

Formal Certification of Code-Based Cryptographic Proofs

Gilles Barthe^{1,2} Benjamin Grégoire^{1,3} Santiago Zanella^{1,3}

¹ Microsoft Research - INRIA Joint Centre, France

² IMDEA Software, Madrid, Spain ³ INRIA Sophia Antipolis - Méditerranée, France

Gilles.Barthe@imdea.org {Benjamin.Gregoire,Santiago.Zanella}@sophia.inria.fr

Abstract

As cryptographic proofs have become essentially unverifiable, cryptographers have argued in favor of developing techniques that help tame the complexity of their proofs. Game-based techniques provide a popular approach in which proofs are structured as sequences of games, and in which proof steps establish the validity of transitions between successive games. Code-based techniques form an instance of this approach that takes a code-centric view of games, and that relies on programming language theory to justify proof steps. While code-based techniques contribute to formalize the security statements precisely and to carry out proofs systematically, typical proofs are so long and involved that formal verification is necessary to achieve a high degree of confidence. We present CertiCrypt, a framework that enables the machine-checked construction and verification of code-based proofs. CertiCrypt is built upon the general-purpose proof assistant Coq, and draws on many areas, including probability, complexity, algebra, and semantics of programming languages. CertiCrypt provides certified tools to reason about the equivalence of probabilistic programs, including a relational Hoare logic, a theory of observational equivalence, verified program transformations, and game-based techniques such as reasoning about failure events. The usefulness of CertiCrypt is demonstrated through classical examples, including a proof of semantic security of OAEP (with a bound that improves upon [9]), and a proof of existential unforgeability of FDH signatures. Our work provides a first yet significant step towards Halevi's ambitious programme [21] of providing tool support for cryptographic proofs.

1. Introduction

Provable security [34], whose origins can be traced back to the pioneering work of Goldwasser and Micali [20], advocates a mathematical approach based on complexity theory in which the goals and requirements of cryptosystems are specified precisely, and where security proofs are carried out rigorously and make all underlying assumptions explicit. In a typical provable security setting, one reasons about effective adversaries, modeled as arbitrary probabilistic polynomial-time Turing machines, and about their probability of thwarting a security objective, e.g. secrecy; in general, provable security statements do not refer directly to the probability of the adversary breaking security, but to its advantage over a *blind*, uninformed adversary. In a similar fashion, security assumptions about cryptographic primitives bound the probability of polynomial algorithms to solve hard problems, e.g. computing discrete logarithms. The security proof is performed by reduction by showing that the existence of an effective adversary with a certain advantage in breaking security implies the existence of an effective algorithm contradicting the security assumptions. Although the adoption of provable security has significantly enhanced confidence in security proofs, several published proofs have been found

incorrect (cf. [31]), and the cryptographic community is increasingly wary that the field may be approaching a crisis of rigor [9, 21].

The game-playing technique [9, 21, 32] is a general method to structure and unify cryptographic proofs, thus making them less error-prone. Its central idea is to view the interaction between an adversary and the cryptosystem as a game, and to study transformations that preserve security. In a typical game-based proof, one considers transitions of the form $G, A \rightarrow^h G', A'$, where G and G' are games, A and A' are events, and h is a monotonic function such that $\Pr_G[A] \leq h(\Pr_{G'}[A'])$. One can obtain an upper bound for the probability of an event A_0 in some initial game G_0 by successively refining G_0, A_0 into a game/event pair G_n, A_n ,

$$G_0, A_0 \rightarrow^{h_1} G_1, A_1 \rightarrow \dots \rightarrow^{h_n} G_n, A_n$$

and then bounding the probability of event A_n in G_n .

The game-playing technique is widely applicable, and has been extensively used. Code-based techniques [9] is an instance of the game-playing technique whose distinguishing feature is to take a code-centric view of games, security hypotheses and computational assumptions, that are expressed using (probabilistic, imperative, polynomial-time) programs. Under this view, game transformations become program transformations, and can be justified rigorously by semantic means; in particular, many transformations can be viewed as common program optimizations, and are justified by proving that the original and transformed programs are observationally equivalent. Although code-based proofs are easier to verify, they go far beyond established theories of program equivalence and exhibit a surprisingly rich and broad set of reasoning principles that draws on program verification, algebraic reasoning, and probability and complexity theory. Thus, despite the beneficial effect of their underlying framework, code-based proofs remain inherently complex. Whereas Bellare and Rogaway [9] already observed that code-based proofs could be more easily amenable to machine-checking, Halevi [21] argued that formal verification techniques should be used to improve trust in cryptographic proofs, and set up a programme for building a tool that could be used by the cryptographic community to mechanize their proofs.

This article reports on a first yet significant step towards Halevi's programme. We describe CertiCrypt, a framework to construct machine-checked code-based proofs in the Coq proof assistant [35], supporting:

Faithful and rigorous encoding of games. In order to be readily accessible to cryptographers, we have chosen a formalism that is commonly used to describe games. Concretely, the lowest layer of CertiCrypt is the formalization of pWHILE, an imperative programming language with random assignments, structured datatypes, and procedure calls. We provide a deep and dependently-typed embedding of the syntax; thanks to dependent types, the typability of pWHILE programs is obtained for free. We also provide a small-step operational semantics us-

ing the distribution monad of Audebaud and Paulin [4]. The semantics is instrumented to calculate the cost of running programs; this offers the means to define complexity classes, and in particular to define formally the notion of effective (probabilistic polynomial-time) adversary. In addition, we also model nonstandard features, such as policies on variable accesses and procedure calls, and use them to formulate the notion of well-formed adversary, which captures many assumptions left informal in cryptographic proofs.

Exact security. Many security proofs establish an asymptotic behavior for adversaries and show that the advantage of any effective adversary is negligible w.r.t. a security parameter (which typically determines the length of keys or messages). However, the cryptographic community is increasingly focused on exact security, a much more useful result since it gives hints as to how to choose system parameters in practice to satisfy a security guarantee. The goal of exact security is to provide concrete bounds both for the advantage of the adversary and for its execution time. CertiCrypt supports the former (but for the time being, not the latter)¹.

Full and independently verifiable proofs. CertiCrypt adopts a formal semanticist perspective and goes beyond Halevi’s vision in two respects. First, it provides a unified framework to carry out full proofs; all intermediate steps of reasoning can be justified formally, including complex side conditions that justify the correctness of transformations (about probabilities, groups, polynomials, etc). Second, one notable feature of Coq, and thus CertiCrypt, is to support independent verifiability of proofs, which is an important motivation behind game-based proofs. More concretely, every proof is represented by a proof object, that can be checked automatically by a (small and trustworthy) proof checking engine. In order to trust a cryptographic proof, one only needs to check its statement, and not its details.

Powerful and automated reasoning methods. CertiCrypt formalizes a Relational Hoare Logic and a theory of observational equivalence, and uses them as stepping stones to support the main tools of code-based reasoning through certified, reflective tactics. In particular, CertiCrypt shows that many transformations used in code-based proofs, including common optimizations, are semantics-preserving. One of its specific contributions is to prove formally the correctness of a variant of lazy sampling, which is used ubiquitously in cryptographic proofs. In addition, CertiCrypt supports methods based on failure events (the so-called fundamental lemma of game-playing).

We have successfully conducted nontrivial case studies that validate our design, show the feasibility of formally verifying cryptographic proofs, and conform the plausibility of Halevi’s programme.

Contents The purpose of this article is to provide an overview of the CertiCrypt project, and to stir further interest in machine-checked cryptographic proofs. Additional details on design choices, on formalizing the semantics of probabilistic programs, and on case studies, will be provided elsewhere². In consequence, the paper is organized as follows: we begin in Section 2 with two introductory examples of game-based proofs, namely the semantic security of ElGamal encryption, and the PRP/PRF switching lemma; in Section 3 we introduce the language we use to represent games and its semantics, and we discuss the notions of complexity and termination; in Section 4.1 we present a probabilistic relational Hoare

logic that forms the core of our framework; in Sections 5 and 6 we overview the formulation and automation of game transformations in CertiCrypt; in Section 7 we report on two significant case studies we have formalized in CertiCrypt: existential unforgeability of the FDH signature scheme, and semantic security of OAEP; we finish with a discussion of related work and concluding remarks.

2. Basic examples

This section illustrates the principles of CertiCrypt on two basic examples of game-based proofs: semantic security of ElGamal encryption and the PRP/PRF switching lemma. The language used to represent games is formally introduced in the next section. We begin with some basic definitions.

The Random Oracle Model is a model of cryptography extensively used in security proofs in which some cryptographic primitives, e.g. hash functions, are assumed to be indistinguishable from random functions (despite the fact that no real function can implement a truly random function [14]). Such primitives are modeled by oracles that return random values in response to queries. The sole condition is that queries are answered consistently: if some value is queried twice, the same response must be given. Our formalism captures the notion of random oracle using stateful procedures that store queries and their results, e.g.

Oracle $\mathcal{O}(x)$: if $x \notin \text{dom}(L)$ then $y \leftarrow_{\$} \{0, 1\}^\eta$; $L \leftarrow (x, y) :: L$;
return $L[x]$

An asymmetric encryption scheme is composed of three algorithms: key generation $\text{KG}(\eta)$, where η is the security parameter; encryption $\text{Enc}(pk, m)$ where pk is a public key and m a plaintext; and decryption—not relevant here. An asymmetric encryption scheme is said to be semantically secure (equivalently, IND-CPA secure) if it is infeasible to gain significant information about a plaintext given only a corresponding ciphertext and the public key. This is formally defined using the following game, where \mathcal{A} and \mathcal{A}' are allowed to share state via global variables and thus are regarded as a single adaptive adversary:

Game IND-CPA :
 $(sk, pk) \leftarrow \text{KG}(\eta)$;
 $(m_0, m_1) \leftarrow \mathcal{A}(pk)$;
 $b \leftarrow_{\$} \{0, 1\}$; $\gamma \leftarrow \text{Enc}(pk, m_b)$;
 $b' \leftarrow \mathcal{A}'(pk, \gamma)$

The game first generates a new key pair and gives the public key to the adversary, who returns two plaintexts m_0, m_1 of his choice. Then, the challenger tosses a fair coin b and gives the encryption of m_b back to the adversary, whose goal is to guess which message has been encrypted. The scheme is IND-CPA if for every effective adversary $\mathcal{A}, \mathcal{A}'$, $|\text{Pr}_{\text{IND-CPA}}[b = b'] - \frac{1}{2}|$ is negligible in the security parameter, i.e. the adversary cannot do much better than a blind guess. Formally, a function $\nu : \mathbb{N} \rightarrow \mathbb{R}$ is negligible iff $\forall c. \exists n_c. \forall n. n \geq n_c \Rightarrow |\nu(n)| \leq n^{-c}$.

