

Proved Generation of Implementations from Computationally Secure Protocol Specifications*

David Cadé and Bruno Blanchet
INRIA Paris-Rocquencourt
23 avenue d'Italie, 75013 Paris, France
{david.cade,bruno.blanchet}@inria.fr

August 23, 2013

Abstract

In order to obtain implementations of security protocols proved secure in the computational model, we previously proposed the following approach: we write a specification of the protocol in the input language of the computational protocol verifier CryptoVerif, prove it secure using CryptoVerif, then generate an OCaml implementation of the protocol from the CryptoVerif specification using a specific compiler that we have implemented. However, until now, this compiler was not proved correct, so we did not have real guarantees on the generated implementation. In this paper, we fill this gap. We prove that this compiler preserves the security properties proved by CryptoVerif: if an adversary has probability p of breaking a security property in the generated code, then there exists an adversary that breaks the property with the same probability p in the CryptoVerif specification. Therefore, if the protocol specification is proved secure in the computational model by CryptoVerif, then the generated implementation is also secure.

Keywords: cryptographic protocol, computational model, implementation, compiler, CryptoVerif, OCaml, verification

1 Introduction

The verification of security protocols is an important research area since the 1990s: the design of security protocols is notoriously error-prone, and errors can have serious consequences. Formal verification first focused on verifying formal specifications of protocols. However, verifying a specification does not guarantee that the protocol is correctly implemented from this specification.

*This paper is an extended version of the work originally presented at the 2nd Conference on Principles of Security and Trust (POST 2013), Rome, Italy, March 2013 [9].

It is therefore important to make sure that the implementation is secure, and not only the specification. Moreover, two models were considered for verifying protocols. In the symbolic model, the so-called Dolev-Yao model, messages are terms. This abstract model facilitates automatic proofs. In contrast, in the computational model, typically used by cryptographers, messages are bitstrings and attackers are polynomial-time probabilistic Turing machines. Proofs in the latter model are more difficult than in the former, but yield a much more precise analysis of the protocol. Therefore, we would like to obtain implementations of protocols proved secure in the computational model.

To reach this goal, we have proposed the following approach in [8]. We start from a formal specification of the protocol. In order to prove the specified protocol secure in the computational model, we rely on the automatic protocol verifier CryptoVerif [5, 7, 6]. This verifier can prove secrecy and authentication properties. The generated proofs are proofs by sequences of games, like the manual proofs written by cryptographers. These games are formalized in a probabilistic process calculus. In order to obtain a proved implementation from the specification, we have written a compiler that takes a CryptoVerif specification and returns an implementation in the functional language OCaml (<http://caml.inria.fr>). This compiler starts from a CryptoVerif specification annotated with implementation details: the annotations specify how to divide the protocol in different roles, for example, key generation, server, and client, and how to implement the various cryptographic primitives and types. The compiler then generates an OCaml module for each role in the input file. In order to get a full implementation of the protocol, this module is combined with manually written network code, responsible in particular for sending and receiving messages from the network. From the point of view of security, the network code can be considered as part of the adversary, so we do not need to prove its security.

To make sure that the generated implementation is actually secure, we need to prove the correctness of our compiler. This proof was still missing in [8]. It is the topic of this paper. To make this proof, we needed a formal semantics of OCaml. We adapted the operational small-step semantics of a core part of OCaml by Owens et al. [15]. We added to this language support for simplified modules, multiple threads where only one thread can run at any given time, and communication between threads by a shared part of the store.

An adversary against the generated implementation is an OCaml program using the modules generated by our compiler. On the CryptoVerif side, an adversary is a process running in parallel with the verified protocol. In our proof, for each OCaml adversary, we construct a corresponding CryptoVerif adversary that simulates the behavior of the OCaml adversary. When the OCaml adversary calls one of the functions generated by our compiler, which comes from an oracle in the CryptoVerif process, the CryptoVerif adversary calls this oracle. Then we establish a precise correspondence between the traces of the CryptoVerif process with that CryptoVerif adversary and the traces of the OCaml program. This correspondence allows us to show that the probability of success of an attack is the same on the CryptoVerif side and on the OCaml side. There-

fore, if CryptoVerif proves that the protocol is secure, then the generated OCaml implementation is also secure, and the bound on the probability of success of an attack computed by CryptoVerif is also valid for the implementation.

