables used for arguments $P(f)_{param}$, variables used for returning results $P(f)_{res}$, and a command, i.e., a sequence of instructions $P(f)_{body}$ for the body of the function.

Instructions include assignments, operators, conditionals, for and while loops, and function calls. Expressions occurring in instructions include variable and array access, arithmetic and logical operators, as well as assembly operations such as shifts, increments, *etc.*

3.2.2 Jasmin State. Jasmin features both global and local state, denoted by a pair (m, ρ) of a *global memory* m and local *variable map* ρ . A variable is local when it is declared within a function, and global when declared at the top level. We will write $\rho[\cdot]$ and $\rho[\cdot \leftarrow \cdot]$ respectively for local variable map lookup and update. For global state, we will write $m[\cdot]_i$ and $m[\cdot \leftarrow \cdot]_i$ for lookup and storage of *size* i , given in bits (possible values are 8, 16, 32, 64, 128, 256). Global state is indexed by integers (pointers) and local state by variables (strings). Note that looking up memory in Jasmin can fail, so we will abuse notation by denoting by $m[p]_i = v$ that v is stored at p in m and that it is valid to make a read of size i at p in m . We will do the same for writes.

3.2.3 Jasmin Operational Semantics. The operational semantics of Jasmin is mostly standard. A judgment of the form $\langle c \mid (m, \rho) \rangle \Downarrow (m', \rho')$ means that for an initial state (m, ρ) , execution of the command c terminates in the final state (m', ρ') , and $\langle e \mid (m, \rho) \rangle \Downarrow_{exp} v$ means that the expression e evaluates to the value v under state (m, ρ) (expressions can only read, not modify the state). All judgments are implicitly parametrized by an ambient program (i.e., list of function definitions), which will not be mentioned explicitly unless required. For instance, in the rule for assigning a local variable in Figure 2 we start by evaluating the expression e to v . We then look up the type α of the variable x , and perform

a truncation⁵ of v at type α , yielding v' compatible with the type of x . Finally, we update the local state to $\rho[x \leftarrow v']$, while the global state remains unchanged.

$$\begin{array}{c}
 \text{ASSGN} \\
 \frac{\langle e \mid (m, \rho) \rangle \Downarrow_{exp} v \quad \alpha = ty(x) \quad v' = \|v\|^\alpha}{\langle x = e \mid (m, \rho) \rangle \Downarrow (m, \rho[x \leftarrow v'])} \\
 \\
 \text{FUNCALL} \\
 \frac{\begin{array}{l} \langle e_i \mid (m, \rho_0) \rangle \Downarrow_{exp} v_i \quad \text{for } i = 1, \dots, k \\ \langle f(v_1, \dots, v_k) \mid m \rangle \Downarrow_{call} \langle (w_1, \dots, w_n) \mid m' \rangle \\ \langle x_j = w_j \mid (m', \rho_{j-1}) \rangle \Downarrow (m', \rho_j) \quad \text{for } j = 1, \dots, n \end{array}}{\langle x_1, \dots, x_n = f(e_1, \dots, e_k) \mid (m, \rho_0) \rangle \Downarrow (m', \rho_n)} \\
 \\
 \text{CALLRUN} \\
 \frac{\begin{array}{l} \text{let } \rho_0 = \emptyset \text{ and } c = P(f)_{body} \\ \text{and let } y_i = (P(f)_{param})_i \text{ and } x_j = (P(f)_{res})_j \\ \langle y_i = v_i \mid (m, \rho_{i-1}) \rangle \Downarrow (m, \rho_i) \quad \text{for } i = 1, \dots, k \\ \langle c \mid (m, \rho_k) \rangle \Downarrow (m', \rho') \\ w_j = \|\rho'[x_j]\|^{ty(x_j)} \quad \text{for } j = 1, \dots, n \end{array}}{\langle f(v_1, \dots, v_k) \mid m \rangle \Downarrow_{call} \langle (w_1, \dots, w_n) \mid m' \rangle}
 \end{array}$$

Figure 2. Excerpt of Jasmin operational semantics

The main subtlety for translating Jasmin to SSProve arises from function calls and their treatment of local state. The execution of function calls in Jasmin is split into two rules. The perspective of the caller is captured by `FUNCALL`: We evaluate the arguments e_i and perform the call to the function f according to the callee’s perspective. We obtain a new global state m' and store the resulting values w_j in the caller-local variables x_j . Jasmin’s type checker guarantees that the number of returned values equals the number of variables. Crucially, when switching from caller to callee, we *retain the local state* ρ_0 and pass only the global state m to `CALLRUN` as witnessed by the use of $\Downarrow_{call}$ relating pairs of instructions and global memories and values and global memories.

To describe the callee perspective, we write ρ_0 for the empty local state, and c , y_i , and x_j for the body, parameter-, and result-variables of f respectively. Each argument v_i is stored in the local variable y_i according to the definition of parameters of $P(f)_{param}$. We then execute c from state (m, ρ_k) , yielding (m', ρ') . We obtain the values $w_1, \dots, w_n$ by reading the result variables x_j from the local state ρ' and truncating as necessary. Finally, the local state ρ' is discarded, and the result values and updated global state m' are returned.

⁵This truncation only exists at the high level to mimic the implicit truncations happening at the assembly level. In practice, the types of v and x mostly agree and the truncation can be simplified away.

3.3 SSProve

SSProve is a Coq library for modular cryptographic proofs introduced by Abate et al. [1]. We only review the concepts needed to understand the current paper. More details can be found in the extended version of the SSProve paper [25].

3.3.1 Code. In this paper, we de-emphasize the probabilistic capabilities of SSProve, as they are not currently reflected in Jasmin. Thus, for our purposes, SSProve essentially embeds a stateful language inside Coq using a monad called `raw_code`. In `raw_code` A one can (1) embed any pure value x of type A using `ret x`, (2) read from a memory location ℓ to a variable x , and use x in a continuation k , written $x \leftarrow \text{get } \ell \ ; \ k \ x$, (3) write a value v to a memory location ℓ and then continue with k , written `put  $\ell$  ; k`, (4) sequentially combine $u : \text{raw_code } X$ and $k : X \rightarrow \text{raw_code } A$ using the bind operator that we write $x \leftarrow u \ ; \ k \ x$. It is also possible to sample from a distribution D in this monad using $x \leftarrow \text{sample } D \ ; \ k \ x$ as shown in Section 2.

