Reflection, Rewinding, and Coin-Toss in EasyCrypt

Denis Firsov
deins@cs.ioc.ee
Tallinn University of Technology
Estonia
Guardtime
Tallinn, Estonia

Dominique Unruh
unruh@ut.ee
University of Tartu
Estonia

Abstract

In this paper we derive a suite of lemmas which allows users to internally reflect EasyCrypt programs into distributions which correspond to their denotational semantics (probabilistic reflection). Based on this we develop techniques for reasoning about rewinding of adversaries in EasyCrypt. (A widely used technique in cryptology.) We use our reflection and rewindability results to prove the security of a coin-toss protocol.

1 Introduction

Handwritten cryptographic security proofs are inherently error-prone. Humans will make mistakes both when writing and when checking the proofs. To ensure high confidence in cryptographic systems, we use frameworks for computer-aided verification of cryptographic proofs. One widely used such framework is the EasyCrypt tool [5]. In EasyCrypt, a cryptographic proof is represented by a sequence of "games" (simple probabilistic programs), and the relationship between programs are analyzed in a probabilistic relational Hoare logic (pRHL). EasyCrypt has been successfully used to verify a variety of cryptographic schemes: electronic voting [13], digital signatures [14], differential privacy [3], security of IPsec [16], and many others.

However, there is one cryptographic proof technique that seems very difficult to implement in EasyCrypt, namely rewinding. Rewinding is ubiquitous in more advanced proofs, especially when involving zero-knowledge proofs, multi-party computation, but also in relatively simple cases such as building a coin-toss from a commitment. (The latter case we will explore in Sec. 5.) In a nutshell, rewinding refers to the proof technique in which we take a given (usually unknown) program A (an adversary), and convert it into an adversary B that performs the following or similar steps:

1. Remember the initial state of A.
2. Run A.
3. Restore the original initial state of A.
4. Run A again.
5. Combine the results from the runs and/or repeat this until it yields a desired outcome.

While the above steps seem simple, we run into numerous challenges when trying to implement rewinding in Easy-Crypt, both due to restrictions in the type system, and due to the necessity for reasoning about probability distributions of program outputs in a way that is not directly supported by EasyCrypt's tactics.

To the best of our knowledge, rewinding has not been implemented in EasyCrypt, nor in other frameworks for reasoning about cryptographic proofs, see Sec. 1.4.

Our contribution. In this work, we design a set of tools to address rewindability in the EasyCrypt framework, and for reasoning about the probabilistic semantics of programs inside EasyCrypt (we call this probabilistic reflection). We validate our results by developing a formal proof of a coin-toss protocol based on rewinding. The EasyCrypt code of our framework is found in [15].

1.1 Challenges

To understand the motivation behind our project, let us look at the example of a pen-and-paper derivation using rewinding. When analyzing a coin-toss protocol based on a commitment, we are faced with the following situation: We have an adversary A that can open the commitment to contain given boolean value b (with some non-zero probability). We want to show that this means that A could also, in the same run, produce openings to both boolean values (i.e., *false* and *true*). This is done by defining a different adversary B that runs $A.commit$ (to produce the commitment), stores the state of A,

runs *A.open(false)* (to produce the first opening), restores the state of *A*, and runs *A.open(true)* (to produce the second opening). Then we show that the probability that *B* produces two valid openings is lower-bounded in terms of the probability that *A* is successful in producing one valid opening. As we will see in more detail in Sec. 5.1, the core theorem for showing this is the following.

Theorem 1.1. *Let A be a probabilistic program and let* **m** *denote a memory configuration which represents an initial state of A. We write* $\Pr[\, r \leftarrow X.p() \,@\, \mathbf{m} : M \,]$ *to denote the probability of a predicate M being satisfied by the result of running the procedure p() of a module X on the initial memory* **m**. *In this case, the following inequality holds:*

$$
\Pr \left[\begin{array}{l} A.init();\ s \leftarrow A.getState(); \\ r_1 \leftarrow A.main();\ A.setState(s); \\ r_2 \leftarrow A.main()\ @\ \mathbf{m} : r_1 \land r_2 \end{array} \right]
$$
$$
\geq\ \Pr[\, A.init();\ r \leftarrow A.main()\ @\ \mathbf{m} : r\,]^2.
$$

Proof. Step (1) applies "the averaging technique" by representing *A.init()* as a family of distributions $D_A^{\mathbf{m}}$. ($D_A^{\mathbf{m}}$ denotes the distribution in the family corresponding to memory **m**.) We write $\mu_1(D_A^{\mathbf{m}}, \mathbf{n})$ for the probability that $D_A^{\mathbf{m}}$ assigns to memory configuraion **n**. Then $\mu_1(D_A^{\mathbf{m}}, \mathbf{n})$ is the probability of *A.init*() terminating in the memory state **n** given that it starts in the initial state **m**. The rest of the computations are run starting from memory configuration **n**.

$$
\Pr \left[\begin{array}{l} A.init();\ s \leftarrow A.getState(); \\ r_1 \leftarrow A.main();\ A.setState(s); \\ r_2 \leftarrow A.main()\ @\ \mathbf{m} : r_1 \land r_2 \end{array} \right]
$$
$$
\overset{(1)}{=} \sum_{\mathbf{n}} \mu_1(D_A^{\mathbf{m}}, \mathbf{n}) \cdot \Pr \left[\begin{array}{l} s \leftarrow A.getState(); \\ r_1 \leftarrow A.main(); A.setState(s); \\ r_2 \leftarrow A.main()\ @\ \mathbf{n} : r_1 \land r_2 \end{array} \right]
$$
$$
\overset{(2)}{=} \sum_{\mathbf{n}} \mu_1(D_A^{\mathbf{m}}, \mathbf{n}) \cdot \Pr[\, r \leftarrow A.main()\ @\ \mathbf{n} : r\,]^2
$$
$$
\overset{(3)}{\geq} \left(\sum_{\mathbf{n}} \mu_1(D_A^{\mathbf{m}}, \mathbf{n}) \cdot \Pr[\, r \leftarrow A.main()\ @\ \mathbf{n} : r\,] \right)^2
$$
$$
\overset{(4)}{=} \Pr[\, A.init();\ r \leftarrow A.main()\ @\ \mathbf{m} : r\,]^2.
$$

Step (2) makes use of the fact that the probability of a success (i.e., *A.main* returning *true*) in both of two independent runs equals to the square of a probability of a success in a single run. Step (3) is an application of Jensen's inequality. The final step undoes the averaging. □

There are number of challenges in performing this proof formally in EasyCrypt:

1. Our proof turns the program *A.init()* into a parameterized distribution of final memories (memories after *A.init()* terminates). While EasyCrypt has a notion $\Pr[\, A.init()\ @\ \mathbf{m} : M\,]$ it does not formally recognize

that this defines a distribution, not withstanding the suggestive syntax.
2. Also, our proof makes use of results about probability distributions (e.g., Jensen's inequality) which can be easily stated and proved in terms of probabilistic distributions while they would be much harder (or even impossible) to express and prove directly using program logics (e.g., probabilistic Hoare logic).
3. Another challenge is that EasyCrypt does not have a type "memory"; we cannot define a distribution over memories because we cannot even assign it a type. In EasyCrypt, memories are recognized purely syntactically by prepending the variable names with &.
4. Finally, it is not immediate how one can generically specify the interface of programs (modules) which can return their own state. In other words, what should be the type of the return value of the function *A.getState()*?

1.2 Probabilistic Reflection

As explained above, one of the challenges is that we cannot access the distribution which corresponds to the semantics of the program. To enable this, we introduce a suite of lemmas which allows us to access the distribution corresponding to a program.

This turned out to be a powerful tool for rewinding proofs, but we also believe that it can be useful when one needs to derive facts for programs based on their denotational semantics. For example, in a situation when a particular tactic is not available in EasyCrypt, but in a pen-and-paper proof one would show it simply based on probability theory reasoning (e.g., Jensen's inequality, averaging). In the following, we sketch our solution for what we call probabilistic reflection.

(A word on terminology: The term *reflection* is used in many different ways in the literature. For example, in programming languages it often refers to the ability of a program to see its own structure at runtime. In proof assistants, reflection often means that one translates a term into abstract syntax and can reason about it inside the logic. Our probabilistic reflection is different from these, but we believe the term "reflection" is still justified because it allows us to expose the internal denotational semantics inside the program logic.)

Recall that there are no valid types which refer to distribution of final memories. These would be needed to give a type to the denotational semantics of a program, let alone define those semantics. However, in EasyCrypt each program has an associated variable $\mathcal{G}_A^{\mathbf{m}}$ (the type of $\mathcal{G}_A^{\mathbf{m}}$ is $\mathcal{G}_A$) which refers to the part of the memory **m** accessible by module *A*. It is guaranteed that running a program *A* will never change anything outside $\mathcal{G}_A^{\mathbf{m}}$. So "effectively", the semantics of a program can be described by looking only at the $\mathcal{G}_A$-part of a memory. So, we define a family of distributions D_A^g for g of type $\mathcal{G}_A$ such that $\mu_1(D_A^g, h)$ is the probability that we get h

in $\mathcal{G}_A$-part of the final memory configuration when starting with g in $\mathcal{G}_A$-part of the initial memory configuration.