We have made several assumptions to obtain this proof; the most important ones are:

- A1. The cryptographic primitives are correct with respect to the assumptions made on them in the specification.
- A2. The roles are executed in the order specified in CryptoVerif (e.g., in a key-exchange protocol, the key generation is called before the servers and clients).
- A3. The adversary and the network code do not access files created by our implementation (e.g. private key files).
- A4. The network code is a well-typed OCaml program, which does not use unsafe OCaml functions to bypass the type system.
- A5. The network code does not mutate bitstrings passed to or received from generated code. This property can be guaranteed by representing bitstrings by an immutable OCaml type. However, the most natural type for representing bitstrings is the OCaml type `string`, which is mutable. Immutable strings can be implemented in OCaml using an abstract type instead of `string`. In our semantics, strings are immutable values.
- A6. Our semantics of threads is obeyed, which implies that two processes that read or write the same file are not run concurrently (which can be enforced using locks), and that one cannot fork in the middle of a role.

Related work. Several approaches have been considered in order to obtain proved implementations of security protocols. In the symbolic model, several approaches generate protocols from specifications, e.g. [14, 16]. Other approaches analyze implementations by extracting a specification verified by a symbolic protocol verifier, e.g. [4, 1], or analyze them by other tools such as the model-checker ASPIER [10], the general-purpose C verifier VCC [11], or typing [3, 17].

In contrast, the following approaches provide computational security guarantees, by analyzing implementations. The tool FS2CV [13] translates a subset of $F\#$ to the input language of CryptoVerif, which can then prove the protocol secure. The tool F7 [3], which uses a dependent type system to prove security properties on protocols implemented in $F\#$, has been adapted to the computational model in [12]; it uses type annotations to help the proof. The symbolic execution approach of [1] provides computational security guarantees by applying a computational soundness result, which however restricts the class of protocols that can be considered. The tool of [2] generates a CryptoVerif model from a C implementation; however, it can analyze only a single execution path.

To the best of our knowledge, our approach is the first one for generating implementations with a computational proof. The work of [2] and ours are the only ones to provide an explicit bound on the probability of success of an attack against the verified protocol implementation.

Outline. Section 2 describes the common input language of CryptoVerif and of our compiler. Section 3 describes OCaml, the output language of our compiler. Section 4 describes the instrumentation of the OCaml semantics that we use to facilitate our proof. Section 5 describes the compiler itself. Section 6 presents our proof. Because of the length of the proof, the details of the proof of some lemmas are postponed to the appendix.

Notations. Let us introduce some basic notations. When f is a function, we denote by $\text{Dom}(f)$ the domain of f , that is, the set of elements x such that $f(x)$ is defined. We denote by $f[x \mapsto y]$ the function f' defined by $f'(x) = y$ and $f'(x') = f(x')$ for $x' \neq x$. When f_1 and f_2 are functions with disjoint domains, we denote by $f_1 \cup f_2$ the function f' defined by $f'(x) = f_1(x)$ if $x \in \text{Dom}(f_1)$ and $f'(x) = f_2(x)$ if $x \in \text{Dom}(f_2)$. When f_1 and f_2 are functions, we write $f_1 \subseteq f_2$ (or $f_2 \supseteq f_1$) when $\text{Dom}(f_1) \subseteq \text{Dom}(f_2)$ and, for all $x \in \text{Dom}(f_1)$, we have $f_2(x) = f_1(x)$. We denote by \emptyset the function whose domain is the empty set \emptyset .

2 The CryptoVerif Input Language

This section presents the syntax and semantics of the CryptoVerif input language, as well as the annotations that specify implementation details. CryptoVerif supports two input languages: the channel and oracle front-ends. The channel front-end [5] uses channels to pass data between the adversary and the protocol, and the oracle front-end [7] defines oracles that can be called by the adversary. In this paper, we focus on the oracle front-end, which is closer to the syntax of games used by cryptographers; oracles are also easier to translate into OCaml functions. (Our compiler also supports the channel front-end.) We adapt the semantics given in [5] for the channel front-end to the oracle front-end.