3.3.2 Memory Model. Memory locations consist of a natural number and a type that together serve as an index in a global shared memory. This global state is represented as a map from locations to values. We say that a state is valid for a set of (typed) locations when all locations point to values of the matching type. Note that to be able to use the type in the key of the memory, we must in fact use codes of types; since SSProve is built for probabilistic programs, these codes represent types on which one may build (discrete) distributions. In type-theoretic terms, they encode a universe of datatypes `choice_type` which represents a subset of `mathcomp's choiceType` [30, §8.3]. For the purposes of our translation, we use a modified version of SSProve where `choice_type` is extended to include sums, words and lists. This allows us to encode all the types needed to represent Hacspect and Jasmin programs. Memory is simulated using a structure we call `heap`, essentially a map from locations to values. We would like to stress the fact that in SSProve the memory is *global*, in contrast to Jasmin's function local state. Thus, one must take care to generate code without overlapping locations. We address this in Section 5.

For a heap h , location ℓ and value v , we will write $h[\ell]$ and $h[\ell \leftarrow v]$ for heap lookup and storage (as for Jasmin state).

3.3.3 Packages. Another defining feature of SSProve is that of packages. Packages are used extensively to compose modular security games in the style of state-separating proofs [17]. Since our methodology allows us to reuse existing security proofs [25], we will not get into the details of security proofs, so we only introduce packages briefly. Packages are collections of procedures that can all refer to the same set of locations and invoke certain procedures that are part of an *import interface*. The signature of this collection defines the *export interface* of the package. Packages can thus be combined modularly to create bigger packages. For

instance, a package can be linked to another that implements its import interface or they can be composed in parallel to export the union of their respective export interfaces.

3.3.4 Relational Hoare Logic. Finally, SSProve features a (probabilistic) relational Hoare logic that allows us to prove the relational properties of programs. Once again, we will focus on the stateful but deterministic fragment. In this program logic, we prove judgments of the form

$$\vdash \{\phi\} \ c_0 \ \sim \ c_1 \ \{\psi\}$$

where c_0 and c_1 are two code pieces we compare and ϕ and ψ are respectively a pre- and a postcondition relating (1) the initial heaps (for ϕ); (2) the final heaps and final return values of both code pieces (for ψ). For deterministic code, this is equivalent to: for all initial memory states m_0 and m_1 such that $\phi(m_0, m_1)$ holds, running c_i in state m_i will yield final state m'_i and final value v_i such that $\psi(v_0, m'_0)(v_1, m'_1)$ holds.

SSProve comes with a number of rules for this logic and provides tactics to facilitate writing proofs. Moreover, one can fall back on the semantics above to prove judgments [25].

3.4 Interoperability

For the sake of getting Hacspect, Jasmin, and SSProve to interact smoothly, we had to extend each of them in a minor way. We did not, however, make any modifications to the core projects that would change the interpretation of any of the statements that can be found in the published literature.

Specifically, besides the translations which constitute the core contributions of this work, we made the following additions. For SSProve, we added sum types to represent the result types of Hacspect. We also added bitwidth-indexed machine words as well as lists. For Jasmin, we added the ability to pretty-print the Coq abstract syntax tree of a parsed Jasmin program, and we added the definition of the Intel AES-NI instructions [23] to Jasmin's x86 semantics, since the AES-NI instructions are used in the AES case study. Hacspect remained unchanged.

4 Hacspect & SSProve

Hacspect facilitates proving the correctness of efficient implementations with respect to a specification by translating it to multiple proof assistants. We further this goal by adding a translation from any valid Hacspect specification to SSProve. This imperative translation is accompanied by a pure translation, which adds a wrapper around the existing Coq translation to embed it into SSProve. We can thus compare the imperative and pure translations using SSProve's relational logic and automatically generate a proof stating that they return the same values.

4.1 The Functional Translation

The pure translation constitutes a minor modification of Hacspecc's existing Coq backend [35] that we undertook to facilitate the connection to Jasmin. Coq does not provide a standard library for machine integers, so the existing backend chose the CompCert library to model machine integers [28]. Jasmin uses its own word library. In the long run, we would hope for a unified word library in the Coq ecosystem. Meanwhile, we changed the backend to use Jasmin words.

We translate for-loops as a fixed point with an accumulator of all the mutable variables changed inside the loop. Hacspecc has support for early return of option or result types. We model these early returns using the option and error monad. We thus need a fold operation that respects the monadic operations to allow early returns in for-loops.

4.2 The Imperative Translation

Since we provide the first translation from Hacspecc to an imperative programming language, we need to extend the information gathered in the translation from Hacspecc to the various backends. SSProve needs information about what memory locations and functions are used in a given scope. To compute this, we add static dependency analysis to the Hacspecc pipeline. This is done by walking the AST for every block of code and adding a unique memory location for each mutable variable. In a second pass, we then unify the memory locations used by all the local function calls, to get the total set of memory locations a function might change.

The translation evaluates arguments passed to function calls or operators before evaluating the function or operator. This is done by binding the arguments to temporary values, which are then passed to the function. This makes it easier to prove equality to another SSProve implementation, as we can first prove that all the arguments are equal, and then show that the functions agree on equal input.

A subtlety arises from the fact that Hacspecc supports early return statements: $x = e?$ is operationally equivalent to

```
x = match e { Some(v) => v, None => return None }
```

In particular, if e evaluates to **None**, the ambient function in which the statement $x = e?$ occurs returns early with the result **None**. Since SSProve's `raw_code` does not support control effects, we cannot directly represent this `return`. We instead embed Rust code with early returns into the option monad. To ensure that this encoding interacts well with the effectful operations of SSProve which manipulate state, we define a special bind operation, combining the two monads.

```
Definition obind (x : raw_code (option A))
  (f : A → raw_code (option B)) : raw_code (option B)
:= t_x ← x ;;
   match t_x with Some s => f s | None => ret None end.
```

The Hacspecc code we translate carries sufficient typing information to determine whether a function may return early.

We leverage this information to select between this custom bind operator and SSProve's standard bind. For example:

```
x = f(v)? ; y = g(x) ; y + 2
```

is translated to the following SSProve code:

```
temp_x ← f(v) ;;
obind temp_x (λ x, temp_y ← g(x) ;; temp_y .+ 2)
```

4.3 Equivalence Between the Hacspecc Translations

On the one hand, it is often easier to define and prove properties for a functional specification. On the other hand, it is easier to show an efficient imperative implementation equivalent to an imperative specification. So, it is desirable to derive an equality between the imperative and functional translations. We automatically generate such a proof, as part of the translation from the Hacspecc specification. To achieve this we first define a record `both`, which has projections to a piece of code for the functional translation and for the imperative translation. It also contains the proof of equivalence for the two pieces of code. We traverse the AST building the functional translation, the imperative translation and their equivalence at the same time. This is achieved by using compositional blocks for the control structures of Hacspecc.