2.1 The ElGamal encryption scheme

ElGamal [18] is a widely used asymmetric encryption scheme, and an emblematic example of game-based proofs, as it embodies many of the techniques described in Sections 4 and 5. The proof follows [32]; all games are defined in Fig. 1.

Given a cyclic group of order q , and a generator g , we define:³

- Key generation: $\text{KG}() \stackrel{\text{def}}{=} x \leftarrow_{\$} \mathbb{Z}_q$; return (x, g^x)
- Encryption: $\text{Enc}(\alpha, m) \stackrel{\text{def}}{=} y \leftarrow_{\$} \mathbb{Z}_q$; return $(g^y, \alpha^y \times m)$

ElGamal is IND-CPA secure under the Decisional Diffie-Hellman (DDH) assumption, which states that it is hard to distinguish be-

¹The limitation is reflected in the statements of our case studies.

²For review purposes, the complete development is available from <http://www-sop.inria.fr/everest/certcrypt-popl.tar.gz>

³The security parameter, implicit in this presentation, determines this cyclic group by indexing a family of groups where the DDH problem is believed intractable.

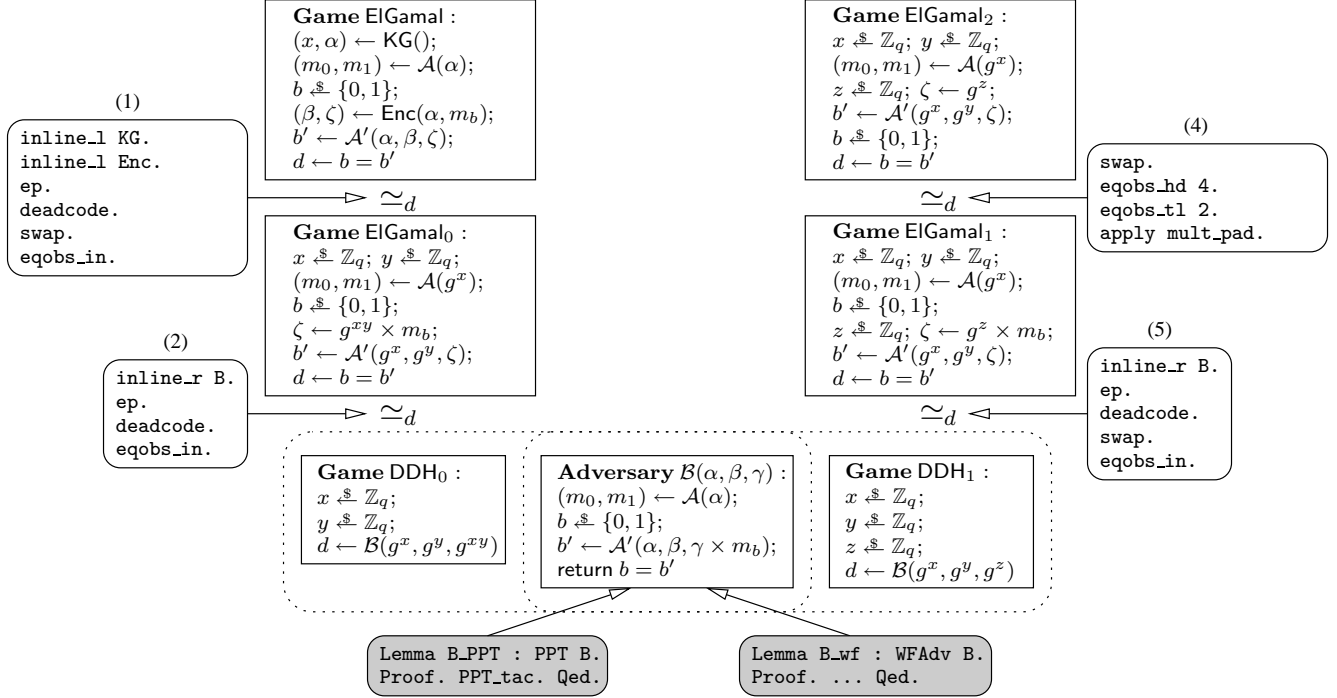


Figure 1. Code-based proof of ElGamal semantic security

tween triples of the form (g^x, g^y, g^{xy}) and (g^x, g^y, g^z) where x, y, z are uniformly sampled in \mathbb{Z}_q . In our setting, DDH is formulated precisely by stating that for any polynomial-time and well-formed adversary \mathcal{B} , $|\Pr_{\text{DDH}_0}[d] - \Pr_{\text{DDH}_1}[d]|$ is negligible in the security parameter. Figure 1 presents a high level view of the proof: the square boxes represent games, whereas the rounded boxes represent proof sketches of the transitions between games; the tactics that appear in these boxes hopefully have self-explanatory names, but are explained in more detail in Section 5. The rounded grey boxes represent proof sketches of side conditions that guarantee that the DDH assumption is correctly applied. The proof proceeds by constructing an adversary \mathcal{B} against DDH such that the distribution of $b = b'$ (i.e. d) after running the IND-CPA game ElGamal is exactly the same as the distribution of d after running DDH₀. Furthermore we show that the probability of d being true in DDH₁ is $\frac{1}{2}$ for the same adversary \mathcal{B} . The proof is summarized by the following equations:

$$\begin{aligned}
|\Pr_{\text{ElGamal}}[b = b'] - \frac{1}{2}| &= |\Pr_{\text{ElGamal}_0}[d] - \frac{1}{2}| & (1) \\
&= |\Pr_{\text{DDH}_0}[d] - \frac{1}{2}| & (2) \\
&= |\Pr_{\text{DDH}_0}[d] - \Pr_{\text{ElGamal}_2}[d]| & (3) \\
&= |\Pr_{\text{DDH}_0}[d] - \Pr_{\text{ElGamal}_1}[d]| & (4) \\
&= |\Pr_{\text{DDH}_0}[d] - \Pr_{\text{DDH}_1}[d]| & (5)
\end{aligned}$$

Equation (1) is justified because ElGamal and ElGamal₀ induce the same distribution on d (ElGamal \simeq_d ElGamal₀). To prove this, we inline the calls to KG and Enc, and then perform expression propagation and dead code elimination (ep, deadcode). At this point we are left with two almost equal games, except the sampling of y is done later in one game than in the other. The tactic swap is used to hoist instructions whenever is possible in order to obtain a common prefix, and allows us to hoist the sampling of y to the right place. We conclude by applying eqobs_in that decides observational equivalence of a program with itself. Equations (2) and (5) are obtained similarly, while (3) holds because b' is independent

from the sampling of b in ElGamal₂. Finally, to prove equation (4) we begin by removing the common part of the two games with the exception of the instruction $z \stackrel{\$}{\leftarrow} \mathbb{Z}_q$ (eqobs_hd, eqobs_tl). We then apply an algebraic property of cyclic groups (mult_pad): when multiplying a uniformly distributed element of the group by another element, the result is uniformly distributed. This allows to prove that $z \stackrel{\$}{\leftarrow} \mathbb{Z}_q; \zeta \leftarrow g^z \times m_b$ and $z \stackrel{\$}{\leftarrow} \mathbb{Z}_q; \zeta \leftarrow g^z$ induce the same distribution on ζ .

The proof concludes by applying the DDH assumption. We prove that the adversary \mathcal{B} is strict probabilistic polynomial-time and well-formed (under the assumption that \mathcal{A} and \mathcal{A}' are so). The proof of the former condition is automated in CertiCrypt.

2.2 The PRP/PRF switching lemma

In cryptographic proofs, particularly those dealing with blockciphers, it is often convenient to replace a pseudo-random permutation (PRP) by a pseudo-random function (PRF). The PRP/PRF switching lemma establishes that such a replacement does not change significantly the advantage of an effective adversary. In a code-based setting, the Switching Lemma states that

$$|\Pr_{G_{\text{PRP}}}[d] - \Pr_{G_{\text{PRF}}}[d]| \leq \frac{q(q-1)}{2^{\eta+1}}$$

where the games G_{PRP} and G_{PRF} give the adversary access to an oracle that represents a random permutation and a random function respectively, and where q bounds the number of oracle queries made by \mathcal{A} .

The proof is split in two parts: the first part formalizes the intuition that the probability of the adversary outputting a given value is the same if a PRP is replaced by a PRF and no collisions are observed; it uses the Fundamental Lemma of game-playing (Lemma 2 in Sec. 6). The second part provides an upper bound to the probability of a collision; it uses lazy sampling (Lemma 1 in Sec. 5.2).

3.2 Operational semantics

Programs in pWHILE are given a small-step semantics using the distribution monad $M(X)$, whose type constructor is defined as

$$M(X) \stackrel{\text{def}}{=} (X \rightarrow [0, 1]) \rightarrow [0, 1]$$

and whose operators unit and bind are defined as

$$\begin{aligned} \text{unit} &: X \rightarrow M(X) \stackrel{\text{def}}{=} \lambda x. \lambda f. f x \\ \text{bind} &: M(X) \rightarrow (X \rightarrow M(Y)) \rightarrow M(Y) \\ &\stackrel{\text{def}}{=} \lambda \mu. \lambda M. \lambda f. \mu(\lambda x. M x f) \end{aligned}$$

The monad $M(X)$ was proposed by Audebaud and Paulin [4] as a variant of the expectation monad used by Ramsey and Pfeffer [29], and builds on earlier work by Kozen [24]. The formalization of the semantics heavily relies on Paulin’s axiomatization in Coq of the $[0, 1]$ real interval—for our purposes, it has been necessary to add division to the library.

The semantics of commands and expressions are defined relative to a given memory, i.e. a mapping from variables to values. We let \mathcal{M} denote the set of memories. Expressions are deterministic; their semantics is standard and given by a function $\llbracket \cdot \rrbracket_{\text{expr}}$, that evaluates an expression in a given memory and returns a value. The semantics of distribution expressions is given by a function $\llbracket \cdot \rrbracket_{\text{distr}}$. For a distribution expression d of type T , we have that $\llbracket d \rrbracket_{\text{distr}} : \mathcal{M} \rightarrow M(X)$, where X is the interpretation of type T . For instance, in the previous section we have used $\{0, 1\}^\eta$ to denote the uniform distribution on bit-strings of length η (the security parameter), formally, we have $\llbracket \{0, 1\}^\eta \rrbracket_{\text{distr}} \stackrel{\text{def}}{=} \lambda m f. \sum_{bs \in \{0, 1\}^\eta} \frac{1}{2^\eta} f(bs)$. Thanks to dependent types, the semantics of expressions and distribution expressions is total. In the following, and whenever there is no confusion, we will drop the subscripts in $\llbracket \cdot \rrbracket_{\text{expr}}$ and $\llbracket \cdot \rrbracket_{\text{distr}}$.