Another problem is that EasyCrypt forbids to refer to the type $\mathcal{G}_A$ in the top-level definitions (global definitions of operators/constants). So, in particular, we cannot define a distribution D_A^g parameterized by $\mathcal{G}_A$-values in EasyCrypt. In our workaround to this problem, we prove lemmas of the existence of that family of distributions, but we do not define a constant referring to that family. This works since we only need to refer to the type of D_A^g locally in the theorem statement. Then, when reasoning, one can inside a proof refer to "the" distribution that exists by our lemmas.

In conclusion, our lemma for probabilistic reflection looks roughly as follows:

Theorem 1.2. *For all memories* $\boldsymbol{m}$ *and programs A there exists a family of distributions D_A^g (with g of type $\mathcal{G}_A$) such that for all predicates M on values of type $\mathcal{G}_A$:*

$$\Pr\left[A.main() @ \boldsymbol{m} : M(\mathcal{G}_A^{\boldsymbol{fin}}) \right] = \mu(D_A^{\mathcal{G}_A^{\boldsymbol{m}}}, M).$$

Here, **fin** is the final memory after execution of *A.main()* and $\mu(d, M)$ denotes the probability that the predicate M holds for values distributed according to d. Actual reflection lemma adds generality, e.g., referring also to the inputs/outputs of *A.main*, see Sec. 3.1.

However, being able to reflect the distribution corresponding to a given program is not enough. If we want to reason about composite programs, we will also need to understand how the different constructs in our language operate on the distributions. For example, given a program $A; B$, reflection gives us distributions D_{AB}, D_A, and D_B relating to the semantics of $(A; B)$, A, and B, respectively. However, we do not, a priori, know how D_{AB} is related to D_A and D_B. It is not even a priori clear whether *inside* the logic of EasyCrypt, it is possible to derive that relationship. Thus we prove additional lemmas for this and other cases that allow us to derive the distribution of a more complex program from the distributions of its components (see Sec. 3.3). For example, D_{AB} is shown to be the monadic bind of D_A and D_B.

Altogether this gives us a library for probabilistic reflection in EasyCrypt, independent of the results on rewinding below. See Sec. 3 for details.

1.3 Rewinding

The final challenge in the formal derivation of Thm. 1.1 is that in EasyCrypt, we cannot define a generic interface of modules which return their own state. Morally, we want *A.getState()* to return a value $\mathcal{G}_A^{\boldsymbol{m}}$ of type $\mathcal{G}_A$. However, this is impossible since the type $\mathcal{G}_A$ is only allowed to appear in the logical statements and program code of other modules but not in the code of the module A itself.

We solve the above problem by defining what it means for a module to be rewindable. In essence, a module is rewindable if and only if the state of the program can be encoded

as a bitstring (or equivalently, as any other countable type). In particular, a program with variables of type "real" (which is uncountable) would not be rewindable in that sense.[1] A security proof using rewinding would then only apply to rewindable adversaries which is not a restriction from the cryptographic point of view. (Typically, cryptographic adversaries are assumed to operate on data that is representable in a computer. Such data can always be encoded as a bitstring.)

In conclusion, our definition for rewindable modules (programs) roughly requires a module to have procedures *getState* and *setState*. The execution of *A.getState()* in state $\boldsymbol{m}$ must return the value $f(\mathcal{G}_A^{\boldsymbol{m}})$ where f is an arbitrary injective mapping from the type $\mathcal{G}_A$ to some parameter type *sbits*. The *setState* procedure gets an argument $x : sbits$ and sets $\mathcal{G}_A^{\boldsymbol{m}}$ to $f^{-1}(x)$ if $f^{-1}(x)$ is defined.

Altogether this gives an approach for working with rewindable adversaries in EasyCrypt. See Sec. 4 for details.

1.4 Rewinding in Other Verification Frameworks

To the best of our knowledge, rewinding has not been implemented in EasyCrypt, nor in other frameworks for reasoning about cryptographic proofs such as CryptHOL [7] in Isabelle, FCF [17] in Coq, CryptoVerif [9] in OCaml, Verypto [8] in Isabelle, CertiCrypt [6] in Coq, SSProve [1] in Coq.

A natural question is: Can the methods from the present work be ported to these other frameworks? Can rewinding be implemented in those frameworks in other ways? If our work on rewinding in EasyCrypt is any indication, then it is hard to make a reasonable estimate how difficult it is to implement (or even whether it is possible at all) before actually going through the process of doing it. Nonetheless, we will hazard some guesses:

CryptHOL, FCF, Verypto, CertiCrypt, and SSProve are foundational frameworks. That is, they are implemented inside the general purpose logic of existing frameworks such as Coq and Isabelle. This means that the details of the probabilistic semantics of the language are already directly accessible; probabilistic reflection as we did in EasyCrypt would not be necessary. Then it should be possible to implement the different result on reflection (Sec. 3) directly as theorems in Coq/Isabelle. (How hard or easy the proofs of these theorems would be is, like always in formal verification, hard to predict in advance.)

On the other hand, CryptoVerif is, like EasyCrypt, not foundational. This means that we cannot directly access the semantics, and thus would probably have to implement something akin to our probabilistic reflection. In addition, it is not clear whether the various facts about rewinding (Sec. 3) can even be expressed inside CryptoVerif. If not, the

---

[1]The restriction to countable types is arbitrary. We could choose a different, larger type sbits instead of bitstrings for encoding program states. However, we choose bitstrings because these are more natural in the computational setting – notions of runtime of algorithms apply naturally to bitstrings but not to uncountable types such as reals.

logic of these tools would have to be extended, and the tools rewritten accordingly.

But, as we stated, these are educated guesses only. To know for sure, one would have to actually put in the work and implement rewinding in these various frameworks.

2 Preliminaries

In this section we review the syntax and semantics of the main EasyCrypt constructs. Readers familiar with EasyCrypt can skip this section and just familiarize themselves with our syntactic conventions in this footnote[2]. The top-level definitions in EasyCrypt consist of types, operators, lemmas/axioms, module types, and modules. In EasyCrypt one can specify datatypes and operators, where types intuitively denote non-empty sets of values and operators are typed pure functions on these sets. EasyCrypt provides basic built-in types such as `unit`, `bool`, `int`, etc. The standard library includes formalizations of lists, arrays, finite sets, maps, probability distributions, etc. EasyCrypt also allows users to implement their own datatypes and functions (including inductive datatypes and functions defined by pattern matching). For example, we can give a definition of a polymorphic identity function as follows:

```
op id ['a] : 'a → 'a = λ x. x.
```

In this paper, for ease of readability, we use a more compact notation λ x. x for lambda-abstractions. In the original EasyCrypt code, this would be written as fun x => x.

The ambient logic in EasyCrypt is based on a classical (i.e., non-constructive) set theory which we can use to state and prove properties. (The term ambient logic refers to a built-in logic in EasyCrypt. Ambient logic is not specific to reasoning about programs). For example, we can prove that application of id to any x equals to x. In lemmas and axioms we will use symbols ∀ and ∃ instead of the EasyCrypt syntax which uses keywords forall and exists, respectively.

```
lemma id_prop ['a] : ∀ (x :'a), id x = x.
proof. trivial. qed.
```

In EasyCrypt, a proof starts with the keyword `proof`. The steps of the proof consist of tactic applications (e.g., `auto`, `trivial`, etc.) which either discharge the proof obligation or transform it into subgoal(s). The `qed` finishes the proof.

Types and operators without definitions are abstract and can be seen as parameters to the rest of the development. Parameters can additionally be restricted by axioms. For example, we can parameterize the development by a uniform distribution of elements of type `bits`.

---

[2]We write ← for both <- and <@, $\xleftarrow{\$}$ for <$, ∧ for /\, ∨ for \/, ≤ for <=, ∀ for forall, ∃ for exists, **m** for &m, $\mathcal{G}_A$ for glob A, $\mathcal{G}_A^{\mathbf{m}}$ for (glob A){m}, $\lambda x. x$ for fun x => x, × for *, t $\mathcal{L}$ for t list, t $\mathcal{D}$ for t distr, μ for mu, μ_1 for mu1, t $\mathcal{O}$ for t option. Furthermore, in Pr-expressions, in abuse of notation, we allow sequences of statements instead of a single procedure call. It is to be understood that this is shorthand for defining an auxiliary wrapper procedure containing those statements.

```
type bits.

op bD : bits 𝒟.
axiom bDU : is_uniform bD.
```

The theory containing this axiom can later be "cloned" and the operator bD instantiated with a value for which the axiom bDU is actually provable. This enables modular design of theories.

In this paper we will use notation `bits` $\mathcal{D}$ as a more concise version of the EasyCrypt syntax `bits distr` which denotes the type of distributions of `bits`. Similarly, we will use `bits` $\mathcal{L}$ instead of `bits list`, and `bits` $\mathcal{O}$ instead of `bits option`.