An example of such block is the one used for `let` expression in Hacspecc, where the functional translation is a functional let binding in Coq, while the imperative translation uses `bind` in SSProve. The equivalence can be proven using the `bind` rule in SSProve, since we have a proof of equality of the arguments and a proof of equality of the rest of the code bodies. Other blocks are loops, mutable let bindings (where a location is used, as shown in Section 2.2.1), early returns, operator calls, lifting pure values, etc. We can therefore get the full translation to the imperative and functional code, together with the equality between them, by chaining these compositional blocks. This also requires us to define all the library functions in Hacspecc in the `both` type. Using this combined type, we can write elements in a style where the translation looks close to the original specification and can be made more readable by the notation engine of Coq.

5 Jasmin & SSProve

5.1 Memory

A major difference between the Jasmin and SSProve semantics is how memory is handled: SSProve only has a global notion of memory and Jasmin supports both global and local variables. To model local variables in SSProve, we parameterize all translated code over a “base stack frame ID” which reserves an (*a priori* unbounded) region of SSProve's global memory for local variables. Then instantiating translated code with a concrete base stack frame ID correctly assigns new stack frame IDs to all its called functions. In particular, we prove that variables translated with different stack frame IDs never overlap, i.e., translation of variables is injective

w.r.t. IDs. We store the Jasmin global memory in a map (from integers to bytes) at a static location called MEM.

5.2 Program Translations

We now describe our translation from Jasmin to SSProve, meaning the translation of programs, but also of types, values, expressions and commands. As a first step, we use the Jasmin compiler to pretty-print the internal AST corresponding to a Jasmin source program to Coq syntax. Since this AST datatype was extracted from Coq in the first place, it amounts to ‘de-extracting’ it back to Coq. Our translation thus translates Coq’s datatype of Jasmin programs to SSProve programs (i.e., the `raw_code` monad).

5.2.1 Types and Values. The only base types missing from SSProve’s `choice_type` (the restricted set of types which a `raw_code` can return; see Section 3.3.2) were words and arrays. Following Jasmin, we use the `coqword` library’s type of words, which is based on the `mathcomp` library [30]. We represent arrays as maps from integers to bytes. The only minor difference is our implementation of maps differs from Jasmin. Using similar types makes it easy to embed Jasmin values into SSProve values (via the identity) for all except array values. We denote the function taking Jasmin values to SSProve values by `translate_value`.

5.2.2 Expressions. For the translation of expressions (denoted `translate_pexpr`) we have to be careful and do the right casts and truncations, as dictated by the semantics of Jasmin: e.g., when looking up in an array, the index is always cast to an integer type. For the translation of function applications in expressions (additions, subtractions, *etc.*), we reused the semantics from Jasmin expressions, by transporting values back to Jasmin types, applying the operations, and then transporting back to SSProve types. Note that this transport is only non-trivial for arrays. This simplifies the proof significantly, only requiring us to prove that all operations are invariant under this transport.

5.2.3 Instructions. The main difficulty in translating instructions is translating function calls; for calls to operations we could mostly use the same solution as for expressions and for for-loops we simply iterate the translated body. To be able to call functions, we choose to let our translation keep track of previously translated functions, and only allow these to be called; this avoids cyclic calls and recursion (which are always rejected by the Jasmin compiler). Furthermore, we make sure to call these translated functions with a fresh stack frame ID to avoid collisions between local variables.

Note that we currently do not translate Jasmin while loops, as they do not have a correspondent in SSProve. This does not constitute a conceptual problem in practice, since for-loops are sufficient for most cryptographic routines.

5.2.4 Programs. We translate Jasmin programs, which map function names to function declarations (Section 3.2.1),

to maps from function names to SSProve functions taking an ID and a list of inputs to SSProve code.

5.3 Unary Deterministic Judgments

SSProve originally supported only relational judgments of the form $\vdash \{\phi\} c_0 \sim c_1 \{\psi\}$, as presented in Section 3.3. For the sake of our correctness theorem, we want to relate a translated Jasmin term c_0 to the value v it evaluates to, i.e., c_1 is always of the form `ret v`. Since Jasmin’s semantics is deterministic, we do not need the full power of a probabilistic judgment. We thus extend SSProve and build a new unary judgment on top of the relational logic, to deal with the special case where we relate a `raw_code` with a return value: $\vdash \{\phi\} c \Downarrow v \{\psi\}$. Here ϕ is a precondition on the initial state of c , while ψ is a postcondition on the final state after running c . The postcondition no longer mentions a final state or return value for the right hand side, instead the return value v is part of the judgment. We define $\vdash \{\phi\} c \Downarrow v \{\psi\}$ as the following judgment relating c to `ret v`:

$$\vdash \{(m_0, m_1). \phi m_0\} c \sim \text{ret } v \{(a_0, m'_0), (a_1, m'_1). \psi m'_0 \wedge a_0 = a_1 \wedge a_1 = v\}$$

The precondition only considers the memory of the left-hand side, while the postcondition also states that both sides must produce the value v .

While this unary judgment is conceptually simpler than the relational logic, we have found it beneficial to reuse the existing theory instead of starting from scratch. An advantage of this is that we can easily leverage the rules of the relational program logic and the tactics provided by SSProve to prove unary judgments. Moreover, we establish a precise connection between the two logics by proving that whenever c is free of sampling operations, the judgment above is equivalent to saying that running c on any initial state m such that ϕm will yield return value v and final state m' such that $\psi m'$. For instance, we obtain the expected rules for values, sequential composition, and writing to the heap.

$$\begin{array}{c} \frac{\forall m. \phi m \implies \psi m \wedge v = v'}{\vdash \{\phi\} \text{ret } v \Downarrow v' \{\psi\}} \\ \frac{\vdash \{\phi\} c \Downarrow u \{\xi\} \quad \vdash \{\xi\} k u \Downarrow v \{\psi\}}{\vdash \{\phi\} x \leftarrow c;; k x \Downarrow v \{\psi\}} \\ \frac{\vdash \{\lambda m, \exists m', \phi(m') \wedge m = m'[\ell \leftarrow v]\} r \Downarrow w \{\psi\}}{\vdash \{\phi\} \text{put } \ell v;; r \Downarrow w \{\psi\}} \end{array}$$

Other rules can also be derived straightforwardly from the definition of the unary judgment as analogues of the relational rules, which are detailed by Haselwarter et al. [25].