The semantics of commands relates a deterministic state to a (sub-)probability distribution over deterministic states and uses a frame stack to deal with procedure calls. Formally, a deterministic state is a triple consisting of the current command $c : \mathcal{C}$, a memory $m : \mathcal{M}$, and a frame stack $F : \text{list frame}$. We let \mathcal{S} denote the set of deterministic states. One step execution $\llbracket \cdot \rrbracket^1 : \mathcal{S} \rightarrow M(\mathcal{S})$ is defined by the rules of Fig. 3. In the figure, we use $a \rightsquigarrow b$ as a notation for $\llbracket a \rrbracket^1 = b$ and loc and glob to project memories on local and global variables respectively.

We briefly comment on the transition rules for calling a procedure (3rd rule) and returning from a call (2nd rule). Upon a call, a new frame is appended to the stack, containing the destination variable, the return expression of the called procedure, the continuation to the call, and the local memory of the caller. The state resulting from the call contains the body of the called procedure, the global part of the memory, a local memory initialized to map the formal parameters to the value of the actual parameters just before the call, and the updated stack. When returning from a call with a non-empty stack, the top frame is popped, the return expression is evaluated and the resulting value is assigned to the destination variable after previously restoring the local memory of the caller; the continuation taken from the frame becomes the current command. If the stack is empty when returning from a call, the execution of the program terminates and the final state is embedded into the monad using the unit operator.

Using the monadic constructions, one can define an n -step execution function $\llbracket \cdot \rrbracket_n$:

$$\llbracket S \rrbracket_0 \stackrel{\text{def}}{=} \text{unit } S \quad \llbracket S \rrbracket_{n+1} \stackrel{\text{def}}{=} \text{bind } \llbracket S \rrbracket_n \llbracket \cdot \rrbracket^1$$

Finally, the denotation of a command c in an initial memory m is defined to be the (limit) distribution of reachable final memories:

$$\llbracket c \rrbracket m : M(\mathcal{M}) \stackrel{\text{def}}{=} \lambda f. \sup \{ \llbracket (c, m, []) \rrbracket_n f |_{\text{final}} \mid n \in \mathbb{N} \}$$

where $f|_{\text{final}} : \mathcal{S} \rightarrow [0, 1]$ is the function that when applied to a state (c, m, F) gives $f(m)$ if it is a final state and 0 otherwise. Since the sequence $\llbracket (c, m, []) \rrbracket_n f|_{\text{final}}$ is increasing and upper bounded by 1, this least upper bound always exists and corresponds to the limit of the sequence.

We have shown that the semantics is discrete, which we use to apply a variant of Fubini’s theorem for proving the rules [R-Comp] and [R-Trans] in the next section.

Computing probabilities The advantage of using this monadic semantics is that, if we use an arbitrary function as a continuation to the denotation of a program, what we get (for free) as a result is its expected value w.r.t. the distribution of final memories. In particular, we can compute the probability of an event A in the distribution obtained after executing a command c in an initial memory m by measuring its characteristic function $\mathbb{1}_A : \text{Pr}_{c, m}[A] \stackrel{\text{def}}{=} \llbracket c \rrbracket m \mathbb{1}_A$. For instance, one can verify that the denotation of $x \stackrel{\$}{\leftarrow} [0..1]; y \stackrel{\$}{\leftarrow} [0..1]$ in the memory m is

$$\begin{aligned} &\lambda f. \frac{1}{4} (f(m\{0, 0/x, y\}) + f(m\{0, 1/x, y\}) \\ &\quad + f(m\{1, 0/x, y\}) + f(m\{1, 1/x, y\})) \end{aligned}$$

and conclude that the probability of the event $x \leq y$ after executing the command above is $\frac{3}{4}$.

3.3 Probabilistic polynomial-time programs

In general, cryptographic proofs reason about effective adversaries, which can only use a polynomially bounded number of resources. The complexity notion that captures this intuition, and which is pervasive in cryptographic proofs, is that of *strict probabilistic polynomial-time* complexity. Concretely, a program is said to be strict probabilistic polynomial-time (PPT) whenever there exists a polynomial bound (on some security parameter η) on the cost of each possible execution, regardless of the outcome of its coin tosses. Otherwise said, a probabilistic program is PPT whenever the same program, seen as a non-deterministic program, terminates and the cost of each possible run is bounded by a polynomial.

Termination and efficiency are orthogonal. Consider, for instance, the following two programs:

$$\begin{aligned} &b \leftarrow \text{true}; \text{while } b \text{ do } b \stackrel{\$}{\leftarrow} \{0, 1\} \\ &b \stackrel{\$}{\leftarrow} \{0, 1\}; \text{if } b \text{ then while true do nil} \end{aligned}$$

The former terminates with probability 1 (it terminates within n iterations with probability $1 - 2^{-n}$), but may take an arbitrarily large number of iterations to terminate. The latter terminates with probability $\frac{1}{2}$, but when it does, it takes only a constant time. We deal with termination and efficiency separately.

Definition 2 (Termination). *The probability that a program c terminates starting from an initial memory m is $\llbracket c \rrbracket m \mathbb{1}_{\text{true}}$. We say that a program c is absolutely terminating, and note it $\text{Lossless}(c)$, iff it terminates with probability 1 in any initial memory.*

To deal with efficiency, we non-intrusively instrument the semantics of our language to compute the cost of running a program. The instrumented semantics ranges over $M(\mathcal{M} \times \mathbb{N})$ instead of simply $M(\mathcal{M})$. We recall that our semantics is implicitly parametrized by a security parameter η , on which we base our notion of complexity.

Definition 3 (Polynomially bounded distribution). *We say that a distribution $\mu : M(\mathcal{M} \times \mathbb{N})$ is (p, q) -bounded, where p and q are polynomials, and note it $\text{bounded}(p, q, \mu)$, whenever for every (m, n) occurring with non-zero probability in μ , the size of every value in the memory m is bounded by $p(\eta)$ and $n \leq q(\eta)$. This notion is formally defined by means of the range predicate introduced in Sec. 4.*

$$\begin{array}{l}
(\text{nil}, m, []) \rightsquigarrow \text{unit}(\text{nil}, m, []) \\
(\text{nil}, m, (x, e, c, l) :: F) \rightsquigarrow \text{unit}(c, (l, m.\text{glob})\{\llbracket e \rrbracket m/x\}, F) \\
(x \leftarrow p(\vec{e}); c, m, F) \rightsquigarrow \text{unit}(E(p).\text{body}, (\emptyset\{\llbracket \vec{e} \rrbracket m/E(p).\text{params}\}, m.\text{glob}), (x, E(p).\text{re}, c, m.\text{loc}) :: F) \\
(\text{if } e \text{ then } c_1 \text{ else } c_2; c, m, F) \rightsquigarrow \text{unit}(c_1; c, m, F) \quad \text{if } \llbracket e \rrbracket m = \text{true} \\
(\text{if } e \text{ then } c_1 \text{ else } c_2; c, m, F) \rightsquigarrow \text{unit}(c_2; c, m, F) \quad \text{if } \llbracket e \rrbracket m = \text{false} \\
(\text{while } e \text{ do } c; c', m, F) \rightsquigarrow \text{unit}(c; \text{while } e \text{ do } c; c', m, F) \quad \text{if } \llbracket e \rrbracket m = \text{true} \\
(\text{while } e \text{ do } c; c', m, F) \rightsquigarrow \text{unit}(c', m, F) \quad \text{if } \llbracket e \rrbracket m = \text{false} \\
\\
(x \leftarrow e; c, m, F) \rightsquigarrow \text{unit}(c, m\{\llbracket e \rrbracket m/x\}, F) \\
(x \stackrel{\$}{\leftarrow} d; c, m, F) \rightsquigarrow \text{bind}(\llbracket d \rrbracket m)(\lambda v. \text{unit}(c, m\{v/x\}, F))
\end{array}$$

Figure 3. Probabilistic semantics of pWHILE programs

Definition 4 (Strict probabilistic polynomial-time program). A program c is *strict probabilistic polynomial-time*, and we note it $\text{PPT}(c)$, iff it *terminates absolutely*, and there exist polynomial transformers F, G such that for every (p, q) -bounded distribution μ , the distribution $(\text{bind } \mu \llbracket c \rrbracket)$ is $(F(p), q + G(p))$ -bounded.

We can recover some intuition by taking $\mu = \text{unit}(m, 0)$ in the above definition. In this case, we may paraphrase the condition as follows: if the size of values in m is bounded by some polynomial p , and an execution of the program in m terminates with non-zero probability in memory m' , then the size of values in m' is bounded by the polynomial $F(p)$, and the cost of the execution is bounded by the polynomial $G(p)$. It is in this latter polynomial that bounds the cost of executing the program that we are ultimately interested. The increased complexity in the definition is required for proving compositional results (e.g. the sequential composition of two PPT programs results in a PPT program).

Although our formalizations for termination and efficiency rely on semantic definitions, it is not necessary for users to reason directly about the semantics of a program to prove it meets those definitions. CertiCrypt implements a certified algorithm showing that every program without loops and recursive calls is lossless.⁴ CertiCrypt also provides another algorithm that, together with the first, establishes that a program is PPT provided that, additionally, the program does not contain expressions that might generate values of superpolynomial size or take a superpolynomial time when evaluated in a polynomially bounded memory.

3.4 Adversaries

In order to reason formally about security, we make explicit which variables and procedures are accessible to adversaries, and provide a simple analysis to check whether an adversary respects its policy. Given a set of procedure identifiers \mathcal{O} (the procedures that may be called by the adversary), and sets of global variables \mathcal{G}_A (those that can be read and written by the adversary) and \mathcal{G}_{ro} (those that the adversary can only read), we say that an adversary \mathcal{A} is well-formed in an environment E if the judgment $\vdash_{\text{wf}} \mathcal{A}$ can be derived using the rules in Fig. 4. These rules guarantee that each time a variable is written by the adversary, the adversary has the right to do so; and that each time a variable is read by the adversary, it is either a global variable in $\mathcal{G}_A \cup \mathcal{G}_{ro}$ or a local variable previously initialized. A well-formed adversary is free to call oracles, but any other procedure it calls must be a well-formed adversary itself.