Distributions. In EasyCrypt, every type t is associated with the type t $\mathcal{D}$ of discrete distributions. A discrete distribution over type t is fully defined by its mass function. i.e. by a non-negative function f from elements of type t to reals so that $\Sigma_x f(x) \leq 1$. The probability mass function of distribution d can be accessed by writing μ_1 d (i.e., μ_1 d x is the probability assigned to element x by distribution d). Also, for any predicate P : t → bool we can write μ d P to get the total probability assigned by d to elements which satisfy P.

Modules. In EasyCrypt, modules consist of typed global variables and procedures. The set of all global variables of a module is the union of the set of global variables that are declared in that module and the set of all global variables (declared in other modules) which the module could read or write by a series of procedure calls beginning with a call of one of its procedures. In EasyCrypt, the whole memory (state) of a program is referred to by &m (or &n etc.). We can refer to the tuple of all global variables of the module A in &m as (glob A){m}. The type of all global variables of A (i.e., the type of (glob A){m}) is denoted by glob A. For readability, we will use syntax $\mathcal{G}_A$ for the type glob A. Memories &m will be typed in bold without the & (i.e., **m** for &m). And $\mathcal{G}_A^{\mathbf{m}}$ will denote the EasyCrypt value (glob A){m}.

For illustration, we implement the following example of a guessing-game module GG:

```
module GG = {
  var c,win,q,r
  proc init(x : int) = {
   (c,win,q) ← (0,false,x); // simultaneous assignment
  }
  proc guess(x : bits) : bool = {
    if (c < q){
      r $← bD;
      win ← win || r = x;
      c ← c + 1;
    }
    return win;
  }
}.
```

The module GG has three global variables: c and q of type int, and win of type bool. Hence, for any memory $\mathbf{m}$, $\mathcal{G}_{GG}^{\mathbf{m}}$ has type $\mathcal{G}_{GG}$ which equals to a product bool × int × int. The GG module allows a player to guess (call the GG.guess procedure) the next value sampled from distribution bD. The player has at most q attempts (set during initialization by procedure GG.init). The player wins if they guess correctly at least once.

Module types. In EasyCrypt, *module types* specify the types of a set of module procedures [4]. Therefore, module types in EasyCrypt are similar to interfaces in other programming languages (e.g., Java). We can specify the module type of GG as follows:

```
module type GuessGame = {
  proc init(x : int) : unit
  proc guess(x : bits) : bool
}.
```

Note that module types say nothing about the global variables a module could have and only specify the input and output types of the module procedures.

Next, we define a module type of protocol parties (adversaries), who receive an instance G of a guessing game as a module parameter. An adversary must have a play procedure which starts the game:

```
module type Adversary(G : GuessGame) = {
  proc play() : unit {G.guess}
}.
```

To forbid adversaries to reinitialize the game the play procedure can only execute the guess procedure of the parameter game G. This is optionally expressed by listing the allowed method(s) in the curly braces next to the procedure.

Probability expressions. EasyCrypt has Pr-constructs which can be used to refer to the probabilities of events in program executions: Pr[r ← X.p() @ m: M r] denotes the probability that the return value r of procedure p of module X given initial memory $\mathbf{m}$ satisfies the predicate M. (I.e., the general form is Pr[program @ initial memory: event].) The Pr-notation in EasyCrypt is somewhat restrictive, the program can only be a single procedure call. In our presentation, we relax this notation and allow multiple statements; it is to be understood that in the actual EasyCrypt code this is implemented by defining an auxiliary wrapper procedure that contains those statements.

For example, we can express that for any adversary A the probability of winning the guessing-game is equal to or smaller than $\frac{q}{n}$, where n is the size of the support of distribution used by GG and q is the maximal allowed number of guesses.

```
lemma winPr : ∀ (A <:Adversary {GG}) m q, 0 ≤ q
⇒ Pr[GG.init(q); A(GG).play() @ m: GG.win]
    ≤ q / (support_size bD).
```

(In EasyCrypt, X <: T states that the module X satisfies the module type T.) Note that the module type Adversary also includes adversaries who simply set the value GG.win to true. EasyCrypt allows us to write Adversary{GG} to denote a subset of adversaries who has disjoint set of global variables from the module GG.

3 Toolkit for Probabilistic Reflection

In this section, we discuss a derivation of probabilistic reflection for programs (i.e., modules) in EasyCrypt. Recall that by probabilistic reflection, we mean tools to get access to probabilistic denotational semantics of imperative programs inside EasyCrypt proofs. (Without needing any meta-reasoning.) Also, we use the probabilistic reflection to derive a powerful toolkit of lemmas which are common in pen-and-paper proofs when arguing about distributions underlying programs.

3.1 Probabilistic Reflection

Recall that in Sec. 1.2, we introduced Thm. 1.2 that proves the existence of a distribution corresponding to a program's denotational semantics. In EasyCrypt, we formally state this theorem as follows:

```
lemma reflection_simple : ∃ (D : 𝒢_A → 𝒢_A 𝒟),
  ∀ m M, μ (D 𝒢_A^m) M = Pr[A.main() @ m: M 𝒢_A^fin].
```

Here, D g corresponds to the family D_A^g from Thm. 1.2, i.e., D $\mathcal{G}_A^{\mathbf{m}}$ is the supposed distribution of final states of the program after running on initial memory $\mathbf{m}$.

Inside an EasyCrypt proof, this lemma could be used as elim (reflection_simple A) ⇒ D H_D; this will introduce a variable D in the environment of the proof, together with its defining property H_D stating the relationship between D and Pr[A.main()...].

However, reflection_simple as stated is not general enough for many purposes. In particular, if A.main takes an argument i or returns a value r then we cannot reason about the distribution of r and express how D depends on r. The following more general reflection lemma removes these limitations:

```
lemma reflection : ∃ (D : 𝒢_A → at → (rt × 𝒢_A) 𝒟),
  ∀ m M i, μ (D 𝒢_A^m i) M
    = Pr[r ← A.main(i) @ m: M (r, 𝒢_A^fin)].
```

The intuition behind this lemma and the previous one is the same. The only difference is that D now has an additional argument i, referring to the input of A.main, and the resulting distribution (D $\mathcal{G}_A^{\mathbf{m}}$ i) is a distribution over pairs (r, $\mathcal{G}_A^{\mathbf{fin}}$) of output and final memory.

Note that while EasyCrypt allows us to all-quantify over the module A and over the argument and input types (i.e., at and rt), we cannot quantify over the name main of the procedure. Similarly, the number of arguments of procedure main is fixed in the lemma. Fortunately, this only constitutes a minor inconvenience, not a real restriction because

we can always define a wrapper module A′ that has a procedure `main` with a single argument i (possibly of a tuple type). Then `main` can untuple i and invoke the procedure that we actually want to investigate. A simple call to the tactic `inline A′.main` in the proof will then unwrap this wrapper procedure.

In EasyCrypt, we directly prove `reflection` and derive `reflection_simple` as an immediate corollary. For readability, we proceed with describing a direct proof of `reflection_simple` instead.

Proof. We start the proof by defining a predicate P on probabilities which is parameterized by an initial state g of type $\mathcal{G}_A$ and an element x of type $\mathcal{G}_A$. Below we use the Easy-Crypt tactic `pose` to give a definition which is local to the proof.

```
pose P g x := λ p. ∀ n,
  𝒢ₐⁿ = g ⟹ Pr[r ← A.main() @ n: 𝒢ₐᶠⁱⁿ = x] = p.
```

The probability p satisfies the predicate (P g x) if it equals to the probability of A.main terminating in the state x by starting its run from a memory **n** which has $\mathcal{G}_A$ variables equal to g.

Before continuing with the proof, we explain that the standard library of EasyCrypt provides the formalization of the Axiom of Choice in the form of the operator `choiceb` and its corresponding property `choicebP`:

```
op choiceb ['a] : ('a → bool) → 'a.

axiom choicebP ['a] : ∀ (P : 'a → bool),
  (∃ (x : 'a), P x) ⟹ P (choiceb P).
```

It states that for any predicate P if there exists an element which satisfies it then the element denoted by (`choiceb P`) satisfies P. Here, it is worth mentioning that all propositions in EasyCrypt have type `bool`.

The next step of the proof is to define a function (Q g x) which uses the choice operator on the predicate (P g x) to assign a probability to x.

```
pose Q g x := choiceb (P g x).
```

The intuition is that (Q g x) returns "the" probability `Pr[A.main() @ n: 𝒢ₐᶠⁱⁿ = x]` for all **n** with $\mathcal{G}_A^n$ = g. Note that a priori we do not know that there is such a probability, because probability could depend on **n**. To show that (Q g x) is a well-defined we need to prove that the value (Q g x) satisfies the predicate (P g x). Because in the lemma, (D g) is only used for g of the form $\mathcal{G}_A^m$, we specifically need to show the following claim:

```
have Q_well_def : P 𝒢ₐᵐ x (Q 𝒢ₐᵐ x).
```

(Here we use the EasyCrypt tactic `have name : fact` which allows to locally prove the `fact` and call it `name`.)