5.4 Correctness Theorem

We prove that our translation preserves the semantics of well-defined programs. To do this we define a relation between Jasmin memory states and SSProve memory states. First, we relate the global Jasmin memory to the “global memory map” stored on the heap in SSProve. We say that the global Jasmin state m is related to the heap h when, if one can successfully read a single byte at an address from the Jasmin memory, then one can look up the corresponding value in the “global memory map” stored at MEM on the SSProve heap:

$$m \sim h := \forall p v. m[p]_8 = v \Rightarrow h[\text{MEM}][p] = v$$

To relate the local memory of Jasmin and our encoding of local memory in SSProve, we define a relation between a variable map ρ and a heap h relative to a **stack frame ID** ι . We write $h[x]'$ for the lookup of the variable x on the heap relative to ID ι . A variable map ρ is related to the heap h w.r.t. ι if successfully looking up a variable x in ρ implies that looking up x on h relative to ι yields the same value:

$$\rho \sim_\iota h := \forall x v. \rho[x] = v \Rightarrow h[x]' = v$$

Now, the relation between a Jasmin memory pair (m, ρ) (of global and local state) and an SSProve heap is not just the conjunction over all these relations, since we need to know that a function can make an arbitrary number of function calls, each with their own local state, and not run out of space on the heap. To state this we need some terminology: We say that a stack frame ID ι is **fresh** w.r.t. a heap h when $\rho_0 \sim_\iota h$ holds, where ρ_0 is the empty variable map. We assume that we have a prefix order $\leq$ on stack frame IDs and say that a stack frame ID s is **valid** w.r.t. a heap h when all strict successors of s are fresh w.r.t. h , i.e., for all $s' > s$, $\rho_0 \sim_{s'} h$. Furthermore, we say that two IDs s_1 and s_2 are **disjoint**, when there is no ID which they are both a prefix of. Concretely, we require for all IDs s that $s_1 \leq s$ and $s_2 \leq s$ do not both hold simultaneously. We assume that storing at disjoint ID locations preserves values: if s_1 and s_2 are disjoint then $\forall x, y. h[y \leftarrow v]^{s_2}[x]^{s_1} = h[x]^{s_1}$.

For a variable map ρ , two stack frame IDs ι, σ (main and sub-ID) and a set I of IDs we say that the tuple (ρ, ι, σ, I) is a **stack frame**. We say that a stack frame (ρ, ι, σ, I) is **valid** w.r.t. a heap h when the following conditions hold: (1) σ is valid w.r.t. h , (2) $\rho \sim_\iota h$, (3) $\sigma \notin I$, (4) for all $\sigma' \in I$, $\iota < \sigma'$, σ' is disjoint from σ and σ' is valid w.r.t. h , (5) for all $\sigma', \sigma'' \in I$, σ' and σ'' are disjoint.

The intuition for a valid stack frame (ρ, ι, σ, I) is that ρ should be related to the main stack frame ID ι , and the sub stack frame ID σ should be a valid ID from which the current function can spawn new functions with fresh memory; I is there to keep track of which IDs are currently in use and to which variable maps they relate. Note that the set I is only needed for the proof of correctness, and is not actually used in the translation of a given program.

A **stack** is then a list of stack frames. The empty stack is denoted by S_0 . A stack frame (ρ, ι, σ, I) is **disjoint** from a stack S when ι is disjoint from all sub IDs and IDs occurring in sets of the stack frames on S . A stack S is **valid** w.r.t. a heap h when either S is empty or $S = F :: S'$ where S' is a valid stack and F is a valid stack frame disjoint from S' .

Using these constructions we can finally define our relation on Jasmin and SSProve states. A Jasmin state pair (m, ρ) is related to the heap h w.r.t. the stack S , which we write $(m, \rho) \sim_S h$, when the following conditions hold: (1) S is valid w.r.t. h , (2) $m \sim h$, (3) ρ is the variable map at the top of the stack, i.e., the top of the stack is of the form (ρ, ι, σ, I) . This relation satisfies two key lemmas, which are needed to prove the correctness of our translation.

Lemma 1 (Push empty stack frame). *If $(m, \rho) \sim_{(\rho, \iota, \sigma, I) :: S} h$ and σ_1, σ_2 are two disjoint IDs with $\sigma < \sigma_1, \sigma_2$, then*

$$(m, \rho_0) \sim_{(\rho_0, \sigma_1, \sigma_2, \emptyset) :: (\rho, \iota, \sigma_2, I) :: S} h.$$

Lemma 2 (Pop stack frame). *Let $F_i = (\rho_i, \iota_i, \sigma_i, I_i)$, then if $(m, \rho_2) \sim_{F_2 :: F_1 :: S} h$ then $(m, \rho_1) \sim_{F_1 :: S} h$.*

These two lemmas correspond to (1) calling a function and assigning it a fresh region of memory for local state and (2) returning from a function call to its caller, accounting for the operational semantics of Jasmin function calls according to Figure 2. Note in Lemma 1 that the sub-ID of the calling stack frame, σ , is updated to a fresh ID σ_2 , and that we initialize the callee frame with the same main and sub ID σ_1 , since when the frame gets pushed in a function call, the callee has not invoked any further functions yet.

Using this relation, we show how our translation of Jasmin code relates to its source. For example, if we consider the function `translate_pexpr`, which translates Jasmin expressions to `raw_code`, we get the following correctness lemma.

Lemma 3. *Let v be a value, e an expression, s a Jasmin state pair and S a stack. If $\langle e \mid s \rangle \Downarrow_{\text{exp}} v$ then*

$$\vdash \{h. s \sim_S h\} \quad \text{translate_pexpr } S \ e \ \Downarrow \quad \text{translate_value } v \ \{h. s \sim_S h\}$$

As evaluating expressions does not have memory side effects, the relation between Jasmin and SSProve states is preserved under expression translation.

We now prove the main theorem, which establishes the connection between *function calls* in Jasmin and in SSProve:

Theorem 1. *Let P be a Jasmin program, (m, ρ) a Jasmin state-pair, f a function name, and v_i, w_i values for $i = 1, \dots, k$. Furthermore, let $\iota, \sigma, \sigma_1, \sigma_2$ be IDs such that σ_1 and σ_2 are disjoint and strict successors of σ . If P' is the result of translating P and $\langle f(v_1, \dots, v_k) \mid m \rangle \Downarrow_{\text{call}} \langle (w_1, \dots, w_n) \mid m' \rangle$ then*

$$\begin{aligned} &\vdash \{h. (m, \rho) \sim_{(\rho, \iota, \sigma, I)} h\} \\ &\quad P' \ f \ \sigma_1 \ \text{translate_values } (v_1, \dots, v_k) \\ &\Downarrow \quad \text{translate_values } (w_1, \dots, w_n) \\ &\{h. (m', \rho) \sim_{(\rho, \iota, \sigma_2, I)} h\} \end{aligned}$$

The theorem states that if calling the function f in the Jasmin program P and global memory m with arguments $\vec{v}$ results in the new global memory m' and returns the values $\vec{w}$, then we can conclude two things:

1. Calling the function at a fresh ID (σ_1) and with the translation of the arguments $\vec{v}$ evaluates to the translation of the return values $\vec{w}$.
2. After calling the translated function, the global memory m' is related to heap where we have updated the sub-ID to a fresh one (from σ to σ_2).