⁴It is of course a weak result in terms of termination of probabilistic programs, but nevertheless sufficient as regards cryptographic applications. Extending our formalization to a certified termination analysis for loops is interesting, but orthogonal to our main goals, and left for future work.

$$\begin{array}{c}
I \vdash \text{nil} : I \quad \frac{I \vdash i : I' \quad I' \vdash c : O}{I \vdash i; c : O} \\
\frac{\text{Writable}(x) \quad \text{fv}(e) \subseteq I}{I \vdash x \leftarrow e : I \cup \{x\}} \quad \frac{\text{Writable}(x) \quad \text{fv}(d) \subseteq I}{I \vdash x \stackrel{\$}{\leftarrow} d : I \cup \{x\}} \\
\frac{\text{fv}(e) \subseteq I \quad I \vdash c_i : O_i \quad i = 1, 2}{I \vdash \text{if } e \text{ then } c_1 \text{ else } c_2 : O_1 \cap O_2} \quad \frac{\text{fv}(e) \subseteq I \quad I \vdash c : I}{I \vdash \text{while } e \text{ do } c : I} \\
\frac{\text{fv}(\vec{e}) \subseteq I \quad \text{Writable}(x) \quad o \in \mathcal{O}}{I \vdash x \leftarrow o(\vec{e}) : I \cup \{x\}} \\
\frac{\text{fv}(\vec{e}) \subseteq I \quad \text{Writable}(x) \quad \mathcal{A} \notin \mathcal{O} \quad \vdash_{\text{wf}} \mathcal{A}}{I \vdash x \leftarrow \mathcal{A}(\vec{e}) : I \cup \{x\}} \\
\frac{\mathcal{G}_A \cup \mathcal{G}_{ro} \cup \mathcal{A}_E.\text{params} \vdash \mathcal{A}_E.\text{body} : O \quad \text{fv}(\mathcal{A}_E.\text{re}) \subseteq O}{\vdash_{\text{wf}} \mathcal{A}} \\
\text{Writable}(x) \stackrel{\text{def}}{=} \text{Local}(x) \vee x \in \mathcal{G}_A \quad \mathcal{A}_E \stackrel{\text{def}}{=} E(\mathcal{A}).
\end{array}$$

Figure 4. Static analysis for well-formedness of adversaries

Additional constraints may be imposed on adversaries. For example, exact security proofs usually impose an upper bound to the number of calls adversaries can make to a given oracle, whereas for some properties such as IND-CCA2 there are some restrictions on the parameters with which the oracles may be called. Likewise, some proofs impose extra conditions such as forbidding repeated or malformed queries. These kinds of properties can be formalized using lists that records the parameters of the calls to oracles, and verifying as postcondition that the calls were legitimate.

4. Relational Hoare Logic

Shoup [32] classifies proof steps into three categories: transitions based on indistinguishability—which typically involve applying a security hypothesis, e.g. the DDH assumption—; transitions based on failure events—which typically amount to bound the probability of bad, as in the Switching Lemma—; and bridging steps—which correspond to replacing or reorganizing code in a way that is not observable by adversaries. In some circumstances, a bridging transition from G_1 to G_2 may replace a program fragment P by another fragment P' observationally equivalent to P . In general, however, P and P' are only observationally equivalent in the context where the replacement is done. Such transitions are supported through a relational Hoare logic, that generalizes observational equivalence through preconditions and postconditions which we use to characterize the context where the replacement is valid. Besides, we use relational Hoare logic to establish (in)equalities between probabilities of two events, as shown by the lemmas $(=_{\square})$ and (\leq_{\square}) below, and to establish program invariants, e.g. in the proof of the Switching Lemma in Sec. 2.2.

4.1 Probabilistic Relational Hoare Logic (pRHL)

Our logic pRHL elaborates on and extends to probabilistic programs Benton’s Relational Hoare Logic [10]. Benton’s logic uses judgments of the form $\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi$, and relates the evaluation of a program G_1 to the evaluation of a program G_2 w.r.t. a precondition Ψ and a postcondition Φ , both defined as relations on deterministic states. Such a judgment states that for any initial memories m_1 and m_2 satisfying the precondition $m_1 \Psi m_2$, if the evaluations of G_1 in m_1 and G_2 in m_2 terminate with final memories m'_1 and m'_2 respectively, then $m'_1 \Phi m'_2$ holds. In a probabilistic setting, the evaluation of a program w.r.t. an initial memory yields a (sub-)distribution. In order to give a meaning to the above judgment, one therefore needs to lift relations over memories into relations over distributions.⁵ We follow early work on probabilistic bisimulations [23]. The lifting to distributions of a unary predicate P and of a binary relation Φ are respectively defined as

$$\begin{aligned} \text{range } P \ \mu &\stackrel{\text{def}}{=} \forall f. (\forall a. P a \Rightarrow f a = 0) \Rightarrow \mu f = 0 \\ \mu_1 \sim_{\Phi} \mu_2 &\stackrel{\text{def}}{=} \exists \mu. \pi_1(\mu) = \mu_1 \wedge \pi_2(\mu) = \mu_2 \wedge \text{range } \Phi \ \mu \end{aligned}$$

where the projections of μ are defined as

$$\pi_1(\mu) \stackrel{\text{def}}{=} \text{bind } \mu (\lambda x y. \text{unit } x) \quad \pi_2(\mu) \stackrel{\text{def}}{=} \text{bind } \mu (\lambda x y. \text{unit } y)$$

Definition 5 (pRHL judgments). *Programs G_1 and G_2 are equivalent w.r.t. precondition Ψ and postcondition Φ iff*

$$\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi \stackrel{\text{def}}{=} \forall m_1 m_2. m_1 \Psi m_2 \Rightarrow \llbracket G_1 \rrbracket m_1 \sim_{\Phi} \llbracket G_2 \rrbracket m_2$$

Our approach slightly departs from Benton’s: rather than defining the rules for pRHL and proving them sound w.r.t. the meaning of judgments, we place ourselves in a semantic setting and derive the rules as lemmas. This allows to easily extend the system by deriving extra rules, or even to resort to the semantic definition if the system turns out to be insufficient.

Figure 5 gathers some representative derived rules. To improve readability, in the figure and in the remainder of the paper we let $e\langle i \rangle$ denote $\lambda m_1 m_2. \llbracket e \rrbracket m_i = \text{true}$, where e is a boolean expression. As pRHL allows for arbitrary relations, we freely use higher-order logic; in particular, PER and SYM are predicates over relations that stand for *partial equivalence relation* and *symmetric relation* respectively. There are two points worth noting. First, most rules admit, in addition to their symmetrical version of Fig. 5, one-sided (left and right) variants, e.g. for assignments

$$\frac{m_1 \Phi' m_2 \stackrel{\text{def}}{=} (m_1 \{ \llbracket e_1 \rrbracket m_1 / x_1 \}) \Phi m_2}{\vdash E_1, x_1 \leftarrow e_1 \sim E_2, \text{nil} : \Phi' \Rightarrow \Phi}$$

Second, some rules of pRHL do not appear in RHL, or generalize existing rules. The rule [R-Case] allows to do a case analysis on the evaluation of an arbitrary relation in the initial memories. Together with simple rules in the spirit of

$$\frac{\vdash E_1, c_1 \sim E_2, c : \Psi \wedge e\langle 1 \rangle \Rightarrow \Phi}{\vdash E_1, \text{if } e \text{ then } c_1 \text{ else } c_2 \sim E_2, c : \Psi \wedge e\langle 1 \rangle \Rightarrow \Phi}$$

it subsumes [R-Cond] and allows to prove judgments that would otherwise not be derivable, such as the equivalence between (if e then c_1 else c_2) and (if $\neg e$ then c_2 else c_1). We also use [R-Case] to prove the correctness of dataflow analyses that exploit the information provided by entering branches.

In addition, we often use the rule [R-Inv] that generalizes the rule [R-Sym] to inverse of relations

$$\frac{\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi}{\vdash G_2 \sim G_1 : \Psi^{-1} \Rightarrow \Phi^{-1}} [\text{R-Inv}]$$

⁵ An alternative would be to develop a logic in which Ψ and Φ are relations over distributions of states. However, it is not clear whether such a logic would allow a similar level of proof automation. This is left for future work.

and we make an extensive use of the rule [R-Comp] that generalizes the rule [R-Tr] to composition of relations⁶

$$\frac{\vdash G_1 \sim G : \Psi' \Rightarrow \Phi' \quad \vdash G \sim G_2 : \Psi'' \Rightarrow \Phi''}{\vdash G_1 \sim G_2 : \Psi' \circ \Psi'' \Rightarrow \Phi' \circ \Phi''} [\text{R-Comp}]$$

The benefits of the rule [R-Comp], as opposed to [R-Tr], are illustrated by considering “independent” preconditions and postconditions of the form

$$\Psi \stackrel{\text{def}}{=} \lambda x y. \Psi_1 x \wedge \Psi_2 y \quad \Phi \stackrel{\text{def}}{=} \lambda x y. \Phi_1 x \wedge \Phi_2 y$$

In order to apply the rule [R-Tr] to G_1 and G_2 , we are essentially forced to have $\Psi_1 = \Psi_2$ and $\Phi_1 = \Phi_2$, and furthermore we must also choose the same pre and postcondition for the intermediate game G . This constraint makes the rule [R-Tr] impractical, we use instead the rule [R-Comp] to introduce intermediate games that do not satisfy the same specification as G_1 or G_2 .

The rule [R-Rand] is also (obviously) not present in RHL. Let $\mathbb{I}_x \stackrel{\text{def}}{=} (\lambda v. \text{if } x = v \text{ then } 1 \text{ else } 0)$, and define the support of a distribution, $\text{supp}(\llbracket d \rrbracket m)$, by the clause

$$v \in \text{supp}(\llbracket d \rrbracket m) \Leftrightarrow \llbracket d \rrbracket m \mathbb{I}_v \neq 0$$

Finally, let $\llbracket d_1 \rrbracket m_1 =_g \llbracket d_2 \rrbracket m_2$ iff there exists a set X and a bijection $g : X \rightarrow X$ such that $\text{supp}(\llbracket d_1 \rrbracket m_1) = \text{supp}(\llbracket d_2 \rrbracket m_2) = X$ and $\llbracket d_1 \rrbracket m_1 \mathbb{I}_a = \llbracket d_2 \rrbracket m_2 \mathbb{I}_{(g a)}$ for all a in X . To apply rule [R-Rand], it is necessary to exhibit a function f such that for all memories m_1 and m_2 , $\llbracket d_1 \rrbracket m_1 =_f m_1 m_2 \llbracket d_2 \rrbracket m_2$. Thus, if $d_1 = d_2 = [0..n]$ for some constant n , and we take f to be the identity function, the premise simplifies to the expected,

$$m_1 \Psi m_2 \stackrel{\text{def}}{=} \forall v \in [0..n]. (m_1 \{v/x_1\}) \Phi (m_2 \{v/x_2\})$$

Section 5.3 shows that the generality of the rule is required for applications such as optimistic sampling.