If we can show that there exists a probability q, so that (P $\mathcal{G}_A^m$ x q) then the proof Q_well_def amounts to a simple application of the `choicebP` property. The obvious candidate for this probability q is the `Pr`-expression:

```
Pr[A.main() @ m: 𝒢ₐᶠⁱⁿ = x].
```

To show that this candidate satisfies (P $\mathcal{G}_A^m$ x), we must prove the following (by definition of P):

```
have good_q : ∀ n, 𝒢ₐⁿ = 𝒢ₐᵐ ⟹
    Pr[A.main() @ m: 𝒢ₐᶠⁱⁿ = x]
  = Pr[A.main() @ n: 𝒢ₐᶠⁱⁿ = x].
```

The `good_q` is intuitively simple and we prove it using the pRHL which is available in EasyCrypt.

Now, as we know that (Q g) assigns adequate probabilities to elements of type $\mathcal{G}_A$, we use the standard Easy-Crypt constructor `mk` which turns any function of type 'a → real into distribution of type 'a $\mathcal{D}$.

```
pose D g := mk (Q g).
```

The above defines a parameterized distribution D typed as $\mathcal{G}_A → \mathcal{G}_A$ $\mathcal{D}$. We skip the technical details of a proof which shows that D is a well-formed probability distribution. We only show the final derivation which proves that D is the denotation of A.main. To achieve this we show the pointwise equality of distribution D and `Pr`-expression. Let x be an element of type $\mathcal{G}_A$:

```
have pointwise :
    μ₁(D 𝒢ₐᵐ) x = Pr[A.main() @ m: 𝒢ₐᶠⁱⁿ = x].
```

The proof is as follows:

```
μ₁(D 𝒢ₐᵐ) x = μ₁(mk (Q 𝒢ₐᵐ)) x
            = Q 𝒢ₐᵐ x
            = choiceb (P 𝒢ₐᵐ x)
            = Pr[A.main() @ m: 𝒢ₐᶠⁱⁿ = x].
```

The first equality is by definition of D, in the second equality μ_1 cancels application of mk since D is a well-defined distribution (see muK from EasyCrypt standard library), the third equality is by definition of Q, the fourth equality is an application of the previously proved property Q_well_def.

At the first glance, it seems that this implies that (D $\mathcal{G}_A^m$) indeed describes the probability distribution corresponding to A.main. That is, we want:

```
have onsubs : μ (D 𝒢ₐᵐ) M = Pr[A.main() @ m: M 𝒢ₐᶠⁱⁿ].
```

meaning that the probability that a value sampled from (D $\mathcal{G}_A^m$) satisfies M equals the probability that the final state of A.main satisfies M. Unfortunately, this is not immediate. For example, hypothetically, the function M ↦ `Pr[A.main() @ m: M 𝒢ₐᶠⁱⁿ]` might not be a discrete probability measure and thus it might not be determined by its values on singleton sets. To show `onsubs`, we need to use one more trick: We define an auxiliary module and procedure P.sampleFrom such that P.sampleFrom(d) simply returns some x $\overset{\$}{\leftarrow}$ d. Then (μ (D $\mathcal{G}_A^m$) M) equals to `Pr[x ← P.sampleFrom(D 𝒢ₐᵐ) @ m: M x]` and we get:

```
have aux1 : Pr[x ← P.sampleFrom(D 𝒢ₐᵐ) @ m: M x]
  = Pr[A.main() @ m: M 𝒢ₐᶠⁱⁿ].
```

For goals of this shape, we use a combination of the `byequiv` and `bypr` tactics from EasyCrypt; `byequiv` changes this goal into a pRHL judgment relating the programs `A.main` and `P.sampleFrom`. And `bypr` converts such a pRHL judgment back to an equality of probabilities. It seems that we are back at `onsubs` now. However, the final equality is actually:

```
have aux2 : ∀ x,
  Pr[r ← P.sampleFrom(D 𝒢ᵐₐ) @1: r = x]
  = Pr[A.main() @2: 𝒢fin_A = x].
```

(for memories **1** and **2** that are equal to **m** in global variables of $\mathcal{G}_A$.) But this follows from `pointwise` proven above. □

We clarify that reflection results are proved "once and for all" – by using the EasyCrypt's module cloning mechanism the `reflection` lemma could be instantiated for arbitrary adversaries.

Note that in our proof we rely on the fact that the tactics `byequiv` and `bypr` in combination imply that the probability `Pr[...: M x]` (even for infinite M) can be related to λ `y. Pr[...: x = y]`.

3.2   Probabilistic Toolkit

In this section, we present well-known results from probability theory which we formalized and used extensively in our EasyCrypt development. In particular, recall that in Thm. 1.1 we used averaging and Jensen's inequality in the derivation of the key lemma needed for proving the security of a coin-toss protocol (see Sec. 5). In turn, proofs of Jensen's inequality and averaging depend on finite probabilistic approximation. (To the best of our knowledge, none of these results were previously formalized in EasyCrypt.)

Finite Pr-approximation. We prove that the support of a distribution can be finitely approximated with arbitrary precision. We formally prove finite approximation for distributions and then use the probabilistic reflection to extend this result to programs.

Let d be a distribution of type $'a$ $\mathcal{D}$. Then there exists a sequence of lists (L n) so that the probability that an element sampled from d is not in the list (L n) converges to 0 (for $n \to \infty$). (This holds for discrete distributions only.)

```
lemma fin_pr_approx _distr _conv ['a] :
  ∀ (d : 'a 𝒟), ∃ (L : int → 'a ℒ),
    convergeto (λ n. μ d (λ x. x ∉ L n)) 0.
```

Having the finite probabilistic approximation for distributions allows us to use the probabilistic reflection mechanism to extend the finite probabilistic approximation to programs. More specifically, let `A.main` be a procedure which takes an argument of type `at` and produces the result of type `rt`. In this case, there exists a sequence of lists (L n), so that the result and the final state produced by `A.main(i)` are not in (L n) with probability converging to 0.

```
lemma fin_pr_approx_prog_conv : ∀ m i,
```

```
  ∃ (L : int → (rt × 𝒢ₐ) ℒ),
  convergeto
   (λ n. Pr[r ← A.main(i) @m: (r,𝒢ₐ) ∉ L n]) 0.
```

Averaging. Averaging allows us to express the probability corresponding to a program `x ←$ d; A.main(x)` in terms of probabilities corresponding to the program `A.main(x)` and the probability assigned to x by distribution d. In this sense, the averaging technique can be seen as a generalized version of case-analysis.

In our EasyCrypt formalization we state and prove a general version of averaging for an arbitrary distribution d which also might have an infinite support:

```
lemma averaging : ∀ m M i d,
 Pr[x ←$ d; r ← A.main(x,i) @m: M r]
 = sum (λ x. (μ₁ d x) · Pr[r ← A.main(x,i) @m: M r]).
```

Here, the value (sum f) denotes $\Sigma_x f(x)$. (Note that the sum in EasyCrypt does not have to range over a finite set.)

Jensen's inequality. Jensen's inequality is another well-known result which is widely used in cryptography. In general, it relates the value of a convex function of an integral/sum to the integral/sum of the convex function. In the context of probability theory, it is generally stated in the following form: if X is a distribution, g maps elements of X to reals, and f is convex then f (E X g) ≤ E X (f ∘ g). Here, E X h is an expected value and ∘ denotes function composition.

We prove a slightly restricted version of Jensen's inequality. In particular, we assume that on the support of X the function g takes values in an interval between some parameter-values a and b and that f x takes values in an interval between parameter-values c and d if a ≤ x ≤ b. Also, the standard assumptions are that f is convex, that the distribution X is lossless (i.e., μ X (λ x. true) = 1), and that the expectations E X g and E X (f ∘ g) exist.

```
lemma Jensen_inf ['a] :
  ∀ (X : 'a 𝒟) g f (a b c d : real),
  is_lossless X ⟹ hasE X g ⟹ hasE X (f ∘ g)
 ⟹ (∀ (a b : real), (convex f a b))
 ⟹ (∀ x, a ≤ x ≤ b ⟹ c ≤ f x ≤ d)
 ⟹ (∀ x, x ∈ X ⟹ a ≤ g x ≤ b)
 ⟹ f (E X g) ≤ E X (f ∘ g).
```

(The EasyCrypt standard library derives Jensen's inequality for distributions with *finite* support only.)

3.3   Reflection of Composition

In this section we address the probabilistic reflection of the sequential composition of programs. For example, let us analyze the program: `r₁ ← A.ex₁(); A.ex₂(r₁)`. We can use the `reflection_simple` lemma from Sec. 3.1 to get access to a distribution D_{12} such that:

```
∀ m M, μ (D₁₂ 𝒢ᵐₐ) M
  = Pr[r₁ ← A.ex₁(); A.ex₂(r₁) @m: M 𝒢fin_A].
```

The distribution D_{12} corresponds to a composite program as a whole. However, being able to reflect the distribution corresponding to a composite program is not enough to enable reasoning about composite programs based on the properties of its components; we do not know how D_{12} is related to $A.ex_1$ and $A.ex_2$ separately.