This is the expected behavior: calling a function can change the global but not the local state. We have to update our sub-ID because the previous one is no longer fresh, as we might have stored local state inside the function call.

6 AES Example

As a larger case study of our framework, we verify the security of a Jasmin implementation of a PRF-based encryption scheme using AES and prove it equivalent to a Hacspeg reference implementation. The Jasmin implementation and the general methodology for proving security are similar to the presentation in EasyCrypt [8], but we use our toolchain based on SSProve to conduct the formalization.

The workflow for proving security of our AES implementation is as follows:

1. Implement the encryption scheme in Hacspeg and Jasmin.
2. Translate the two implementations to SSProve code.
3. Prove the two translations equivalent and prove security properties of the Jasmin translation.⁶

We skip implementing the Jasmin code by reusing the implementation from the EasyCrypt and Jasmin tutorial [8], which relies on the Intel AES-NI hardware acceleration instructions [23]. Our reference implementation in Hacspeg is based on the NIST standard [20], and it successfully passes the corresponding public test vectors [20, 23].

For the security analysis, we prove indistinguishability under chosen plaintext attack (IND-CPA) of the AES implementation of the PRF-based symmetric encryption scheme described below. Concretely, we prove that the advantage of an adversary in distinguishing the encryption of a message from the encryption of a random message is (linearly) bounded by the advantage of the same adversary in distinguishing AES from a PRF. For details on the concrete bounds, see the SSProve journal paper [25, §2.3].

As was the case in Section 2, we do not have to write a security proof of the abstract encryption scheme from scratch, since such a proof, for an abstract PRF, is already present

in the SSProve library [25, §2.3]. To connect this with our efficient implementation, we need to prove that an adversary cannot distinguish between the efficient implementation and the abstract implementation given in SSProve [25, §2.3] instantiated with a Coq implementation of AES.

As in *loc. cit.*, our definitions follow SSP methodology [17]. The PRF-based encryption scheme is given by the code:

```
Definition PRF_ENC f m :=
  k_val ← kgen ;; enc m k_val.
```

Here, `kgen` is a key generation code that uniformly samples a key on its first invocation and returns a fixed key on subsequent calls. The `enc` function is given by the code:

```
Definition enc m k :=
  r ← sample uniform N ;;
  let pad := f r k in let c := m ⊕ pad in
  ret (r, c).
```

Here f is the function which we assume to be a PRF and which we will instantiate with AES. The PRF is used to generate a pad from a uniformly sampled nonce r ; the ciphertext is computed as the xor of the message and the pad. For all functions $f : \text{word} \rightarrow \text{word} \rightarrow \text{word}$ we denote the game consisting of the single export `PRF_ENC f` by `PRF_real f`.

We reuse the SSProve proof [1, §2.3] by showing that `PRF_real aes` is perfectly indistinguishable from the same scheme with `enc` replaced by the translated Jasmin code.

The high-level structure of the security analysis of the implementation is as follows:

1. Write an intermediate imperative implementation directly in SSProve code.
2. Write a functional implementation directly in Coq.
3. Prove the equivalence between the intermediate implementation and the functional implementation.
4. Prove the equivalence between the translated implementation and the intermediate implementation.
5. Connect the equivalences to the existing security proof of the abstract encryption scheme.

Steps (1) and (2) can also be copied almost verbatim from the EasyCrypt development: the syntactic similarities of the EasyCrypt and SSProve codes make the translation very straightforward. For the proofs in steps (3) and (4) we can reuse some parts, e.g., the loop invariants, but in general the differences in the programming languages and the underlying proof assistants require new proofs.

6.1 Translation

As mentioned in Section 2, we start by printing the Coq ASTs of all the involved functions during Jasmin compilation. Then we use the translation described in Section 5 to obtain SSProve code for each function used in the implementation.

⁶Here we deviate slightly from the intended workflow from Section 2 by doing the security proof on the implementation instead of the specification. The reason for this is simply that parts of security proof about the Jasmin implementation were already completed when the Hacspeg specification was added to the project.

6.2 Specification

Next, we write intermediate specifications for the Jasmin functions. Compared to the example in Section 2, these correspond to the pure Coq XOR function. As mentioned, we take inspiration from the specifications in the EasyCrypt and Jasmin tutorial [8]. This step removes translation artefacts (e.g., compiler-generated memory locations) and allows us to focus on proving the underlying logical statements.

6.3 Equivalences for Intermediate Code

Then we prove that our intermediate implementations are equivalent to functional (stateless) Coq functions. The statements we prove are generally of the form:

$$\vdash \{(m_0, m_1). \phi(m_0, m_1)\} \ c \ i \ \sim \ \text{ret } (f \ i) \ \{(a_0, m'_0), (a_1, m'_1). \phi(m'_0, m'_1) \wedge a_0 = a_1\}$$

where i is arbitrary input, c is the intermediate SSProve code and f is the pure Coq function. Note that we also prove that these equivalences preserve the precondition ϕ ; for the equivalences to hold we usually have to assume that ϕ is stable w.r.t. memory locations used by c .

Even though f is usually stateless, we have to keep the heap of the right-hand side in mind, since it might be relevant in certain contexts; otherwise we could have used the unary judgments of Section 5.3.

6.4 Equivalences for Translated Code

When reasoning about the code generated by our translation from Jasmin to SSProve, we have to prove equivalences of the following form:

$$\vdash \{(m_0, m_1). \phi(m_0, m_1)\} \ P' \ F \ id \ i \ \sim \ c \ i \ \{(a_0, m'_0), (a_1, m'_1). \phi(m'_0, m'_1) \wedge a_0 = a_1\}$$

where P' is the translated Jasmin program, i is an arbitrary input, id is a stack frame ID, F is the function name in the Jasmin program and c is the intermediate code.