It is often fruitful to understand pRHL judgments in terms of the inability of the postcondition to separate between the two commands of the judgment. Define two functions f and g to be equivalent w.r.t. a predicate Φ iff

$$f =_{\Phi} g \stackrel{\text{def}}{=} \forall m_1 m_2. m_1 \Phi m_2 \Rightarrow f(m_1) = g(m_2)$$

The definition of pRHL judgments entails

$$\left. \begin{aligned} &\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi \\ &f =_{\Phi} g \\ &m_1 \Psi m_2 \end{aligned} \right\} \Rightarrow \llbracket G_1 \rrbracket m_1 f = \llbracket G_2 \rrbracket m_2 g \quad (=_{\square})$$

By instantiating f and g to $\mathbb{1}$, one can observe that observational equivalence enjoys some form of termination sensitivity

$$(\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi) \wedge m_1 \Psi m_2 \Rightarrow \llbracket G_1 \rrbracket m_1 \mathbb{1} = \llbracket G_2 \rrbracket m_2 \mathbb{1}$$

This interpretation of pRHL judgments is strongly connected to the relation between relational logics and information flow [3, 10]. We extensively use $(=_{\square})$, and its variant (\leq_{\square}) below, to fall back from the world of pRHL into the world of probabilities, in which security statements are expressed;

$$\left. \begin{aligned} &\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi \\ &f \leq_{\Phi} g \\ &m_1 \Psi m_2 \end{aligned} \right\} \Rightarrow \llbracket G_1 \rrbracket m_1 f \leq \llbracket G_2 \rrbracket m_2 g \quad (\leq_{\square})$$

where $f \leq_{\Phi} g \stackrel{\text{def}}{=} \forall m_1 m_2. m_1 \Phi m_2 \Rightarrow f(m_1) \leq g(m_2)$.

We conclude with an example that nicely illustrates some of the intricacies of pRHL. Let $c = b \stackrel{\text{def}}{=} \{0, 1\}$ and define $m_1 \Phi m_2 \stackrel{\text{def}}{=} m_1 b = m_2 b$. Then $\vdash c \sim c : \text{true} \Rightarrow \Phi$. Indeed, take μ such that $\mu \mathbb{1}_{(x,y)} = 1/2$ if $x = y$ and $\mu \mathbb{1}_{(x,y)} = 0$

⁶ The machine-checked rule requires that Φ is decidable, and uses **Set**-valued existential quantification \exists_{Set} in the composition for preconditions, i.e. $x (\Psi \circ \Psi') z \stackrel{\text{def}}{=} \exists_{\text{Set } y}. x \Psi y \wedge y \Psi' z$.

$$\begin{array}{c}
\vdash E_1, \text{nil} \sim E_2, \text{nil} : \Phi \Rightarrow \Phi \text{ [R-Skip]} \quad \frac{\vdash E_1, c_1 \sim E_2, c_2 : \Phi \Rightarrow \Phi' \quad \vdash E_1, c'_1 \sim E_2, c'_2 : \Phi' \Rightarrow \Phi''}{\vdash E_1, c_1; c'_1 \sim E_2, c_2; c'_2 : \Phi \Rightarrow \Phi''} \text{[R-Seq]} \\
\vdash E_1, x_1 \leftarrow e_1 \sim E_2, x_2 \leftarrow e_2 : (\lambda m_1 m_2. (m_1 \{ \llbracket e_1 \rrbracket m_1 / x_1 \} \Phi (m_2 \{ \llbracket e_2 \rrbracket m_2 / x_2 \}))) \Rightarrow \Phi \text{ [R-Ass]} \\
\frac{m_1 \Psi m_2 \stackrel{\text{def}}{=} \llbracket d_1 \rrbracket m_1 =_f m_1 m_2 \llbracket d_2 \rrbracket m_2 \wedge \forall v \in \text{supp}(\llbracket d_1 \rrbracket m_1). (m_1 \{ v / x_1 \}) \Phi (m_2 \{ f m_1 m_2 v / x_2 \})}{\vdash E_1, x_1 \stackrel{\text{def}}{=} d_1 \sim E_2, x_2 \stackrel{\text{def}}{=} d_2 : \Psi \Rightarrow \Phi} \text{[R-Rand]} \\
\frac{\forall m_1 m_2. m_1 \Psi m_2 \Rightarrow \llbracket e_1 \rrbracket m_1 = \llbracket e_2 \rrbracket m_2 \quad \vdash E_1, c_1 \sim E_2, c_2 : \Psi \wedge e_1 \langle 1 \rangle \Rightarrow \Phi \quad \vdash E_1, c'_1 \sim E_2, c'_2 : \Psi \wedge \neg e_1 \langle 1 \rangle \Rightarrow \Phi}{\vdash E_1, \text{if } e_1 \text{ then } c_1 \text{ else } c'_1 \sim E_2, \text{if } e_2 \text{ then } c_2 \text{ else } c'_2 : \Psi \Rightarrow \Phi} \text{[R-Cond]} \\
\frac{\forall m_1 m_2. m_1 \Phi m_2 \Rightarrow \llbracket e_1 \rrbracket m_1 = \llbracket e_2 \rrbracket m_2 \quad \vdash E_1, c_1 \sim E_2, c_2 : \Phi \wedge e_1 \langle 1 \rangle \Rightarrow \Phi}{\vdash E_1, \text{while } e_1 \text{ do } c_1 \sim E_2, \text{while } e_2 \text{ do } c_2 : \Phi \Rightarrow \Phi \wedge \neg e_1 \langle 1 \rangle} \text{[R-Whl]} \\
\frac{\vdash G_1 \sim G_2 : \Psi' \Rightarrow \Phi' \quad \forall m_1 m_2. m_1 \Psi m_2 \Rightarrow m_1 \Psi' m_2 \quad \forall m_1 m_2. m_1 \Phi' m_2 \Rightarrow m_1 \Phi m_2}{\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi} \text{[R-Sub]} \\
\frac{\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi \quad \text{SYM}(\Psi) \quad \text{SYM}(\Phi)}{\vdash G_2 \sim G_1 : \Psi \Rightarrow \Phi} \text{[R-Sym]} \quad \frac{\vdash G_1 \sim G : \Psi \Rightarrow \Phi \quad \vdash G \sim G_2 : \Psi \Rightarrow \Phi \quad \text{PER}(\Psi) \quad \text{PER}(\Phi)}{\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi} \text{[R-Tr]} \\
\frac{\vdash G_1 \sim G_2 : \Psi \wedge \Psi' \Rightarrow \Phi \quad \vdash G_1 \sim G_2 : \Psi \wedge \neg \Psi' \Rightarrow \Phi}{\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi} \text{[R-Case]}
\end{array}$$

Figure 5. Selection of derived rules of pRHL

otherwise. One can check that μ ensures $\llbracket c \rrbracket m \sim_{\Phi} \llbracket c \rrbracket m'$ for all m and m' . This example shows why the lifting of a binary relation involves an existential quantification, and why it is not possible to always instantiate μ as the product distribution in the definition of \sim_{Φ} (one cannot establish the above judgment using the product distribution). Perhaps more surprisingly, $\vdash c \sim c : \text{true} \Rightarrow \neg \Phi$ also holds. Indeed, take μ such that $\mu \mathbb{1}_{\langle x, y \rangle} = 1/2$ if $x \neq y$ and $\mu \mathbb{1}_{\langle x, y \rangle} = 0$ otherwise. One can check that μ ensures $\llbracket c \rrbracket m \sim_{\neg \Phi} \llbracket c \rrbracket m'$ for all m and m' . Thus, the “obvious” rule

$$\frac{\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi \quad \vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi'}{\vdash G_1 \sim G_2 : \Psi \Rightarrow \Phi \wedge \Phi'}$$

is unsound. While this example may seem unintuitive or even inconsistent if one reasons in terms of deterministic states, its intuitive significance in a probabilistic setting is that neither of the relations Φ and $\neg \Phi$ are enough to tell apart the two final distributions.

4.2 Observational equivalence

Observational equivalence is derived as an instance of relational Hoare judgments in which pre and postconditions are restricted to relations based on equality over a subset of variables. Given a set X of variables, we define $=_X$ as

$$m_1 =_X m_2 \stackrel{\text{def}}{=} \forall x \in X. m_1 x = m_2 x$$

Then, observational equivalence of G_1 and G_2 w.r.t. an input set I and an output set O is defined as

$$\vdash G_1 \simeq_O^I G_2 \stackrel{\text{def}}{=} \vdash G_1 \sim G_2 : =_I \Rightarrow =_O$$

All derived rules for pRHL can be specialized to the case of observational equivalence. For example, we have

$$\frac{\vdash e_1 \simeq^I e_2 \quad \vdash E_1, c_1 \simeq_O^I E_2, c_2 \quad \vdash E_1, c'_1 \simeq_O^I E_2, c'_2}{\vdash E_1, \text{if } e_1 \text{ then } c_1 \text{ else } c'_1 \simeq_O^I E_2, \text{if } e_2 \text{ then } c_2 \text{ else } c'_2}$$

where $\vdash e_1 \simeq^I e_2$ iff for every memories m_1 and m_2 , $m_1 =_I m_2$ implies $\llbracket e_1 \rrbracket m_1 = \llbracket e_2 \rrbracket m_2$.

To support automation, CertiCrypt implements a calculus of variable dependencies and provides two tactics, `eqobsIn` and `eqobsOut`, that given a command c and a set O (respectively I) of output (input) variables compute a set I (O) of input (output) variables such that $\vdash E_1, c \simeq_O^I E_2, c$. CertiCrypt also provides tactics for two variants of observational equivalence that are widely used in game-based proofs, namely

$$\begin{array}{l}
\vdash_{\varphi} G_1 \simeq_O^I G_2 \stackrel{\text{def}}{=} \vdash G_1 \sim G_2 : =_I \wedge \varphi \Rightarrow =_O \wedge \varphi \\
\vdash_{\Psi, \Phi} G_1 \simeq_O^I G_2 \stackrel{\text{def}}{=} \vdash G_1 \sim G_2 : =_I \wedge \Psi \Rightarrow =_O \wedge \Phi
\end{array}$$

These tactics use a (sound but incomplete) weakest precondition calculus for relational judgments.