2.1 Syntax and Informal Semantics

Let us first introduce the syntax of the CryptoVerif language in Figure 1. The language is typed, and types T are subsets of $\text{bitstring}_\perp \stackrel{\text{def}}{=} \text{bitstring} \cup \{\perp\}$ where bitstring is the set of all bitstrings and \perp is a symbol that is not a bitstring, used, for example, to represent the failure of a decryption. The boolean type $\text{bool} \stackrel{\text{def}}{=} \{\text{true}, \text{false}\}$, where true is the bitstring 1 and false 0, and the types bitstring and bitstring_\perp are predefined.

Variables $x[i_1, \dots, i_m]$ represent arrays of bitstrings of a given type T indexed by the values of the indices i of the replications `foreach $i \leq N$ do Q` present above the definition of the variable. We call these indices *replication indices*, and we

$M ::=$	terms
$x[\tilde{i}]$	variable access
$f(M_1, \dots, M_m)$	function application
$Q ::=$	oracle definitions
0	nil
$Q \mid Q'$	parallel composition
foreach $i \leq N$ do Q	replication N times
$O[\tilde{i}](x_1[\tilde{i}] : T_1, \dots, x_k[\tilde{i}] : T_k) := P$	oracle definition
$P ::=$	oracle bodies
return $(M_1, \dots, M_k); Q$	return
end	end
$x[\tilde{i}] \stackrel{R}{\leftarrow} T; P$	random number
$x[\tilde{i}] \leftarrow M; P$	assignment
if M then P else P'	conditional
insert $Tbl(M_1, \dots, M_k); P$	insert in table
get $Tbl(x_1[\tilde{i}], \dots, x_k[\tilde{i}])$ suchthat M in P else P'	get from table
event $e(M_1, \dots, M_k); P$	event
let $(x_1[\tilde{i}] : T_1, \dots, x_{k'}[\tilde{i}] : T_{k'}) = O[M_1, \dots, M_i](M'_1, \dots, M'_k)$ in P else P'	oracle call
let $x[\tilde{i}] : T = \text{loop } O[M_1, \dots, M_i](M')$ in P else P'	loop

Figure 1: Syntax of the CryptoVerif language

abbreviate i_1, \dots, i_m by \tilde{i} . Each function f comes with its type $T_1 \times \dots \times T_m \rightarrow T$; all CryptoVerif functions are deterministic and efficiently computable. Some functions are predefined, and some are infix, like the equality test $=$ and boolean operations. The cryptographic primitives used in the protocol are represented by CryptoVerif functions. Terms M represent computations over bitstrings: they can be variable accesses $x[i_1, \dots, i_m]$ or function applications $f(M_1, \dots, M_m)$.

The oracle definitions Q represent the oracles that will become available to the adversary at this point. The nil construct 0 provides no oracle. The parallel composition $Q \mid Q'$ provides oracles in Q and Q' . The replication **foreach** $i \leq N$ **do** Q provides N copies of Q , indexed by $i \in \{1, \dots, N\}$. The bound N is unspecified and is used by CryptoVerif to express the maximum probability of breaking the protocol, which typically depends on the number of calls to the various oracles. The oracle definition $O[\tilde{i}](x_1[\tilde{i}] : T_1, \dots, x_k[\tilde{i}] : T_k) := P$ makes available the oracle $O[\tilde{i}]$; when $O[\tilde{i}]$ is called by the adversary with arguments a_1, \dots, a_k , it executes the oracle body P with $x_j[\tilde{i}]$ set to a_j .

The oracle bodies P represent the behavior of the oracle. A return statement **return** $(M_1, \dots, M_k); Q$ returns the result of M_1, \dots, M_k to the caller, and makes

available oracles in Q . An end statement `end` returns to the caller with an error. A random number assignment $x[\tilde{i}] \stackrel{R}{\leftarrow} T; P$ stores a uniformly chosen random value of type T in variable $x[\tilde{i}]$, and continues by executing P . The type T must consist of all bitstrings of a given size; in this case, we say that T is a *fixed-length* type. An assignment $x[\tilde{i}] \leftarrow M; P$ puts the result of M in the variable $x[\tilde{i}]$, and continues by executing P . A conditional statement `if M then P else P'` executes P if M evaluates to true and P' otherwise.