Once we have proven such an equivalence, we can reuse it in proofs where F appears as a called function. It is therefore important that the equivalences are parametric in id . We also want to preserve the precondition ϕ and again we have to assume that ϕ is stable w.r.t. the locations of F and c . However, there is one issue here: the locations set of F is not straightforward to compute and might also be rather large. Instead we require that ϕ is stable w.r.t. *all possible* locations used by ϕ , i.e., locations stored using an id' with prefix id ($id \leq id'$). This turns out to be a sufficient and reasonably manageable invariant to preserve.

6.5 Connecting AES to the PRF Security Proof

The encryption function of which we want to prove the security can be implemented in Jasmin as:

```
fn enc(reg u128 n, reg u128 k, reg u128 p) -> reg u128 {
  reg u128 mask, c;
  mask = aes(n, k);
```

```
  c = xor(mask, p);
  return(c);
}
```

We translate it into SSProve as JENC and use it in the following security game, supplying the random nonce r :

```
Definition JPRF_real id0 m :=
  k_val ← kgen ;;
  r ← sample uniform N ;;
  res ← JENC id0 k_val r m ;;
  ret (r, res)
```

We then prove it perfectly indistinguishability from a similar scheme CPRF_real which uses an intermediate, simplified SSProve encryption function, ENC, in place of JENC.

We establish the indistinguishability by applying Theorem 1 of the SSProve paper [1]. We thus have to find a stable invariant that is preserved by a run of each of these schemes and prove that they return equal values. We prove a slight generalization of the version that theorem. Before, the invariant was required to be stable w.r.t. the *finite sets* of locations used by the program. Moreover, these sets were assumed to be disjoint from the state of the adversary. We now only require the invariant to be stable w.r.t. some *arbitrary* sets of locations assumed to be disjoint from the state of the adversary. In particular, the sets can be infinite.

Thanks to this generalization we can apply the theorem when one of the programs is the output of our translation, since we do not have to provide the concrete set of locations used by the program, but instead we can use an infinite over-approximation. We thus obtain the following.

Theorem JPRF_perf_ind id : JPRF_real id $\approx_0$ CPRF_real.

We prove that CPRF_real is perfectly indistinguishable from PRF_real aes using the original SSProve Theorem 1 as we have better control over which locations are used.

Theorem CPRF_perf_ind : CPRF_real $\approx_0$ PRF_real aes.

Combining these two theorems, we get the following: the advantage of any adversary, which uses locations disjoint from JENC and from the intermediate encryption schemes, in distinguishing between JPRF_real and PRF_ENC is 0. This we can then combine with the result from the SSProve paper [1, Section 2.3] which states that PRF_ENC is IND-CPA secure up to the advantage of an adversary against aes as a PRF.

7 Related Work

The use of formal verification for cryptography has been intensely investigated, and Barbosa et al. [7] give an overview. More narrowly, work related to SSProve can be found in the extended version of the SSProve paper [25]. In this section, we survey the closest related work to ours in this space.

CertiCrypt [11] is the earliest framework for reasoning about cryptographic code in Coq, but is no longer maintained. FCF [32] is a more recent foundational Coq framework for

cryptographic proofs. It was used together with VST to verify the C implementations of HMAC in OpenSSL [14] and mbedTLS [41]. Our work is similar in that we prove the security and correctness of the Jasmin implementation of AES. While FCF could have been a reasonable option for us, we chose SSProve because it is under active development, uses the well-developed mathcomp [30] and mathcomp-analysis libraries [3], and supports modular proofs.

EasyCrypt [9, 10] is a proof assistant and verification tool specifically designed for game-based cryptographic proofs. Its good integration with automatic theorem provers (e.g., SMT solvers) is helpful for large proofs, even though it comes at a cost in terms of trusted computing base. The program logics of CertiCrypt and EasyCrypt come with native support for reasoning about function calls. This was not available in SSProve before and addressing this is one of the contributions of the present work (see Section 5.1).

In the fundamental ‘Last Mile’ paper [5] Jasmin programs are given semantics in Coq and the correctness of the Jasmin compiler is proved in Coq with respect to this semantics. As a realistic case study, they use EasyCrypt to prove the security and correctness of a Jasmin implementation of SHA3, relying on an unverified translation from Jasmin to EasyCrypt. In the present work, we bridge this gap by providing a verified translation from Jasmin to SSProve.

CryptHOL [12] is a foundational framework for game-based proofs that uses the theory of relational parametricity to achieve automation in Isabelle/HOL. However, unlike FCF and EasyCrypt, CryptHOL has so far not been used for the verification of efficient programs, as far as we are aware.

Schwabe et al. [37] prove the correctness of the C implementation of X25519 in TweetNaCl using VST. Protzenko et al. [34] verify an impressive library of cryptographic code in F*. Fiat-Crypto [21] is a foundational tool that can generate verified efficient implementations of finite field arithmetic. These works are focused on correctness though and do not consider cryptographic security.

Currently, there is no formal specification for the complete Rust language. The Hacspect semantics can be seen as a precise semantics for a non-controversial subset of Rust. Similar proposals, but for much larger subsets of Rust, include those of Ho and Protzenko [26] and Denis et al. [18].

8 Future Work

Jasminify [40] is a python tool that simplifies the process of calling Jasmin code from Rust. After compiling a program, the Rust object file is replaced with the Jasmin object file. However, Jasminify does not come with any correctness guarantees. We have shown how to prove the equivalence of a Rust (Hacspect) implementation for AES with a Jasmin program. Hacspect is expressive enough to implement high-level cryptographic protocols. For such protocols, we now have a safe way to replace its cryptographic primitives by optimized

Jasmin ones, as we know that their source-level semantics agree. As future work, one could try to test this toolchain, by using Jasminify, proving equivalence between the Hacspect and Jasmin implementations and then benchmarking to see what kind of performance gains one can achieve.

In concurrent work, libcrux [27] provides a library of verified implementations from different frameworks; and combines them with a safe Rust API. For example, it starts with a Hacspect reference implementation of HMAC and HKDF, and replaces their hash functions with optimized Jasmin implementations. It was proved [5] in EasyCrypt that the SHA3 implementation indeed implements a hash-function, but a formal connection with Hacspect is still missing. It would be exciting to use our framework to formally verify some of the replacements done in libcrux.

The Jasmin language is still under active development. In the present work, we devised a verified translation for the published version of the language [4]. It would be interesting to extend our work with language features that were added to Jasmin concurrently to our work.

Acknowledgements

We are very grateful to François Dupressoir for feedback on an earlier version of this article.