5. Proof methods for bridging steps

CertiCrypt provides a powerful set of tactics and algebraic equivalences to automate bridging steps. Most tactics rely on an implementation of a certified optimizer for pWHILE, with the exception of *lazy sampling* which has an ad hoc implementation. Algebraic equivalences are provided as lemmas that follow from algebraic properties of the interpretation of language constructs.

5.1 Certified program transformations

CertiCrypt provides automated support for transformations that consist in applying compiler optimizations. More precisely, it supports transformations based on dependencies and dataflow analyses; we briefly discuss them below. Additionally, CertiCrypt provides support for inlining procedure calls and performing register allocation (not discussed here).

Transformations based on dependencies The functions `eqobsIn` and `eqobsOut` presented in Sec. 4, provide the foundations to support transformations such as dead code elimination, code motion, and common context elimination.

First, CertiCrypt features a function context that strips off two commands c and c' their maximal common context relative to sets I and O of input and output variables. The correctness of context

is expressed by the following rule

$$\frac{\text{context}(I, c_1, c_2, O) = (I', c'_1, c'_2, O') \quad \models E_1, c'_1 \simeq_{O'}^I E_2, c'_2}{\models E_1, c_1 \simeq_O^I E_2, c_2}$$

Using the same idea, CertiCrypt provides algorithms for removing only a common prefix (eqobs_hd) or suffix (eqobs_tl).

Second, CertiCrypt provides a tactic that given two commands repeatedly tries to hoist their common instructions to obtain a maximal common prefix⁷, which can then be eliminated using the previous rule. Its correctness is based on the rule

$$\frac{\models E, c_1 \simeq_{O_1}^{I_1} E, c_1 \quad \models E, c_2 \simeq_{O_2}^{I_2} E, c_2 \quad \text{Modify}(E, c_i, O_i) \quad O_1 \cap O_2 = \emptyset \quad I_1 \cap O_2 = \emptyset \quad I_2 \cap O_1 = \emptyset}{\models E, c_1; c_2 \sim E, c_2; c_1 := \Rightarrow =}$$

where $\text{Modify}(E, c, X)$ is a semantic predicate expressing that only variables in X are modified by the command c in the environment E . This is formally expressed by $\forall m. \text{range}(\lambda m'. m =_{V \setminus X} m') (\llbracket E, c \rrbracket m)$ which ensures that the final memories are equal to the initial ones except maybe on variables in X . The tactic swap is based on the rule above and on an algorithm that over-approximates the set of modified variables (by computing those that may be written, without taking into account self-assignments).

Third, CertiCrypt allows performing dead code elimination relative to a set O of output variables (deadcode). The algorithm behaves more like an aggressive slicing algorithm, i.e. it removes portions of code that do not affect variables in O , and performs at the same time branch prediction (replacing if true then c_1 else c_2 by c_1), branch coalescing (replacing if e then c else c by c), and self-assignment elimination. Its correctness relies on the rule

$$\frac{\text{Modify}(E_1, c, X) \quad \text{Lossless}(E_1, c) \quad \text{fv}(\varphi) \cap X = \emptyset}{\models E_1, c \sim E_2, \text{nil} : \varphi \Rightarrow \varphi}$$

Optimizations based on dataflow analyses CertiCrypt has builtin, generic, support for such optimizations: given an abstract domain D (a semi-lattice) for the analysis, transfer functions for assignment and branching instructions, and an operator transforming expressions in the language into their optimized versions (using the result of the analysis), CertiCrypt automatically constructs the certified optimization function $\text{optim} : \mathcal{C} \rightarrow D \rightarrow \mathcal{C} \times D$. When given a command c and an element $\delta \in D$, this function transforms c into its optimized version c' assuming the validity of δ . In addition, it returns an abstract postcondition $\delta' \in D$ which is valid after executing c (or c'). We use these abstract postconditions to state the correctness of the optimization, and to apply the optimization recursively.

The correctness of optim is proved using a mixture of the techniques of [10] and [11, 26]: we express the validity of the information contained in the analysis domain using a predicate $\text{Valid}(\delta, m)$ that states the agreement between the compile time abstract values in δ and the runtime memory m . Then, correctness is expressed in terms of a pRHL judgment (universally quantified over c and δ):

$$\text{let } (c', \delta') := \text{optim}(c, \delta) \text{ in } \models E, c \sim E, c' : \asymp_\delta \Rightarrow \asymp_{\delta'}$$

where $m_1 \asymp_\delta m_2 \stackrel{\text{def}}{=} m_1 = m_2 \wedge \text{Valid}(\delta, m_1)$. The following useful rule is derived using [R-Comp]

$$\frac{\forall m_1 m_2. m_1 \Psi m_2 \Rightarrow \text{Valid}(\delta, m_1) \quad \text{optim}(c_1, \delta) = (c'_1, \delta') \quad \models E_1, c'_1 \sim E_2, c_2 : \Psi \Rightarrow \Phi}{\models E_1, c_1 \sim E_2, c_2 : \Psi \Rightarrow \Phi} \text{[R-Opt]}$$

⁷One could also provide a complementary tactic that hoists instructions to obtain a maximal common suffix.

Our case studies extensively use instantiations of [R-Opt] to expression propagation (ep). In contrast, we found that common subexpression elimination is seldom used.

5.2 Lazy sampling

It is sometimes convenient to defer random choices in games until they are actually needed, or conversely, to make random choices as early as possible. The lazy sampling technique, allows to delay the random sampling of a value until the point in the program where it is first used. Conversely, eager sampling allows to choose a random value, which would be otherwise sampled later, at the beginning of a game. These techniques are presented in [9], where the authors discuss some of its subtleties. In this section we present a syntax-oriented criterion for the correctness of lazy or eager sampling that can be applied provided the sampling is adequately guarded. Here by *context* we mean a program context with multiple holes that may appear either in the main program or any of the procedures in the environment.

Lemma 1 (Lazy/eager sampling). *Let $C[\cdot]$ be a context, c_1 and c_2 commands, e a boolean expression, d a distribution expression, and z a variable, such that $C[\cdot]$ does not modify $\text{fv}(e) \cup \text{fv}(d)$ and does not use z . Assume*

1. $\models z \stackrel{\$}{=} d; c_1; \text{if } e \text{ then } z \stackrel{\$}{=} d \sim z \stackrel{\$}{=} d; c_1 : (= \wedge e\langle 1 \rangle) \Rightarrow =$
2. $\models c_2 \sim c_2 : (= \wedge \neg e\langle 1 \rangle) \Rightarrow (= \wedge \neg e\langle 1 \rangle)$

Let $c = \text{if } e \text{ then } z \stackrel{\$}{=} d; c_1 \text{ else } c_2$ and $c' = \text{if } e \text{ then } c_1 \text{ else } c_2$. Then, $\models C[c]; \text{if } e \text{ then } z \stackrel{\$}{=} d \sim z \stackrel{\$}{=} d; C[c'] : (= \wedge e\langle 1 \rangle) \Rightarrow =$

Intuitively, in the above lemma e indicates whether z has not been used in the game since it was last sampled. If it has not been used, then it is perfectly fine to resample it. The first two hypotheses ensure that e has exactly this meaning, c_1 must set it to false if it has used the value sampled in z , and c_2 must not reset e if it is false. The first hypothesis is the one that allows to swap c_1 with $z \stackrel{\$}{=} d$, provided the value of z is not used in c_1 . Note that, for clarity, we have omitted environments in the above lemma, and so the second hypothesis is not as trivial as it may seem because both programs may have different environments.

Suppose we want to eagerly sample the answer that a random oracle

$$\mathcal{O}_1(x) \stackrel{\text{def}}{=} \text{if } x \notin \text{dom}(L) \text{ then } y \stackrel{\$}{=} d; L \leftarrow (x, y) :: L; \text{return } L[x]$$

gives to a particular query x' , i.e. we want to transform \mathcal{O}_1 into

$$\mathcal{O}_e(x) \stackrel{\text{def}}{=} \text{if } x \notin \text{dom}(L) \text{ then} \\ \text{if } x = x' \text{ then } y \leftarrow y' \text{ else } y \stackrel{\$}{=} d; \\ L \leftarrow (x, y) :: L \\ \text{return } L[x]$$

Define $\mathcal{O}'_1(x) \stackrel{\text{def}}{=} \text{if } x' \notin \text{dom}(L) \text{ then } y' \stackrel{\$}{=} d; \mathcal{O}_e(x) \text{ else } \mathcal{O}_1(x)$, $\mathcal{O}'_e(x) \stackrel{\text{def}}{=} \text{if } x' \notin \text{dom}(L) \text{ then } \mathcal{O}_e(x) \text{ else } \mathcal{O}_1(x)$, both returning $L[x]$. Oracles \mathcal{O}_1 and \mathcal{O}_e result semantically equivalent to \mathcal{O}'_1 and \mathcal{O}'_e , respectively. Lemma 1 can be applied taking $e = x' \notin \text{dom}(L)$, $z = y'$, and $c_1 = \mathcal{O}_e(x)$, $c_2 = \mathcal{O}_1(x)$ to safely replace oracle \mathcal{O}_1 by oracle \mathcal{O}_e in the environment of a program, sampling y' at the beginning. Whenever a bound for the number of queries to a random oracle is known in advance, the above trick can be iteratively applied to completely remove randomness from the oracle code, as it is done in the proof of the Switching Lemma in Sec. 2.2.