An insert statement `insert $Tbl(M_1, \dots, M_k)$` ; P inserts the result of M_1, \dots, M_k into the table Tbl . Tables are lists of tuples, used for example to store tables of keys. Each table Tbl has a type $T_1 \times \dots \times T_k$, which means that Tbl contains k -uples a_1, \dots, a_k such that a_j is of type T_j for all $j \leq k$. A get statement `get $Tbl(x_1[\tilde{i}], \dots, x_k[\tilde{i}])$ suchthat M in P else P'` searches for an element a_1, \dots, a_k in the table Tbl such that the term M evaluates to true when $x_1[\tilde{i}] = a_1, \dots, x_k[\tilde{i}] = a_k$. If there is no such element, we continue by executing P' , and otherwise we choose randomly one of the elements that correspond, store it in the variables $x_1[\tilde{i}], \dots, x_k[\tilde{i}]$, then execute P . An event statement `event $e(M_1, \dots, M_k)$` ; P is used to log events. Events serve for specifying security properties of protocols, but do not change the execution of the process.

An oracle call `let $(x_1[\tilde{i}] : T_1, \dots, x_{k'}[\tilde{i}] : T_{k'}) = O[M_1, \dots, M_l](M'_1, \dots, M'_k)$ in P else P'` calls oracle $O[M_1, \dots, M_l]$ with arguments M'_1, \dots, M'_k , stores its returned values in the variables $x_1[\tilde{i}], \dots, x_{k'}[\tilde{i}]$, and continues by executing P if the oracle terminates with a `return` statement, or continues by executing P' if the oracle terminates with `end`.

A loop `let $x[\tilde{i}] : T = \text{loop } O[M_1, \dots, M_l](M')$ in P else P'` calls oracle O in a loop. Oracle O takes a unique argument (the internal state of the loop) and returns a pair containing the modified internal state of the loop and a boolean indicating whether the loop should continue or not. $O[M_1, \dots, M_l](M')$ is first called. If it returns (a_1, true) , $O[M_1 + 1, M_2, \dots, M_l](a_1)$ is called. If it returns (a_2, true) , $O[M_1 + 2, M_2, \dots, M_l](a_2)$ is called, and so on, until $O[M_1 + k, M_2, \dots, M_l](a_k)$ returns (a_{k+1}, false) . Then we run P with $x[\tilde{i}]$ set to a_{k+1} . If O terminates with `end`, we run P' . Oracle call and loop statements cannot appear in the CryptoVerif process representing the protocol, but are used for representing the adversary.

Example 1 Let us consider a simple protocol in which the first participant Alice generates a nonce m , sends it to the second participant Bob with a signature of the nonce under Alice's signature key sk . Bob then verifies that the signature is correct using Alice's public key pk . This protocol can be described by the

following CryptoVerif process:

```

Ostart() :=
   $rk \stackrel{R}{\leftarrow} \textit{keyseed}; pk \leftarrow \textit{pkgen}(rk); sk \leftarrow \textit{skgen}(rk);$ 
  return( $pk$ ); (foreach  $i_1 \leq N_1$  do  $P_A$  | foreach  $i_2 \leq N_2$  do  $P_B$ )
 $P_A \stackrel{\text{def}}{=} \text{OA}() :=$ 
   $m \stackrel{R}{\leftarrow} \textit{nonce}; s \stackrel{R}{\leftarrow} \textit{seed}; \text{return}(m, \textit{sign}(m, sk, s))$ 
 $P_B \stackrel{\text{def}}{=} \text{OB}(m' : \textit{nonce}, s' : \textit{signature}) :=$ 
  if check( $m', pk, s'$ ) then return() else end

```

The only callable oracle at the beginning is the oracle Ostart, which generates the signature key pair (pk, sk) by first generating a random seed rk and applying the key generation algorithms \textit{pkgen} and \textit{skgen} to it. We return to the attacker the public key, so that the attacker can check whether a signature signed with the signature key sk is correct. When the oracle Ostart returns, one can call the oracle OA N_1 times, and the oracle OB N_2 times.

The oracle OA generates a random nonce m and a random seed s , and returns the nonce m and the signature of the nonce m under the signature key sk with the random seed s .