This work was in part supported by the Concordium Blockchain Research Center at Aarhus University, by a Villum Investigator grant (no. 25804), Center for Basic Research in Program Verification (CPV), from the VILLUM Foundation, by the German Federal Ministry of Education and Research BMBF (grant 16KISK038, project 6GEM), and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as part of the Excellence Strategy of the German Federal and State Governments – EXC 2092 CASA - 390781972.

References

- [1] Carmine Abate, Philipp G. Haselwarter, Exequiel Rivas, Antoine Van Muylder, Théo Winterhalter, Cătălin Hrițcu, Kenji Maillard, and Bas Spitters. 2021. SSProve: A Foundational Framework for Modular Cryptographic Proofs in Coq. (2021). <https://eprint.iacr.org/2021/397>
- [2] AbsInt. [n. d.]. Factsheet: CompCert C Compiler. Available at https://www.absint.com/factsheets/factsheet_compcert_c_web.pdf. https://www.absint.com/factsheets/factsheet_compcert_c_web.pdf
- [3] Reynald Affeldt, Cyril Cohen, Marie Kerjean, Assia Mahboubi, Damien Rouhling, Kazuhiko Sakaguchi, and Pierre-Yves Strub. 2021. mathcomp-analysis. Analysis library compatible with Mathematical Components. <https://github.com/math-comp/analysis>
- [4] José Bacelar Almeida, Manuel Barbosa, Gilles Barthe, Arthur Blot, Benjamin Grégoire, Vincent Laporte, Tiago Oliveira, Hugo Pacheco, Benedikt Schmidt, and Pierre-Yves Strub. 2017. Jasmin: High-Assurance and High-Speed Cryptography. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017, Dallas, TX, USA, October 30 - November 03, 2017*, Bhavani M. Thuraisingham, David Evans, Tal Malkin, and Dongyan Xu (Eds.). ACM, 1807–1823. <https://doi.org/10.1145/3133956.3134078>
- [5] José Bacelar Almeida, Manuel Barbosa, Gilles Barthe, Benjamin Grégoire, Adrien Koutsos, Vincent Laporte, Tiago Oliveira, and Pierre-Yves

- Strub. 2020. The Last Mile: High-Assurance and High-Speed Cryptographic Implementations. In *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 965–982.
- [6] Basavesh Ammanaghatta Shivakumar, Gilles Barthe, Benjamin Grégoire, Vincent Laporte, and Swarn Priya. 2022. Enforcing Fine-Grained Constant-Time Policies. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security (Los Angeles, CA, USA) (CCS '22)*. Association for Computing Machinery, New York, NY, USA, 83–96. <https://doi.org/10.1145/3548606.3560689>
- [7] Manuel Barbosa, Gilles Barthe, Karthikeyan Bhargavan, Bruno Blanchet, Cas Cremers, Kevin Liao, and Bryan Parno. 2019. SoK: Computer-Aided Cryptography. *IACR Cryptol. ePrint Arch.* 2019 (2019), 1393. <https://eprint.iacr.org/2019/1393>
- [8] Manuel Barbossa, François Dupressoir, Benjamin Grégoire, Vincent Laporte, Pierre-Yves Strub, and Tiago Oliveira. 2022. EasyCrypt and Jasmin Tutorial. https://formosa-crypto.gitlab.io/news/2022-06-07/sibenik_Sibenik.
- [9] Gilles Barthe, François Dupressoir, Benjamin Grégoire, César Kunz, Benedikt Schmidt, and Pierre-Yves Strub. 2013. EasyCrypt: A Tutorial. In *Foundations of Security Analysis and Design VII - FOSAD 2012/2013 Tutorial Lectures (Lecture Notes in Computer Science, Vol. 8604)*. Springer, 146–166. https://doi.org/10.1007/978-3-319-10082-1_6
- [10] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. 2011. Computer-Aided Security Proofs for the Working Cryptographer. In *CRYPTO (Lecture Notes in Computer Science, Vol. 6841)*. Springer, 71–90. https://doi.org/10.1007/978-3-642-22792-9_5
- [11] Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. 2009. Formal certification of code-based cryptographic proofs. In *POPL*. 90–101.
- [12] David A. Basin, Andreas Lochbihler, and S. Reza Sefidgar. 2020. CryptHOL: Game-Based Proofs in Higher-Order Logic. *J. Cryptol.* 33, 2 (2020), 494–566. <https://doi.org/10.1007/s00145-019-09341-z>
- [13] Mihir Bellare and Phillip Rogaway. 2004. Code-Based Game-Playing Proofs and the Security of Triple Encryption. *IACR Cryptol. ePrint Arch.* (2004), 331. <http://eprint.iacr.org/2004/331>
- [14] Lennart Beringer, Adam Petcher, Katherine Q. Ye, and Andrew W. Appel. 2015. Verified Correctness and Security of OpenSSL HMAC. In *24th USENIX Security Symposium*. USENIX Association, 207–221. <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/beringer>
- [15] Karthikeyan Bhargavan, Franziskus Kiefer, and Pierre-Yves Strub. 2018. hacspecc: Towards Verifiable Crypto Standards. In *Security Standardisation Research - 4th International Conference, SSR 2018, Darmstadt, Germany, November 26-27, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 11322)*, Cas Cremers and Anja Lehmann (Eds.). Springer, 1–20. https://doi.org/10.1007/978-3-030-04762-7_1
- [16] Barry Bond, Chris Hawblitzel, Manos Kapritsos, K. Rustan M. Leino, Jacob R. Lorch, Bryan Parno, Ashay Rane, Srinath T. V. Setty, and Laure Thompson. 2017. Vale: Verifying High-Performance Cryptographic Assembly Code. In *26th USENIX Security Symposium, USENIX Security 2017, Vancouver, BC, Canada, August 16-18, 2017*, Engin Kirda and Thomas Ristenpart (Eds.). USENIX Association, 917–934. <https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/bond>
- [17] Chris Brzuska, Antoine Delignat-Lavaud, Cédric Fournet, Konrad Kohbrok, and Markulf Kohlweiss. 2018. State Separation for Code-Based Game-Playing Proofs. In *ASIACRYPT*. Springer International Publishing, Cham, 222–249. <https://eprint.iacr.org/2018/306>
- [18] Xavier Denis, Jacques-Henri Jourdan, and Claude Marché. 2022. Creusot: A Foundry for the Deductive Verification of Rust Programs. In *Formal Methods and Software Engineering*, Adrian Riesco and Min Zhang (Eds.). Springer, 90–105.
- [19] Jason A. Donenfeld. [n. d.]. WireGuard: Formal Verification. Available at <https://www.wireguard.com/formal-verification/>. <https://www.wireguard.com/formal-verification/>
- [20] Morris Dworkin, Elaine Barker, James Nechvatal, James Foti, Lawrence Bassham, E. Roback, and James Dray. 2001. Advanced Encryption Standard (AES). <https://doi.org/10.6028/NIST.FIPS.197>
- [21] A. Erbsen, J. Philipoom, J. Gross, R. Sloan, and A. Chlipala. 2019. Simple High-Level Code for Cryptographic Arithmetic - With Proofs, Without Compromises. In *IEEE S&P*. <https://doi.org/10.1109/SP.2019.00005>
- [22] Armando Faz-Hernandez, Sam Scott, Nick Sullivan, Riad S. Wahby, and Christopher A. Wood. 2022. Hashing to Elliptic Curves. Internet-Draft draft-irtf-cfrg-hash-to-curve-16. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-irtf-cfrg-hash-to-curve/16/> Work in Progress.
- [23] Shay Gueron. 2012. White Paper: Intel® Advanced Encryption Standard (AES) New Instructions Set. <https://www.intel.com/content/www/us/en/developer/articles/tool/intel-advanced-encryption-standard-aes-instructions-set.html>
- [24] Shai Halevi. 2005. A plausible approach to computer-aided cryptographic proofs. *IACR Cryptol. ePrint Arch.* (2005), 181. <http://eprint.iacr.org/2005/181>
- [25] Philipp G. Haselwarter, Exequiel Rivas, Antoine Van Muylder, Théo Winterhalter, Carmine Abate, Nikolaj Sidorenko, Cătălin Hrițcu, Kenji Maillard, and Bas Spitters. 2023. SSProve: A Foundational Framework for Modular Cryptographic Proofs in Coq. *ACM Trans. Program. Lang. Syst.* 45, 3, Article 15 (jul 2023), 61 pages. <https://doi.org/10.1145/3594735>
- [26] Son Ho and Jonathan Protzenko. 2022. Aeneas: Rust Verification by Functional Translation. *Proc. ACM Program. Lang.* 6, ICFP, Article 116 (2022), 31 pages. <https://doi.org/10.1145/3547647>
- [27] Franziskus Kiefer, Karthikeyan Bhargavan, Lucas Franceschino, Denis Merigoux, Lasse Letager Hansen, Bas Spitters, Manuel Barbosa, Antoine Séré, and Pierre-Yves Strub. 2023. HACSPEC: a gateway to high-assurance cryptography. In *RWC23*.
- [28] Xavier Leroy, Sandrine Blazy, Daniel Kästner, Bernhard Schommer, Markus Pister, and Christian Ferdinand. 2016. CompCert – a formally verified optimizing compiler. In *ERTS 2016: Embedded Real Time Software and Systems, 8th European Congress*.
- [29] Andreas Lochbihler, S. Reza Sefidgar, David A. Basin, and Ueli Maurer. 2019. Formalizing Constructive Cryptography using CryptHOL. In *CSF. IEEE*, 152–166. <https://doi.org/10.1109/CSF.2019.00018>
- [30] Assia Mahboubi and Enrico Tassi. 2021. Mathematical components. Online book. <https://math-comp.github.io/mcb/>
- [31] Denis Merigoux, Franziskus Kiefer, and Karthikeyan Bhargavan. 2021. hacspecc: succinct, executable, verifiable specifications for high-assurance cryptography embedded in Rust. Technical Report. Inria. <https://hal.inria.fr/hal-03176482>
- [32] Adam Petcher and Greg Morrisett. 2015. The Foundational Cryptography Framework. In *Principles of Security and Trust - 4th International Conference, POST 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015, Proceedings (Lecture Notes in Computer Science, Vol. 9036)*, Riccardo Focardi and Andrew C. Myers (Eds.). Springer, 53–72. https://doi.org/10.1007/978-3-662-46666-7_4
- [33] Amir Pnueli, Michael Siegel, and Eli Singerman. 1998. Translation Validation. In *Tools and Algorithms for Construction and Analysis of Systems, 4th International Conference, TACAS '98, Held as Part of the European Joint Conferences on the Theory and Practice of Software, ETAPS '98, Lisbon, Portugal, March 28 - April 4, 1998, Proceedings (Lecture Notes in Computer Science, Vol. 1384)*, Bernhard Steffen (Ed.). Springer, 151–166. <https://doi.org/10.1007/BFb0054170>
- [34] Jonathan Protzenko and Bryan Parno. 2019. EverCrypt cryptographic provider offers developers greater security assurances. Microsoft Research Blog. <https://www.microsoft.com/en->