5.3 Algebraic equivalences

Bridging steps frequently use algebraic properties of language operators. The proof of semantic security of ElGamal uses the fact that in a cyclic multiplicative group, multiplication by a uniformly

sampled element acts as a one-time pad:

$$\models x \stackrel{\$}{\leftarrow} \mathbb{Z}_q; \alpha \leftarrow g^x \times \beta \simeq_{\{\alpha\}} y \stackrel{\$}{\leftarrow} \mathbb{Z}_q; \alpha \leftarrow g^y$$

In the proof of security of OAEP we use optimistic sampling:

$$\models x \stackrel{\$}{\leftarrow} \{0, 1\}^k; y \leftarrow x \oplus z \simeq_{\{x, y, z\}} y \stackrel{\$}{\leftarrow} \{0, 1\}^k; x \leftarrow y \oplus z$$

and incremental sampling modulo a permutation f :

$$\models x \stackrel{\$}{\leftarrow} \{0, 1\}^{k-\rho}; y \stackrel{\$}{\leftarrow} \{0, 1\}^\rho; z \leftarrow f(x||y) \simeq_{\{z\}} z \stackrel{\$}{\leftarrow} \{0, 1\}^k$$

We show the usefulness of [R-Rand] by sketching the proof of optimistic sampling, as promised in Section 4. For readability, we let $e_{|i}$ denote $\llbracket e \rrbracket_{m_i}$. Define

$$\Psi \stackrel{\text{def}}{=} z_{|1} = z_{|2} \quad \Phi \stackrel{\text{def}}{=} x_{|1} = x_{|2} \wedge y_{|1} = y_{|2} \wedge z_{|1} = z_{|2}$$

Let $\Phi' \stackrel{\text{def}}{=} \Phi\{x_{|1} \oplus z_{|1}/y_{|1}, y_{|2} \oplus z_{|2}/x_{|2}\}$. Then, by [R-Ass] we have

$$\models y \leftarrow x \oplus z \sim x \leftarrow y \oplus z : \Phi' \Rightarrow \Phi$$

Now take $f \ m_1 \ m_2 \ v \stackrel{\text{def}}{=} v \oplus z_{|2}$, and apply [R-Rand] and [R-Sub] to obtain $\models x \stackrel{\$}{\leftarrow} \{0, 1\}^k \sim y \stackrel{\$}{\leftarrow} \{0, 1\}^k : \Psi \Rightarrow \Phi'$. Note that the precondition we obtain after applying [R-Rand] is equivalent to Ψ because f is a bijection on $\{0, 1\}^k$, and because $\forall v. \Phi'\{v/x_{|1}, v \oplus z_{|2}/y_{|2}\}$ is equivalent to $z_{|1} = z_{|2}$ by algebraic properties of the \oplus operator. We conclude by applying [R-Seq].

6. Proof methods for failure events

A technique used very often to relate two games is based on what cryptographers call *failure events*. This technique relies on a *fundamental lemma* that allows to bound the difference in the probability of a given event in two games: one identifies a failure event and argues that both games behave identically until this event occurs. One can then bound the difference in probability of another event by the probability of occurrence of the failure event in either game.

Lemma 2 (Fundamental lemma). *Let G_1 and G_2 be two games, A an event defined on G_1 , B an event defined on G_2 and F an event defined in both games. If*

$$\Pr_{G_1}[A \wedge \neg F] = \Pr_{G_2}[B \wedge \neg F], \text{ and} \\ \Pr_{G_1}[F] \leq \Pr_{G_2}[F]$$

*then $|\Pr_{G_1}[A] - \Pr_{G_2}[B]| \leq \Pr_{G_2}[F]$.*⁸

In code-based proofs, the failure condition is generally indicated by setting a global flag variable (usually called *bad*) to true. This specialization allows to define a syntactic criterion for deciding whether two games behave equivalently up to the raise of the failure condition: we say that two games G_1 and G_2 are equal up to *bad* and note it $\text{uptobad}(G_1, G_2)$ whenever they are syntactically equal up to every point where the *bad* flag is set to true and they do not reset the *bad* flag to false afterward. For instance, games $G_{\text{PRP}}^{\text{bad}}$ and $G_{\text{PRF}}^{\text{bad}}$ in Fig. 2 satisfy this condition. We have used this syntactic criterion to implement a specialization of the fundamental lemma for game-based proofs.

Lemma 3 (Syntactic criterion for fundamental lemma).

$$\forall G_1 \ G_2 \ A. \ \text{uptobad}(G_1, G_2) \Rightarrow \\ \Pr_{G_1}[A \wedge \neg \text{bad}] = \Pr_{G_2}[A \wedge \neg \text{bad}]$$

The first hypothesis in Lemma 2 may be proved automatically by using this syntactic criterion. To prove the second hypothesis it suffices to show that game G_2 is absolutely terminating, for

⁸The second hypothesis is usually omitted in the literature under the assumption that both games are absolutely terminating. In that case, either G_1 or G_2 will do on the right-hand side.

Game EU_{FDH} : $L \leftarrow []; S \leftarrow [];$ $(m, x) \leftarrow \mathcal{A}();$ $d \leftarrow H(m)$	Oracle $H(m)$: if $m \notin \text{dom}(L)$ then $r \stackrel{\$}{\leftarrow} \{0, 1\}^k;$ $L \leftarrow (m, r) :: L$ return $L[x]$	Oracle $\text{Sign}(m)$: $S \leftarrow m :: S$ $r \leftarrow H(m);$ return $f^{-1}(r)$
---	---	--

Figure 6. Initial game in the proof of FDH unforgeability

which we already have implemented a semi-decision procedure (see Sec. 3.3).

7. Case studies

The purpose of this section is to announce the successful completion of two experiments that validate the design and usability of CertiCrypt. A detailed presentation of these works will be given elsewhere.

7.1 Existential unforgeability of FDH

The Full Domain Hash (FDH) scheme is a hash-and-sign signature scheme based on the RSA family of trapdoor permutations, and in which the message is hashed onto the full domain of the RSA function. However, the same construction—and the reduction that proves its security—remains valid for any family of trapdoor permutations. We have proved that FDH is existentially unforgeable under adaptive chosen-message attacks [8] in the random oracle model. The proof is about 2,700 lines long.

In the following we will consider a generic family of trapdoor permutations f (and their inverses f^{-1}) indexed by the security parameter, and an ideal hash function $H : \{0, 1\}^* \rightarrow \{0, 1\}^k$, which we model as a random oracle. The initial game of the proof is shown in Fig. 6.

Definition 6 (Trapdoor permutation security). *We say that a trapdoor permutation is ϵ' -secure if an effective inverter, when given a challenge y uniformly drawn from $\{0, 1\}^k$ succeeds in finding $f^{-1}(y)$ with probability at most $\epsilon'(k)$, i.e. if for any well-formed and PPT adversary \mathcal{B} we have $\Pr_{G_f}[x = f^{-1}(y)] \leq \epsilon'(k)$, where*

Game G_f : $y \stackrel{\$}{\leftarrow} \{0, 1\}^k; x \leftarrow \mathcal{B}(y)$.

Theorem 1 (FDH exact security). *Assume the underlying trapdoor permutation is ϵ' -secure. Then, any effective forger that makes at most q_{hash} and q_{sign} queries to the hash and signing oracles respectively, succeeds in forging a signature for a new message—different from the ones asked to the signing oracle—with probability at most*

$$\epsilon(k) = (q_{\text{hash}}(k) + q_{\text{sign}}(k) + 1) \epsilon'(k)$$

i.e. for any well-formed and PPT adversaries $\mathcal{A}, \mathcal{A}'$ we have $\Pr_{\text{EU}_{\text{FDH}}}[d = f(x)] \leq (q_{\text{hash}}(k) + q_{\text{sign}}(k) + 1) \epsilon'(k)$.

Our next objective is to formalize Coron's proof [16], which improves the conventional bound of FDH given above by eliminating the dependency on q_{hash} .

7.2 Semantic security of OAEP

OAEP is a padding scheme that is widely used in conjunction with an encryption scheme such as RSA in order to add randomness into plaintexts to achieve a high level of security. The history of OAEP perfectly illustrates the difficulty in achieving a correct proof. Indeed, it was initially believed that OAEP was IND-CCA2 secure [7], but it was later discovered it was only IND-CCA1 secure [31], a weaker security notion (where the adversary does not have access to the decryption oracle after receiving the challenge). It is possible to recover IND-CCA2 security by choosing a suitable encryption scheme, as it is the case for RSA-OAEP [19]. Here

we focus on the game-based proof of [9] which shows IND-CPA security of OAEP in the random oracle model.

The definition of OAEP is parametrized by a trapdoor one-way permutation $f : \{0, 1\}^k \rightarrow \{0, 1\}^k$, and two hash functions $G : \{0, 1\}^\rho \rightarrow \{0, 1\}^{k-\rho}$ and $H : \{0, 1\}^{k-\rho} \rightarrow \{0, 1\}^\rho$. OAEP adds randomness into the plaintext m and uses the functions G and H to mask it before applying f to the result, as formalized by the straightforward code for encryption:

```

R  $\xleftarrow{\$}$   $\{0, 1\}^\rho$ ;
S  $\leftarrow G(R) \oplus m$ ;
T  $\leftarrow H(S) \oplus R$ ;
y  $\leftarrow f(S||T)$ ;
return y

```

We have proved that OAEP is IND-CPA secure. The overall proof is about 3,000 lines of Coq.

Theorem 2 (OAEP semantic security). *For well-formed, PPT adversaries \mathcal{A} and \mathcal{A}' making together at most q_G queries to G ,*

$$|\Pr_{\text{IND-CPA}_{\text{OAEP}}} [b = b'] - \frac{1}{2}| \leq \Pr_{G_f} [x = f^{-1}(y)] + \frac{q_G}{2^\rho}$$

where $\Pr_{G_f} [x = f^{-1}(y)]$ is the probability of an adversary inverting f on a random element of its codomain, as in Definition 6.

Our sequence of games differs from [9]: in their initial transitions, Bellare and Rogaway use the fundamental lemma, whereas we use lazy sampling. As a result, our bound for OAEP is tighter than the bound published in [9], which also involves the number q_H of calls to H :

$$|\Pr_{\text{IND-CPA}_{\text{OAEP}}} [b = b'] - \frac{1}{2}| \leq \Pr_{G_f} [x = f^{-1}(y)] + \frac{2q_G}{2^\rho} + \frac{q_H}{2^{k-\rho}}$$

We consider that our proof of OAEP is highly emblematic, because of its complexity and its history. In retrospect, the bound we prove, which is independent of q_H , shows that formalizing proofs sometimes leads to improvement over previous results (to the best of our knowledge). However, cryptographers are really interested in a proof of IND-CCA2, and thus our next objective is to machine-check the results of [19].

8. Related work

Cryptographic protocol verification is an established area of formal methods, and a wealth of automated and deductive methods have been developed to the purpose of verifying that protocols provide the expected level of security [27]. Traditionally, protocols have been verified in a symbolic model, for which effective decision procedures exist under suitable hypotheses. Although the symbolic model assumes perfect cryptography, soundness results such as [1] relate the symbolic model with the computational model, provided the cryptographic primitives are secure. It is possible to combine symbolic methods and soundness proofs to achieve guarantees in the computational model, as done e.g. in [5, 33]. One drawback of this approach is that the security proof relies on intricate soundness proofs. Besides, it is not clear whether computational soundness results will always exist to allow factoring verification through symbolic methods. Consequently, some authors attempt to provide guarantees directly at the computational level [12, 25, 30].