In the following, we prove a lemma for reflection of composition which allows us to show that there exist distributions D_1 and D_2 which are the probabilistic reflection of procedures $A.ex_1$ and $A.ex_2$, and that the composition of D_1 and D_2 is D_{12}. So, the main goal is to prove lemmas that allow us to derive the distribution of a more complex program from the distributions which correspond to its components.

In EasyCrypt, the composition of distributions is implemented as an operator `dlet` which has the following type: $'a\ \mathcal{D} \rightarrow ('a \rightarrow 'b\ \mathcal{D}) \rightarrow 'b\ \mathcal{D}$. Intuitively, the distribution (`dlet d`$_1$ `d`$_2$) could be described imperatively as: $x_1 \xleftarrow{\$} d_1;\ x_2 \xleftarrow{\$} d_2\ x_1;$ `return` x_2.

We can formally state the theorem of reflection of composition as follows:

```
lemma refl_comp_simple :
 ∃ (D₁ : 𝒢_A → (rt₁ × 𝒢_A) 𝒟) (D₂ : 𝒢_A → rt₁ → 𝒢_A 𝒟),
 (∀ m M, μ (D₁ 𝒢_A^m) M
    = Pr[r₁ ← A.ex₁() @ m: M r₁]) ∧
 (∀ m M r₁, μ (D₂ (r₁,𝒢_A^m)) M
    = Pr[r₂ ← A.ex₂(r₁) @ m: M r₂]) ∧
  ∀ m M, μ (dlet (D₁ 𝒢_A^m) D₂) M
    = Pr[r₁ ← A.ex₁(); A.ex₂(r₁) @ m: M 𝒢_A^fin].
```

We give only a rough sketch of the proof. First, by using the `reflection` lemma from Sec. 3.1 we get distributions D_1 and D_2 which correspond to procedures $A.ex_1$ and $A.ex_2$, respectively. Next, we use pRHL reasoning to prove that the imperative composition of D_1 and D_2 corresponds to composition of $A.ex_1$ and $A.ex_2$:

```
Pr[x₁ ←$ D₁; x₂ ←$ D₂ x₁ @ m: M x₂]
 = Pr[r₁ ← A.ex₁(); A.ex₂(r₁) @ m: M 𝒢_A^fin].
```

Finally, we prove that the imperative composition of D_1 and D_2 corresponds to their declarative composition, namely, `dlet D`$_1$ `D`$_2$:

```
Pr[x₁ ←$ D₁; x₂ ←$ D₂ x₁ @ m: M 𝒢_A^fin]
 = μ (dlet (D₁ 𝒢_A^m) D₂) M.
```

This step uses averaging (see Sec. 3.2).

In EasyCrypt formalization we prove a stronger lemma `refl_comp` which generalizes `refl_comp_simple` in the following aspects:

- The procedures $A.ex_1$ and $A.ex_2$ take all-quantified arguments i_1 of type at_1 and i_2 of type at_2, respectively. As a result, the distributions D_1 and D_2 also become parameterized by values of types at_1 and at_2, respectively.
- The distribution $(D_2\ i\ g)$ is over pairs $(r, \mathcal{G}_A^{fin})$ of output of $A.ex_2$ and final memory (not just the final memory).

- In the event part of the probability expression (i.e., $\text{Pr}[...: M\ (r_1, r_2, \mathcal{G}_A^{fin})]$) we allow the predicate M to depend on the r_1 (output of $A.ex_1$), r_2 (output of $A.ex_2$), and the final memory $\mathcal{G}_A^{fin}$ (not just the final memory).

In EasyCrypt, we prove the following general version of reflection of composition:

```
lemma refl_comp :
 ∃ (D₁ : at₁ → 𝒢_A → (rt₁ × 𝒢_A) 𝒟)
   (D₂ : at₂ → rt₁ × 𝒢_A → (rt₂ × 𝒢_A) 𝒟),
 (∀ m M i₁, μ (D₁ i₁ 𝒢_A^m) M
   = Pr[r₁ ← A.ex₁(i₁) @ m: M (r₁,𝒢_A^fin)]) ∧
 (∀ m M r₁ i₂, μ (D₂ i₂ (r₁,𝒢_A^m)) (M r₁)
   = Pr[r₂ ← A.ex₂(i₂,r₁) @ m: M r₁ (r₂,𝒢_A^fin)]) ∧
  ∀ m M i₁ i₂, μ (dlet (D₁ i₁ 𝒢_A^m)
       (λ r. dmap (D₂ i₂ r) (λ x. (r.1,x)))) M
   = Pr[r₁ ← A.ex₁(i₂); r₂ ← A.ex₂(i₂,r₁)
       @ m: M (r₁,r₂,𝒢_A^fin)].
```

Here, (`dmap d f`) denotes the distribution of (`f x`) for $x \xleftarrow{\$} d$.

4 Rewinding

In Sec. 1, we briefly explained that rewinding is a commonly used technique which allows one module (i.e., program) to save a state of another module and also restore that state at some later time. More precisely, we say that a module A is rewindable iff:

1. There exists an injective mapping f from $\mathcal{G}_A$ to some parameter type `sbits`[3].
2. The module A must have a terminating procedure `getState`, so that whenever `A.getState` is called from the state $g\ :\ \mathcal{G}_A$, the result of the call must be equal to (`f g`) and the state of A must not change.
3. The module A must have a terminating procedure `setState`, so that whenever it is given an argument `x : sbits`, so that `x = f g` for some $g\ :\ \mathcal{G}_A$ then A must be set into a state `g`.

In EasyCrypt, we start formalizing this definition by defining a module type `Rew` for rewindable programs:

```
module type Rew = {
  proc getState() : sbits
  proc * setState(s : sbits) : unit
}.
```

(Here, the symbol $*$ indicates that the procedure (re)initializes all global variables of a module.)

---

[3]Intuitively, `sbits` is the type of bitstrings. To keep our development more general, we do not require this, but only assume the existence of embeddings `nat_sbits: nat` $\rightarrow$ `sbits` (to ensure that `sbits` is infinite) and `pair_sbits: sbits` $\times$ `sbits` $\rightarrow$ `sbits`. The EasyCrypt theory cloning mechanism makes it possible to later replace this `sbits` type by a concrete type such as lists of bits.

We formalize the rewinding properties of `getState` and `setState` procedures as a predicate `RewProp` on modules typeable as `Rew`. Unfortunately, in EasyCrypt we cannot define an operator like `RewProp` because its definition depends on a module which is not allowed. As a result, in the actual EasyCrypt code the workaround is to copy-and-paste the verbose definition of `RewProp(A)`. This reduces readability, but is conceptually the same.

```
op RewProp(A : Rew) : bool =
  ∃ (f : 𝒢_A → sbits), injective f ∧
  (∀ m, Pr[r ← A.getState() @ m:
           𝒢_A^fin = 𝒢_A^m ∧ r = f 𝒢_A^m] = 1) ∧
  (∀ m (g : 𝒢_A), Pr[A.setState(f g) @ m:
           𝒢_A^fin = g] = 1) ∧ islossless A.setState.
```

(In EasyCrypt, `islossless` `X.p` expresses that the procedure p of module X must terminate on all inputs.)

Then, whenever we do a proof using rewinding, we will need to explicitly assume that our adversary A satisfies property `RewProp(A)`, or, equivalently, quantify only over adversaries of module type `Rew` satisfying `RewProp(A)`.

The first litmus test of our definition of rewindability is to show that modules without global variables are (trivially) rewindable. For a module without global variables, $\mathcal{G}_A$ will be a singleton type (i.e., $\forall(x\ y\ :\ \mathcal{G}_A),\ x = y$). Then we can generically state that A will be rewindable, as long as we implement `getState` and `setState` as terminating procedures (it does not matter what they do):

```
lemma no_globs_rew : ∀ (A <: Rew),
 (∀ (x y : 𝒢_A), x = y)
 ⟹ islossless A.getState ∧ islossless A.setState
 ⟹ RewProp(A).
```

On the necessity of the `RewProp` axiom. The reader may wonder whether adding the explicit assumption in a security proof that the adversary A satisfies `RewProp(A)` does not weaken the security proof. After all, it means that security only holds with respect to such adversaries, but not with respect to adversaries that do not satisfy `RewProp(A)`. We argue that `RewProp(A)` is not a true restriction of the adversary, merely a requirement that the adversary has a certain interface with certain properties. The only actual restriction about the inner workings of the adversary that `RewProp(A)` makes is that the adversary's state can be encoded as a sequence of bits (`sbits`). Usually, in cryptography, we make even stronger assumptions about the adversary, namely that its state *is* a sequence of bits (or a Turing machine tape). In contrast, here we only assume that its state *can be encoded* as a sequence of bits.