The oracle OB takes as arguments a nonce m' and a signature s' , which should be the elements returned by a call to oracle OA, and checks using the function check whether the signature s' is indeed a correct signature of the message m' under the signature key sk by using the public key pk . If the signature is correct, the oracle returns normally. Otherwise, the oracle terminates with end.

2.2 Formal Semantics

We present the semantics of the language in Figures 2, 3, and 4. The semantics is defined as a reduction relation on semantic configurations, which are tuples of the form $\mathfrak{C} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$.

- The environment E is a mapping from array cells $x[\tilde{a}]$ to their contents, where x is a variable, \tilde{a} gives the value of its replication indices, and the contents of $x[\tilde{a}]$ is a bitstring value.
- The oracle body P is the oracle body currently running.
- The mapping \mathcal{T} maps table names to their contents, which is the list of elements inserted in the table.
- The set \mathcal{Q} contains the set of the callable oracle definitions.
- The list \mathcal{R} is the call stack, which consists of triplets containing the variables with which the result should be bound and two oracle bodies, the

Terms:

$$\begin{array}{c}
E, a \Downarrow a \quad \text{(Cst)} \\
\\
\frac{x[a_1, \dots, a_m] \in \text{Dom}(E)}{E, x[a_1, \dots, a_m] \Downarrow E(x[a_1, \dots, a_m])} \quad \text{(Var)} \\
\\
\frac{\forall j \leq m, E, M_j \Downarrow a_j \quad f : T_1 \times \dots \times T_m \rightarrow T \quad \forall j \leq m, a_j \in T_j}{E, f(M_1, \dots, M_m) \Downarrow f(a_1, \dots, a_m)} \quad \text{(Fun)}
\end{array}$$

Oracle definitions:

$$\begin{array}{c}
\text{reduce}(0) \stackrel{\text{def}}{=} \emptyset \quad \text{(Nil)} \\
\text{reduce}(Q_1 \mid Q_2) \stackrel{\text{def}}{=} \text{reduce}(Q_1) \cup \text{reduce}(Q_2) \quad \text{(Par)} \\
\text{reduce}(\text{foreach } i \leq n \text{ do } Q) \stackrel{\text{def}}{=} \bigcup_{a=1}^n \text{reduce}(Q\{a/i\}) \quad \text{(Repl)} \\
\text{reduce}(O[\tilde{i}](x_1[\tilde{i}] : T_1, \dots, x_k[\tilde{i}] : T_k) := P) \stackrel{\text{def}}{=} \\
\{(O[\tilde{i}](x_1[\tilde{i}] : T_1, \dots, x_k[\tilde{i}] : T_k) := P)\} \quad \text{(Oracle)}
\end{array}$$

Figure 2: Semantics (1)

first will be executed if the oracle returns a result with a return statement, and the second will be executed if the oracle terminates with an end statement.

- The list \mathcal{E} is the list of events $e(a_1, \dots, a_k)$ executed so far, by the construct event $e(M_1, \dots, M_k)$.

During execution, terms may be reduced into constant bitstrings, so we add constant bitstrings a to the grammar of terms M . The notation $E, M \Downarrow a$ means that the term M evaluates to the bitstring a under the environment E . This relation is defined by rules (Cst), (Var), and (Fun) in Figure 2. The set $\text{reduce}(Q)$, also defined in Figure 2, contains all oracle definitions provided by the oracle definition Q , with replication indices instantiated to all their possible values.

The semantics is defined by probabilistic reduction rules between configurations: $\mathfrak{C} \rightarrow_p \mathfrak{C}'$ means that \mathfrak{C} reduces into \mathfrak{C}' with probability p . This relation is defined in Figures 3 and 4. We use the following notations for lists. Let $[]$ be the empty list, and $x :: l$ be the list obtained by adding the element x to the list l . Let $[x_1; \dots; x_k]$ be the list $x_1 :: \dots :: x_k :: []$. Let $[x \in l \mid \text{Prop}(x)]$ be the list containing all elements x of l that satisfy the property $\text{Prop}(x)$, in the same