In contrast, the formal verification of cryptographic functionalities is an emerging trend. An early work of Barthe, Cederquist and Tarento [6] proves the security of ElGamal in Coq, but the proof relies on the generic model, a very specialized and idealized model that elides many of the issues that are relevant for cryptography. Corin and den Hartog [15] also prove ElGamal semantic security, using a probabilistic (non-relational) Hoare logic. However, their formalism is not sufficiently powerful to express precisely the se-

curity goals: notions such as negligible advantage or effective adversary are not modeled.

Blanchet and Pointcheval [13] were among the first to use verification tools to carry out game-based proofs of cryptographic schemes. They used CryptoVerif for proving exact security of the FDH signature scheme, for the conventional bound given in Section 7.1. More recently, Courant et al [17] have developed a form of strongest postcondition calculus that can establish automatically asymptotic security (IND-CPA and IND-CCA2) of schemes that use one-way functions and random oracles. They show soundness and provide a prototype implementation that covers many examples of the literature, including OAEP+. We believe the two approaches are complementary to ours: compiling CryptoVerif sequences of games and embedding the type system of [17] in CertiCrypt, are interesting research directions.

In parallel, several authors have initiated formalizations of game-based proofs in proof assistants, and shown the security of basic examples. Nowak [28] gives a game-based proof of ElGamal semantic security in Coq. Nowak uses a shallow embedding to model games; as a result, its framework ignores complexity issues, and it is difficult to provide support for proof automation: because there is no special syntax for writing games, mechanizing syntactic transformations becomes very difficult. Affeldt et al [2] formalize a game-based proof of the switching lemma in Coq. However, their formalization is tailored towards the particular example they consider, which substantially simplifies their task and hinders generality. They deal with a weak (non-adaptive) adversary model and ignore complexity. All in all, both works appear like preliminary experiments that are not likely to scale.

Leaving the realm of cryptography, CertiCrypt relies on diverse mathematical concepts and theories that have been modeled for their own sake. It is not possible to report on these developments here, so we limit ourselves to singling out Paulin's formalization of probabilities, which we use extensively in our work, and the work of Hurd et al. [22], who developed a mechanized theory in the HOL theorem prover for reasoning about pGCL programs, a probabilistic extension of Dijkstra's guarded command language.

9. Conclusion

Summary and perspectives CertiCrypt is a fully formalized framework that supports machine-checked game-based proofs; we have validated its design through formalizing standard cryptographic proofs. Our work shows that machine-checked proofs of cryptographic schemes is not only plausible but indeed feasible. However, constructing machine-checked proofs requires a high-level of expertise in formal proofs and remains time consuming despite the high level of automation provided by CertiCrypt. Thus, CertiCrypt only provides a first step towards the completion of Halevi's programme, in spite of the amount of work invested so far (the project was initiated in June 2006).

A medium-term objective would be to develop a minimalist interface that eases the writing of games and provides a fixed set of mechanisms (tactics, proof-by-pointing) to prove some basic transitions, leaving the side conditions as hypotheses in the proof. We believe that such an interface would help cryptographers ensure that there are no obvious flaws in their definitions and proofs. In fact, it is our experience that the type system and the automated tactics provide valuable information to debug games and proofs.

Future work Numerous research directions remain to be explored. Our first objective is to strengthen our results for OAEP and FDH. Another objective is to formalize the correctness of the symbolic model and of automated methods for proving computational soundness, i.e. formalize results from [1, 17, 25]. We intend to focus first on [17], as its formalization may be useful to increase

automation in CertiCrypt. It would also be beneficial to formalize cryptographic meta-results such as the equivalence between IND-CPA and IND-CCA2 under plaintext awareness. It would also be worthwhile to explore applications of CertiCrypt outside cryptography, in particular to randomized algorithms and complexity.

References

- [1] M. Abadi and P. Rogaway. Reconciling two views of cryptography (the computational soundness of formal encryption). *Journal of Cryptology*, 15(2):103–127, 2002.
- [2] R. Affeldt, M. Tanaka, and N. Marti. Formal proof of provable security by game-playing in a proof assistant. In *Proceedings of International Conference on Provable Security, ProvSec 2007*, volume 4784 of *Lecture Notes in Computer Science*, pages 151–168. Springer-Verlag, 2007.
- [3] T. Amtoft, S. Bandhakavi, and A. Banerjee. A logic for information flow in object-oriented programs. In *Proceedings of the 33rd ACM Symposium on Principles of Programming Languages*, pages 91–102. ACM Press, 2006.
- [4] P. Audebaud and C. Paulin-Mohring. Proofs of randomized algorithms in Coq. *Science of Computer Programming*, 2008.
- [5] M. Backes and P. Laud. Computationally sound secrecy proofs by mechanized flow analysis. In *Proceedings of the 13th ACM Conference on Computer and Communications Security*, pages 370–379. ACM Press, 2006.
- [6] G. Barthe, J. Cederquist, and S. Tarento. A machine-checked formalization of the generic model and the random oracle model. In *2nd International Joint Conference on Automated Reasoning*, pages 385–399. Springer-Verlag, 2004.
- [7] M. Bellare and P. Rogaway. Optimal asymmetric encryption – How to encrypt with RSA. In *Advances in Cryptology – EUROCRYPT’94*, volume 950 of *LNCS*, pages 92–111. Springer-Verlag, 1995.
- [8] M. Bellare and P. Rogaway. The exact security of digital signatures – How to sign with RSA and Rabin. In *Advances in Cryptology – EUROCRYPT’96*, volume 1070 of *Lecture Notes in Computer Science*, pages 399–416. Springer-Verlag, 1996.
- [9] M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *Advances in Cryptology – EUROCRYPT’06*, volume 4004 of *LNCS*, pages 409–426, 2006.
- [10] N. Benton. Simple relational correctness proofs for static analyses and program transformations. In *Proceedings of the 31th ACM Symposium on Principles of Programming Languages*, pages 14–25. ACM Press, 2004.
- [11] Y. Bertot, B. Grégoire, and X. Leroy. A structured approach to proving compiler optimizations based on dataflow analysis. In *International Workshop on Types for Proofs and Programs*, volume 3839 of *Lecture Notes in Computer Science*, pages 66–81. Springer-Verlag, 2006.
- [12] B. Blanchet. A computationally sound mechanized prover for security protocols. In *IEEE Symposium on Security and Privacy*, pages 140–154, 2006.
- [13] B. Blanchet and D. Pointcheval. Automated security proofs with sequences of games. In *Advances in Cryptology – CRYPTO’06*, volume 4117 of *Lecture Notes in Computer Science*, pages 537–554. Springer-Verlag, 2006.
- [14] R. Canetti, O. Goldreich, and S. Halevi. The random oracle methodology, revisited. *Cryptology ePrint Archive*, Report 1998/011, 1998.
- [15] R. Corin and J. den Hartog. A probabilistic Hoare-style logic for game-based cryptographic proofs. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming*, volume 4052 of *LNCS*, pages 252–263, 2006.
- [16] J.-S. Coron. On the exact security of Full Domain Hash. In *Advances in Cryptology*, volume 1880 of *Lecture Notes in Computer Science*, pages 229–235. Springer-Verlag, 2000.
- [17] J. Courant, M. Daubignard, C. Ene, P. Lafourcade, and Y. Lakhnech. Towards automated proofs for asymmetric encryption in the random oracle model. In *Computer and Communications Security*. ACM Press, 2008.
- [18] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. In *Advances in Cryptology – CRYPTO’84*, volume 196 of *LNCS*, pages 10–18. Springer-Verlag, 1985.
- [19] E. Fujisaki, T. Okamoto, D. Pointcheval, and J. Stern. RSA-OAEP is secure under the RSA assumption. *Journal of Cryptology*, 17(2):81–104, 2004.
- [20] S. Goldwasser and S. Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2):270–299, 1984.
- [21] S. Halevi. A plausible approach to computer-aided cryptographic proofs. *Cryptology ePrint Archive*, Report 2005/181, 2005.
- [22] J. Hurd, A. McIver, and C. Morgan. Probabilistic guarded commands mechanized in HOL. *Theor. Comput. Sci.*, 346(1):96–112, 2005.
- [23] B. Jonsson and K. G. Larsen. Specification and refinement of probabilistic processes. In *Proceedings of the 6th Annual IEEE Symposium on Logic in Computer Science*, pages 266–279. IEEE Computer Society Press, 1991.
- [24] D. Kozen. Semantics of probabilistic programs. *J. Comput. Syst. Sci.*, 22:328–350, 1981.
- [25] P. Laud. Semantics and program analysis of computationally secure information flow. In *European Symposium on Programming*, volume 2028 of *Lecture Notes in Computer Science*, pages 77–91. Springer-Verlag, 2001.
- [26] X. Leroy. Formal certification of a compiler back-end, or: programming a compiler with a proof assistant. In *Proceedings of the 33rd ACM Symposium Principles of Programming Languages*, pages 42–54. ACM Press, 2006.
- [27] C. Meadows. Formal methods for cryptographic protocol analysis: Emerging issues and trends. *IEEE Journal on Selected Areas in Communications*, 21(1):44–54, 2003.
- [28] D. Nowak. A framework for game-based security proofs. In *Information and Communications Security*, volume 4861, pages 319–333. Springer-Verlag, 2007.
- [29] N. Ramsey and A. Pfeffer. Stochastic lambda calculus and monads of probability distributions. In *Proceedings of the 29th ACM Symposium on Principles of Programming Languages*, pages 154–165. ACM Press, 2002.
- [30] A. Roy, A. Datta, A. Derek, and J. C. Mitchell. Inductive proofs of computational secrecy. In *European Symposium On Research In Computer Security*, volume 4734 of *Lecture Notes in Computer Science*, pages 219–234. Springer-Verlag, 2007.
- [31] V. Shoup. OAEP reconsidered. In *Advances in Cryptology – CRYPTO’01*, volume 2139 of *Lecture Notes in Computer Science*, pages 239–259. Springer-Verlag, 2001.
- [32] V. Shoup. Sequences of games: a tool for taming complexity in security proofs. *Cryptology ePrint Archive*, Report 2004/332, 2004.
- [33] C. Sprenger and D. Basin. Cryptographically-sound protocol-model abstractions. In *Proceedings of CSF’08*, pages 115–129. IEEE Computer Society, 2008.
- [34] J. Stern. Why provable security matters? In *Advances in Cryptology – EUROCRYPT’03*, volume 2656 of *Lecture Notes in Computer Science*. Springer-Verlag, 2003.
- [35] The Coq development team. The Coq Proof Assistant Reference Manual v8.1, 2006. Available at <http://coq.inria.fr>.