We stress that we only need to make this assumption for abstract (i.e., all-quantified) adversaries. For adversaries that we explicitly construct as part of a reduction, we can actually prove `RewProp`, see the next section.

4.1 Transformations

Cryptographic proofs are commonly based on transformations of adversaries (or reduction of adversaries). In Easy-Crypt, a transformation is a module which receives other modules as parameters, defines its own global variables, and has procedures which can call procedures of its parameter-modules. Typically, one of the parameter-modules will be the original adversary.

In this section, we show how to prove rewindability of a module which is parameterized by rewindable modules and which has at most countable state. We illustrate this by implementing a module T which is parameterized by rewindable modules A and B and has a global variable x of a parameter type `ct`. As a result, the global state of module `T(A,B)` consist of variable `T.x` and all global variables of modules A and B. (i.e., $\mathcal{G}_{T(A,B)} = $ `ct` $\times \mathcal{G}_A \times \mathcal{G}_B$). Since, by the definition of rewindability, we need to embed elements of type $\mathcal{G}_{T(A,B)}$ into `sbits` then we parameterize our development by an injection from `ct` to `sbits`:

```
op ct_sbits : ct → sbits.
axiom bcu   : injective ct_sbits.

module T(A : Rew, B : Rew) : Rew = {
  var x : ct
  // add any other procedures here
  proc getState(): sbits = {
    var stateA, stateB, xsbits : sbits;
    stateA ← A.getState();
    stateB ← B.getState();
    xsbits ← ct_sbits x;
    return pair_sbits
      (xsbits,pair_sbits(stateA,stateB));
  }
  proc setState(state: sbits): unit = {
    var stateA, stateB, xsbits : sbits;
    (xsbits,s) ← unpair_sbits state;
    (stateA,stateB) = unpair_sbits s;
     A.setState(stateA);
     B.setState(stateB);
     x ← unct_sbits xsbits;
  }
}.
```

The procedure `getState` stores the global states of A and B in the local variables `stateA` and `stateB`, respectively. Then the global variable `T.x` is converted into `sbits` and saved in variable `xsbits`. The resulting state is an embedding of a nested tuple (`xsbits,(stateA,stateB)`) into `sbits`. (Recall that `pair_sbits` is an embedding `sbits` $\times$ `sbits` $\rightarrow$ `sbits`.)

The procedure `setState` receives an `sbits` argument which is then "untupled" into `sbits` variables `xsbits`, `stateA`, and `stateB`. The state of A is set by passing argument `stateA` to its implementation of `setState` procedure (similarly for B, *mutatis mutandis*). The variable `T.x`

is set to the preimage of `xsbits`. Finally, we use pHL to prove that `T(A,B)` is also rewindable:

```
lemma trans_rew : ∀ (A :> Rew) (B :> Rew{A}),
   RewProp(A) ∧ RewProp(B) ⇒ RewProp(T(A,B)).
```

Note, that in the statement of the lemma we additionally require a state of B to be disjoint from a state of A (i.e., `B :> Rew{A}`). This is required because in case of possibly overlapping states not all values of type $\mathcal{G}_{T(A,B)}$ are valid states. For example, if $\mathcal{G}_A^m$ and $\mathcal{G}_B^n$ have overlapping variables with different values then the value $(\mathcal{G}_A^m, \mathcal{G}_B^n, x)$ is typeable as $\mathcal{G}_{T(A,B)}$ (for any x of type `ct`), but does not represent a possible state of `T(A,B)`. Unfortunately, in EasyCrypt, it is not possible to express "consistency" of possibly overlapping states of abstract modules.

4.2 Multiplication Rule and Commutativity

The multiplication rule from probability theory states that the probability of independent events occurring simultaneously is found by multiplying the probabilities of each event.

In terms of probabilistic programs it is natural to say that an execution of a procedure `P.run` is independent of an execution of `Q.run` if after termination of `P.run` the state of Q is not affected. In EasyCrypt, it is easy to prove the *multiplication rule* for modules with disjoint states:

```
lemma rew_mult_simple : ∀ (P :> Runnable)
     (Q :> Runnable{P}) m M₁ M₂ i₁ i₂,
 Pr[r₁ ← P.run(i₁); r₂ ← Q.run(i₂) @m:
    M₁ r₁ ∧ M₂ r₂]
 = Pr[r₁ ← P.run(i₁) @m: M₁ r₁]
 · Pr[r₂ ← Q.run(i₂) @m: M₂ r₂].
```

However, if we need independent runs of the procedure(s) of the same module then we need rewinding. Recall, that the main goal of rewindability is to be able to restore the state of a module after running one of its procedures. Let A be a rewindable module (i.e., A satisfies `RewProp(A)`) with procedures ex_1 and ex_2 which take all-quantified arguments i_1: at_1 and i_2: at_2 and compute results r_1: rt_1 and r_2: rt_2, respectively. Let us analyze the following program.

(1) Save the initial state of A by `s ← A.getState()`.
(2) Run the procedure `r₁ ← A.ex₁(i₁)`.
(3) Restore the initial state by calling `A.setState(s)`.
(4) Run the procedure `r₂ ← A.ex₂(i₂)`.

First, we analyze the steps (1)–(3) as a standalone program. In particular we must show that the `getState` and the `setState` calls do not affect the result computed by $A.ex_1$ procedure and also show that the final state of A equals to its initial state.

```
lemma rew_clean : ∀ m M₁ i₁,
  Pr[s ← A.getState(); r₁ ← A.ex₁(i₁);
     A.setState(s) @m: M₁ r₁ ∧ 𝒢ₐᵐ = 𝒢ₐᶠⁱⁿ]
  = Pr[r₁ ← A.ex₁(i₁) @m: M₁ r₁].
```

(The proof requires only basic pHL tactics and rewindability axioms.) This result allows us to derive the *multiplication rule* which states that the probability of a joint event $M_1\ r_1 \wedge M_2\ r_2$ for the program (1)–(4) on memory `m` equals to the product of probabilities of events $M_1\ r_1$ and $M_2\ r_2$ occurring after independent runs of $A.ex_1$ and $A.ex_2$ on `m`, respectively. In EasyCrypt this is stated as follows:

```
lemma rew_mult_law : ∀ m M₁ M₂ i₁ i₂,
  Pr[s ← A.getState(); r₁ ← A.ex₁(i₁);
     A.setState(s); r₂ ← A.ex₂(i₂) @m:
     M₁ r₁ ∧ M₂ r₂]
  = Pr[r₁ ← A.ex₁(i₁) @m: M₁ r₁]
  · Pr[r₂ ← A.ex₂(i₂) @m: M₂ r₂].
```

In its essence, `rew_mult_law` is derived by a single call to the built-in `seq` tactic.

Commutativity. In its turn, the multiplication rule opens for us an easy route to proving commutativity for rewindable modules. Consider a program consisting of steps (1)–(4)–(3)–(2) (i.e., $A.ex_1$ and $A.ex_2$ calls are swapped). We can prove that it computes the same distribution of pairs (r_1, r_2) as the program (1)–(4).

```
lemma rew_comm_law_simple : ∀ m M i₁ i₂,
   Pr[s ← A.getState(); r₁ ← A.ex₁(i₁);
      A.setState(s); r₂ ← A.ex₂(i₂) @m: M (r₁,r₂)]
 = Pr[s ← A.getState(); r₂ ← A.ex₂(i₂);
      A.setState(s); r₁ ← A.ex₁(i₁) @m: M (r₁,r₂)].
```

In our formalization by using averaging we prove even stronger version of commitativity for programs which begin with state initializer (see lemma `rew_comm_law` in file `RewCommutes.ec` of the accompanying code [15]).

4.3 Rewinding with Initialization

In Thm. 1.1 we sketched a derivation of the equation which is needed to prove sum-binding property for commitments (see Sec. 5). More specifically, we analyzed a program which starts with an explicit state initializer, saves the resulting state of module A, runs a procedure `A.run` for the first time, restores the saved state, and then runs the `A.run` procedure for the second time. We proved that the probability of a success (according to some predicate) in two sequential runs of `A.run` is lower-bounded by a square of probability of a success in a "initialize-then-run" case (i.e., initialize the state and execute the `A.run` procedure once).

In EasyCrypt, we derive a similar equation, but for a more general case:

- The initialization is done with a procedure `B.init`, where B is a module with a state which can possibly intersect with the state of module A.
- The initialization produces a result r_0 of a parameter type which is then supplied to `A.run`.
- The procedure `B.init` receives all-quantified argument `i` of a parameter type.

- The procedure `A.run` returns a result of a parameter type `rt`. The success of a run is defined by a parameter predicate `M (r₀,rᵢ)`, where r_0 and r_i are the values returned by `B.init` and `A.run` procedures, respectively.