- [us/research/blog/evercrypt-cryptographic-provider-offers-developers-greater-security-assurances/](https://research.blog/evercrypt-cryptographic-provider-offers-developers-greater-security-assurances/)
- [35] Mikkel Milo Rasmus Holdsbjerg-Larsen, Bas Spitters. 2022. A Verified Pipeline from a Specification Language to Optimized, Safe Rust. CoqPL. <https://cs.au.dk/~spitters/CoqPL22.pdf>
 - [36] Mike Rosulek. 2021. The Joy of Cryptography. Online textbook. <http://web.engr.oregonstate.edu/~rosulekm/crypto/>
 - [37] Peter Schwabe, Benoît Viguier, Timmy Weerwag, and Freek Wiedijk. 2021. A Coq proof of the correctness of X25519 in TweetNaCl. In *2021 34th CSF*. 1–16. <https://doi.org/10.1109/CSF51468.2021.00023>
 - [38] Victor Shoup. 2004. Sequences of games: a tool for taming complexity in security proofs. *IACR Cryptol. ePrint Arch.* (2004), 332. <http://eprint.iacr.org/2004/332>
 - [39] Laurent Simon, David Chisnall, and Ross Anderson. 2018. What you get is what you C: Controlling side effects in mainstream C compilers. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 1–15.
 - [40] Juriaan van Drunen. 2021. Calling Jasmin from Rust. <https://gitlab.com/Jur/jasminify>
 - [41] Katherine Q. Ye, Matthew Green, Naphat Sanguansin, Lennart Beringer, Adam Petcher, and Andrew W. Appel. 2017. Verified Correctness and Security of mbedTLS HMAC-DRBG. In *CCS'17*. ACM, 2007–2020. <https://doi.org/10.1145/3133956.3133974>
 - [42] Jean-Karim Zinzindohoué, Karthikeyan Bhargavan, Jonathan Protzenko, and Benjamin Beurdouche. 2017. HACl*: A Verified Modern Cryptographic Library. In *ACM Conference on Computer and Communications Security*. ACM, 1789–1806. <http://eprint.iacr.org/2017/536>

Received 2023-09-19; accepted 2023-11-25