Oracle bodies (1):

$$\frac{T \text{ fixed-length type} \quad a \in T}{E, x[\tilde{a}'] \stackrel{R}{\leftarrow} T; P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_{\frac{1}{|T|}} E[x[\tilde{a}'] \mapsto a], P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}} \quad (\text{New})$$

$$\frac{E, M \Downarrow a}{E, x[\tilde{a}'] \leftarrow M; P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_1 E[x[\tilde{a}'] \mapsto a], P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}} \quad (\text{Let})$$

$$\frac{E, M \Downarrow \text{true}}{E, \text{if } M \text{ then } P \text{ else } P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_1 E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}} \quad (\text{If1})$$

$$\frac{E, M \Downarrow \text{false}}{E, \text{if } M \text{ then } P \text{ else } P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_1 E, P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}} \quad (\text{If2})$$

$$\frac{\forall j \leq k, \quad E, M_j \Downarrow a_j}{E, \text{insert } Tbl(M_1, \dots, M_k); P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_1 E, P, \mathcal{T}[Tbl \mapsto (a_1, \dots, a_k) :: \mathcal{T}(Tbl)], \mathcal{Q}, \mathcal{R}, \mathcal{E}} \quad (\text{Insert})$$

$$\frac{l = [(a_1, \dots, a_k) \in \mathcal{T}(Tbl) \mid E[x_1[\tilde{a}'] \mapsto a_1, \dots, x_k[\tilde{a}'] \mapsto a_k], M \Downarrow \text{true}] \quad (a_1^0, \dots, a_k^0) \in l \quad S = \{1 \leq j \leq |l| \mid \text{nth}(l, j) = (a_1^0, \dots, a_k^0)\}}{E, \text{get } Tbl(x_1[\tilde{a}'], \dots, x_k[\tilde{a}']) \text{ suchthat } M \text{ in } P \text{ else } P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_{\sum_{j \in S} \text{among}(\{1, \dots, |l|\}, j)} E[x_1[\tilde{a}'] \mapsto a_1^0, \dots, x_k[\tilde{a}'] \mapsto a_k^0], P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}} \quad (\text{Get1})$$

$$\frac{[(a_1, \dots, a_k) \in \mathcal{T}(Tbl) \mid E[x_1[\tilde{a}'] \mapsto a_1, \dots, x_k[\tilde{a}'] \mapsto a_k], M \Downarrow \text{true}] = []}{E, \text{get } Tbl(x_1[\tilde{a}'], \dots, x_k[\tilde{a}']) \text{ suchthat } M \text{ in } P \text{ else } P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_1 E, P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}} \quad (\text{Get2})$$

Figure 3: Semantics (2)

Oracle bodies (2):

$$\begin{array}{c}
\forall i \leq l, E, M_i \Downarrow a'_i \quad \tilde{a}' = a'_1, \dots, a'_l \quad \forall j \leq k, E, N_j \Downarrow b_j \\
\exists x'_1, \dots, x'_k, P'' \text{ such that} \\
Q_0 = (O[\tilde{a}'](x'_1[\tilde{a}'] : T'_1, \dots, x'_k[\tilde{a}'] : T'_k) := P'') \in \mathcal{Q} \\
E' = E[x'_1[\tilde{a}'] \mapsto b_1, \dots, x'_k[\tilde{a}'] \mapsto b_k] \\
\hline
E, \text{ let } (x_1[\tilde{a}] : T_1, \dots, x_{k'}[\tilde{a}] : T_{k'}) = O[M_1, \dots, M_l](N_1, \dots, N_k) \\
\text{ in } P \text{ else } P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_1 \\
E', P'', \mathcal{T}, \mathcal{Q} \setminus \{Q_0\}, ((x_1[\tilde{a}], \dots, x_{k'}[\tilde{a}]), P, P') :: \mathcal{R}, \mathcal{E} \\
\forall j \leq k, E, N_j \Downarrow b_j \quad \mathcal{Q}' = \text{reduce}(Q'') \\
\hline
E, \text{ return}(N_1, \dots, N_k); Q'', \mathcal{Q}, ((x_1[\tilde{a}], \dots, x_k[\tilde{a}]), P, P') :: \mathcal{R}, \mathcal{E} \\
\rightarrow_1 E[x_1[\tilde{a}] \mapsto b_1, \dots, x_k[\tilde{a}] \mapsto b_k], P, \mathcal{T}, \mathcal{Q} \cup \mathcal{Q}', \mathcal{R}, \mathcal{E} \\
E, \text{ end}, \mathcal{T}, \mathcal{Q}, ((x_1[\tilde{a}], \dots, x_{k'}[\tilde{a}]), P, P') :: \mathcal{R}, \mathcal{E} \rightarrow_1 E, P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \quad (\text{End}) \\
\hline
\forall j \leq l, E, M_j \Downarrow a_j \\
E, \text{ event } ev(M_1, \dots, M_l); P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_1 E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, e(a_1, \dots, a_l) :: \mathcal{E} \\
\quad (\text{Event})
\end{array}$$