The EasyCrypt statement of the lemma is as follows:

```
lemma rew_with_init : ∀ m M i,
 Pr[r₀ ← B.init(i); s ← A.getState();
    r₁ ← A.run(r₀); A.setState(s);
    r₂ ← A.run(r₀) @m: M (r₀,r₁) ∧ M (r₀,r₂)]
 ≥ Pr[r₀ ← B.init(i); r ← A.run(r₀) @m: M (r₀,r)]².
```

We skip the proof as it roughly follows the steps sketched in Thm. 1.1.

5 Case Study: Coin-Toss Protocol

As a case study for our techniques we prove the security of a coin-toss protocol based on bit-commitment. Historically, Blum described the problem of coin-toss protocol with the following example: *Alice and Bob are recently divorced, living in two separate cities, and want to decide who gets to keep the car. To decide, Alice wants to flip a coin over the telephone. However, Bob is concerned that if he were to tell Alice the result of his coin toss, she would adjust hers and automatically tell him that she wins.* Thus, the problem with Alice and Bob is that they do not trust each other; the only resource they have is the telephone communication channel, and there is not a third party available to read the coin [10].

In the following, we describe the coin-toss protocol based on a bit-commitment scheme which is similar to the original Blum's solution to the coin-toss problem:

1. Alice chooses a random bit r_1 and then generates a commitment `c` containing that bit (let `d` be the respective opening).
2. Alice sends the commitment `c` to Bob.
3. Bob chooses a random bit r_2 and sends it to Alice.
4. Alice opens her commitment by sending the bit r_1 and the opening `d` to Bob.
5. Bob verifies that `d` is a valid opening of r_1 for `c`. Otherwise Bob aborts.
6. Alice and Bob compute the final bit as $r_1 \oplus r_2$ (xor).

The coin-toss protocol must ensure the following property: if at least one of the parties correctly generates a random bit, then the final bit will be (nearly) random.

Security of the coin-toss is almost immediate if the commitment scheme satisfies a property called "sum-binding" in [18]. This property says that the probability of Alice opening the commitment to `false` and the probability of Alice opening it to `true` add to at most 1 (plus a negligible error). This property in turn is implied by the usual "computationally binding" property which says that Alice cannot open to both `false` and `true` *simultaneously* (except with negligible probability). Showing that "computationally binding"

implies "sum-binding", however, requires rewinding. Therefore that proof is a prime candidate for our case-study. (In the post-quantum setting, for example, computationally binding does *not* imply sum-binding [2]. This illustrates that this seemingly trivial implication is not as easy as it might seem, and that we indeed need rewinding here.)

5.1 Commitments

The standard library of EasyCrypt defines the module type `CommitmentScheme` which requires a scheme S to implement the following procedures:

1. `p ← S.gen()` generates the public key of a commitment scheme (also known as the public parameters).
2. `(c,d) ← S.commit(p,m)` produces commitment-opening pair for a message `m` and a public key `p`.
3. `b ← S.verify(p,m,c,d)` returns `b = true` iff `d` is a valid opening for message `m`, commitment `c`, and public key `p`.

For our development, we additionally require the existence of a verification function `Ver` (an "operator" in EasyCrypt-parlance) which must agree with the procedure `S.verify` on all arguments:

```
op Ver : pubkey × msg × commitment × opening
                                    → bool.

axiom verify_det : ∀ m a,
  Pr[r ← S.verify(a) @m: r = Ver a] = 1.
```

This means that verification is side-effect free (and deterministic). Otherwise, two runs of the verification algorithm could interfere with each other (and with calls to `S.commit`) and give different results.

In cryptography, a commitment scheme is called *computationally binding* iff the probability that adversary A can produce a commitment with openings of two different messages is negligible. The EasyCrypt standard library defines a module type `Binder` with a single procedure `bind`; we can then define the probability of success of adversary `A : Binder` in the "binding-game":

```
op binding_pr(A,m) = Pr[p ← S.gen();
 (c,m₁,d₁,m₂,d₂) ← A.bind(p);
 v₁ ← S.verify(p,m₁,c,d₁);
 v₂ ← S.verify(p,m₂,c,d₂) @m: v₁ ∧ v₂ ∧ m₁ ≠ m₂].
```

(Here, `binding_pr` is only a shortcut notation used in this text.) Hence, scheme is binding iff `binding_pr(A,m)` is negligible for all A and `m`.

Sum-Binding. Next, we define the "sum-binding" property of commitments. Let A be an adversary and p_b be a probability that A can open the commitment to contain b given input `b = false,true`. The commitment scheme is *sum-binding* iff for all such adversaries the $p_f + p_t \leq 1 + \epsilon$, where ϵ is negligible. We define a module type `SumBinder` with procedures `commit` and `open`. Then we define the probability of success of adversary `A : SumBinder` in the "sum-binding-game":

```
op sum_binding_pr(A,m)  =
 Pr[p ← S.gen(); c ← A.commit(p);
    d ← A.open(false); v ← S.verify(p,false,c,d)
   @ m: v]
 + Pr[p ← S.gen(); c ← A.commit(p);
      d ← A.open(true); v ← S.verify(p,true,c,d)
   @ m: v].
```

(Again, `sum_binding_pr` is only a shortcut notation.) Hence, scheme is sum-binding if and only if for all A and m, there exists a negligible ϵ, so that we can show that `sum_binding_pr(A,m)` $\leq$ 1 + ϵ. Before addressing the sum-binding property for commitments, we prove a more generic sum-binding inequality which shows that the sum of probabilities of success of independent runs of arbitrary procedures A.ex_1 and A.ex_2 is related to the probability of joint success in the same run.[4] More specifically, assume that module A is rewindable and B.init is some initialization procedure. We let p_1 be the probability that after initialization the procedure A.ex_1 succeeds according to some predicate M (similarly for p_2 and A.ex_2, *mutatis mutandis*). In this case, we can prove that the sum of probabilities $p_1 + p_2$ is upper-bounded by a sum 1 + 2 · q, where q is the probability that A.ex_1 and A.ex_2 both succeed in the same run (i.e., both starting from the same initial state produced by B.init). In EasyCrypt, we state this equation as follows:

```
lemma sum_binding_generic : ∀ m M i,
   Pr[r₀ ← B.init(i); r ← A.ex₁(r₀) @ m: M r]
 + Pr[r₀ ← B.init(i); r ← A.ex₂(r₀) @ m: M r]
   ≤ 1 + 2 · Pr[r₀ ← B.init(i); s ← A.getState();
              r₁ ← A.ex₁(r₀); A.setState(s);
              r₂ ← A.ex₂(r₀) @ m: M r₁ ∧ M r₂].
```

The proof revolves around `rew_with _init` inequality.

Equipped with the generic sum-binding inequality, we can now finish the proof that binding commitment schemes are also sum-binding. We start by implementing a reduction R(A) which runs A.commit (to produce the commitment), stores the state of A, runs A.open(false) (to produce the first opening), restores the state of A, and runs A.open(true) (to produce the second opening). Then we show that the probability that R(A) produces two valid openings (i.e., breaks binding) is lower-bounded in terms of the probability that A is successful in producing one valid opening.

```
module R(A : SumBinder) : Binder = {
  proc bind(p : pubkey) = {
    var c,s,d₁,d₂;
    c ← A.commit(p);
    s ← A.getState();
    d₁ ← A.open(false);
    A.setState(s);
    d₂ ← A.open(true);
```

---

[4]This generic lemma may also be useful when analyzing extractors for proof of knowledge protocols with two challenges, e.g., the zero-knowledge protocols for Hamiltonian cycles [11] and graph isomorphism [12].

```
    return (c,false,d₁,true,d₂);
  }
}.
```

Next, we implement wrapper-modules B and A′, so that B.init is a wrapper around the "commitment initialization" phase p ← S.gen(); c ← A.commit(p). The procedure A′.ex_1 is defined as A.open(false), and A′.ex_2 as A.open(true). In this case, sum-binding for commitments becomes an immediate consequence of the inequality `sum_binding_generic` and we can conclude:

```
lemma commitment_sum_binding : ∀ m,
  sum_binding_pr(A,m) ≤ 1 + 2 · binding_pr(R(A),m).
```

5.2 Coin-Toss Protocol

Recall, that a coin-toss protocol is considered secure if it is ensured that if at least one of the parties correctly generates a random bit then the final bit will be (nearly) random.

In the first case, we assume that Alice is honest and Bob is cheating. To simplify this case, we additionally assume that the commitment scheme is perfectly hiding. This means that Bob gets no information about r_1 after receiving the commitment c. Therefore, if Alice follows the protocol honestly and r_1 is uniformly random and independent of r_2 (due to the perfect hiding) then the bit ($r_1 \oplus r_2$) is also uniformly random. (The case of cheating Bob does not involve rewinding and is therefore not the focus of this paper.)

In the second case, we are left to show that if Bob honestly follows the protocol, then for any Alice (adversary A : CoinTossAlice) the resulting bit is nearly uniform. Below we assume that module type CoinTossAlice requires a module to have procedures commit and toss, where commit produces a commitment c, and toss gets a Bob's bit r_2 as an argument and then computes a bit together with its opening for c. We write `coin_toss_pr(A,m,b)` to denote a probability of A being able to open the commitment to Boolean b.

```
op coin_toss_pr(A,m,b) =
  Pr[p ← S.gen(); r₂ ←$ {0,1}; c ← A.commit(p);
     (r₁,d) ← A.toss(r₂) @ m:
     Ver (p,r₁,c,d) ∧ r₁ ⊕ r₂ = b].
```

(Once more, `coin_toss_pr` is only a shortcut notation.) We define $B_f(A)$ and $B_t(A)$ as the transformations of coin-toss adversary into an adversary that breaks binding for the cases b = false and b = true, respectively:

```
//getState and setState procedures are skipped
module Bf(A : CoinTossAlice) : SumBinder = {
  proc commit(p : pubkey) = {
    return A.commit(p);
  }
  proc open(x : bool) = {
    var d, r₁;
    (r₁,d) ← A.toss(x);
    return d;
  }
```

```
}.
module Bₜ(A : CoinTossAlice) : SumBinder = {
  proc commit(p : pubkey) = {
    return A.commit(p);
  }
  proc open(x : bool) = {
    var d, r₁;
    (r₁,d) ← A.toss(not x);
    return d;
  }
}.
```

$B_f(A)$ delegates the commitment generation to A and when asked to open a commitment to bit x then x is submitted to A.toss and the resulting opening is returned. $B_t(A)$ is different in that the negation of x is submitted to A.toss. Finally, we can derive that if Bob is honest then for any Alice the resulting bit is nearly uniform.

```
lemma coin_toss_alice: ∀ m b, coin_toss_pr(A,m,b)
  ≤ 1/2 + max binding_pr(R(Bₜ(A)),m)
              binding_pr(R(B_f(A)),m).
```

Proof. We start the proof with analysis of the case when b = true (i.e., $r_1 \oplus r_2 = true$). We prove that this case is upper-bounded by $1/2 + \epsilon$, where ϵ is the probability of breaking the binding of S by $R(B_t(A))$.

```
have coin_toss_alice_t : coin_toss_pr(A,m,true)
                ≤ 1/2 + binding_pr(R(Bₜ(A)),m).
```

We prove this case by arguing as follows:

```
Pr[p ← S.gen(); r₂ ←$ {0,1}; c ← A.commit(p);
   (r₁,d) ← A.toss(r₂) @ m:
   Ver (p,r₁,c,d) ∧ r₁ ⊕ r₂ = true]
⁼⁽¹⁾ 1/2 · (Pr[p ← S.gen(); c ← A.commit(p);
          (r₁,d) ← A.toss(false) @ m:
          Ver (p,r₁,c,d) ∧ r₁ ⊕ false = true]
      + Pr[p ← S.gen(); c ← A.commit(p);
          (r₁,d) ← A.toss(true) @ m:
          Ver (p,r₁,c,d) ∧ r₁ ⊕ true = true])
≤⁽²⁾ 1/2 · (Pr[p ← S.gen(); c ← A.commit(p);
          (r₁,d) ← A.toss(false) @ m:
          Ver (p,true,c,d)]
      + Pr[p ← S.gen(); c ← A.commit(p);
          (r₁,d) ← A.toss(true) @ m:
          Ver (p,false,c,d)])
⁼⁽³⁾ 1/2 · (Pr[p ← S.gen(); c ← A.commit(p);
          d ← B(A).open(true) @ m:
          Ver (p,false,c,d)]
    + Pr[p ← S.gen(); c ← A.commit(p);
        d ← B(A).open(false) @ m: Ver (p,true,c,d)])
≤⁽⁴⁾ 1/2 + binding_pr(R(Bₜ(A)),m).
```

(Here {0,1} denotes a uniform distribution of Booleans.) Step (1) is by case distinction of r_2. Step (2) is by simplification and event-inclusion. Step (3) is by definition of transformation B_t. Step (4) applies commitment_sum_binding from Sec. 5.1. In a similar way we handle the case when b = false and show:

```
have coin_toss_alice_f : coin_toss_pr(A,m,false)
                ≤ 1/2 + binding_pr(R(B_f(A)),m).
```

Finally, coin_toss_alice is a trivial consequence of coin_toss_alice_t and coin_toss_alice_f. □

6 Conclusions

In this paper we focused on probabilistic reflection and rewindability of adversaries. First, we implemented a powerful toolkit for probabilistic reflection which includes finite probabilistic approximation, averaging, and reflection of composition inside EasyCrypt. Second, we described a notion of rewindable adversaries and derived their basic properties: transformations, multiplication rule, commutativity, rewinding with initialization. Third, by combining these results together we were able to derive a generic sum-binding equation for arbitrary rewindable computations. Fourth, we instantiated the sum-binding property for commitments and proved that if a commitment scheme is binding then it is also sum-binding. Finally, we used this result to prove the security of a bit-commitment based coin-toss protocol.

To the best of our knowledge, probabilistic reflection, rewindable adversaries, and security of a coin-toss protocol have not yet been addressed in theorem provers.

Acknowledgments

References

[1] Carmine Abate, Philipp G. Haselwarter, Exequiel Rivas, Antoine Van Muylder, Théo Winterhalter, Cătălin Hriţcu, Kenji Maillard, and Bas Spitters. 2021. SSProve: A Foundational Framework for Modular Cryptographic Proofs in Coq. In *CSF 2021*, Vol. 1. IEEE Computer Society, Los Alamitos, CA, USA, 1–15. https://doi.org/10.1109/CSF51468.2021.00048

[2] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. 2014. Quantum Attacks on Classical Proof Systems: The Hardness of Quantum Rewinding. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*. 474–483. https://doi.org/10.1109/FOCS.2014.57

[3] Gilles Barthe, George Danezis, Benjamin Grégoire, César Kunz, and Santiago Zanella-Béguelin. 2013. Verified Computational Differential Privacy with Applications to Smart Metering. In *2013 IEEE 26th Computer Security Foundations Symposium*. 287–301. https://doi.org/10.1109/CSF.2013.26

[4] 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*. Springer, 146–166.

[5] Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. 2011. Computer-aided security proofs for the working cryptographer. In *Annual Cryptology Conference*. Springer, 71–90.

[6] Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. 2009. Formal certification of code-based cryptographic proofs. In *Proceedings*

of the 36th annual ACM SIGPLAN-SIGACT symposium on Principles of programming languages. 90–101.

[7] David A. Basin, Andreas Lochbihler, and S. Reza Sefidgar. 2020. CryptHOL: Game-based proofs in higher-order logic. *Journal of Cryptology* 33, 2 (2020), 494–566.

[8] Matthias Berg. 2013. *Formal verification of cryptographic security proofs*. Ph.D. Dissertation. Saarland University. https://doi.org/10.22028/D291-26528

[9] Bruno Blanchet. 2008. A computationally sound mechanized prover for security protocols. *IEEE Transactions on Dependable and Secure Computing* 5, 4 (2008), 193–207.

[10] Manuel Blum. 1983. Coin flipping by telephone: a protocol for solving impossible problems. *ACM SIGACT News* 15, 1 (1983), 23–27.

[11] Manuel Blum. 1987. How to prove a theorem so no one else can claim it. In *In: Proceedings of the International Congress of Mathematicians*. 1444–1451.

[12] Manuel Blum, Paul Feldman, and Silvio Micali. 1988. Non-Interactive Zero-Knowledge and Its Applications. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing* (Chicago, Illinois, USA) *(STOC '88)*. Association for Computing Machinery, New York, NY, USA, 103–112. https://doi.org/10.1145/62212.62222

[13] Véronique Cortier, Constantin Cătălin Drăgan, François Dupressoir, Benedikt Schmidt, Pierre-Yves Strub, and Bogdan Warinschi. 2017. Machine-checked proofs of privacy for electronic voting protocols. In *2017 IEEE Symposium on Security and Privacy (SP)*. IEEE, 993–1008.

[14] Denis Firsov, Henri Lakk, and Ahto Truu. 2021. Verified Multiple-Time Signature Scheme from One-Time Signatures and Timestamping. In *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*. IEEE Computer Society, Los Alamitos, CA, USA, 653–665. https://doi.org/10.1109/CSF51468.2021.00051

[15] Denis Firsov and Dominique Unruh. 2021. Reflection, Rewinding, and Coin-Toss in EasyCrypt – accompanying EasyCrypt code. https://github.com/dfirsov/easycrypt-rewinding/tree/cpp2022. https://doi.org/10.5281/zenodo.5760917

[16] Jakob Nussbaumer. 2019. *Security analysis for IPsec with EasyCrypt*. Master's thesis. University of Bonn.

[17] Adam Petcher and Greg Morrisett. 2015. The foundational cryptography framework. In *International Conference on Principles of Security and Trust*. Springer, 53–72.

[18] Dominique Unruh. 2016. Computationally Binding Quantum Commitments. In *Advances in Cryptology – EUROCRYPT 2016*, Marc Fischlin and Jean-Sébastien Coron (Eds.). Springer Berlin Heidelberg, 497–527.