$$\begin{array}{c}
\forall i \leq l, E, M_i \Downarrow a'_i \quad E, M' \Downarrow c \\
\text{the last replication above the definition of } O \text{ is foreach } i_1 \leq N_1 \quad a'_1 \leq N_1 \\
\hline
E, \text{ let } r[\tilde{a}] : T = \text{loop } O[M_1, \dots, M_l](M') \text{ in } P \text{ else } P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \\
\rightarrow_1 E, \\
(\text{let } (r'_{a'_1, r}[\tilde{a}] : T, b_{a'_1, r}[\tilde{a}] : \text{bool}) = O[a'_1, \dots, a'_l](c) \text{ in} \\
\text{ if } b_{a'_1, r}[\tilde{a}] \text{ then} \\
(\text{let } r[\tilde{a}] : T = \text{loop } O[a'_1 + 1, a'_2, \dots, a'_l](r'_{a'_1, r}[\tilde{a}] : T) \\
\text{ in } P \text{ else } P') \\
\text{ else } r[\tilde{a}] \leftarrow r'_{a'_1, r}[\tilde{a}]; P \\
\text{ else } P'), \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \\
\quad (\text{Loop1})
\end{array}$$

$$\begin{array}{c}
\forall i \leq l, E, M_i \Downarrow a'_i \quad E, M' \Downarrow c \\
\text{the last replication above the definition of } O \text{ is foreach } i_1 \leq N_1 \quad a'_1 > N_1 \\
\hline
E, \text{ let } r[\tilde{a}] : T = \text{loop } O[M_1, \dots, M_l](M') \text{ in } P \text{ else } P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \\
\rightarrow_1 E, P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \\
\quad (\text{Loop2})
\end{array}$$

Figure 4: Semantics (3)

order as in l . This construct is defined by induction on lists:

$$\begin{aligned}
[x \in [] \mid Prop(x)] &\stackrel{\text{def}}{=} [], \\
[x \in y :: l \mid Prop(x)] &\stackrel{\text{def}}{=} \begin{cases} [x \in l \mid Prop(x)] & \text{if } \neg Prop(y), \\ y :: [x \in l \mid Prop(x)] & \text{otherwise.} \end{cases}
\end{aligned}$$

The concatenation of lists $l_1 @ l_2$ is the list containing all elements of l_1 followed by all elements of l_2 . The membership test $x \in l$ is true when l contains the element x , and false otherwise. Let $|l|$ be the length of the list l , and $\text{nth}(l, n)$ be the n th element of list l .

The rule (New) evaluates $x[\tilde{a}'] \stackrel{R}{\leftarrow} T$ by choosing an element $a \in T$ and storing it in $E(x[\tilde{a}'])$. The element $a \in T$ is chosen uniformly, so the probability of each choice is $1/|T|$ and this is possible only when T is a fixed-length type. The rule (Let) evaluates the term M and stores its value in $E(x[\tilde{a}'])$. The rules (If1) and (If2) are straightforward.

The rules (Insert), (Get1), and (Get2) deal with tables of keys. The rule (Insert) evaluates the inserted element and adds it to the table Tbl , by adding it to the list $\mathcal{T}(Tbl)$. The rules (Get1) and (Get2) compute the list of elements that satisfy the condition of the `get`. When this list is empty, the `else` branch is taken by rule (Get2). When this list is not empty, the rule (Get1) chooses an element of this list l , stores it in $E(x_1[\tilde{a}']), \dots, E(x_k[\tilde{a}'])$, and takes the `in` branch. The j -th element of the list l is chosen with probability $\text{among}(\{1, \dots, |l|\}, j)$, where $\text{among}(S, b)$ is the probability that the element $b \in S$ is chosen among elements of the set S , according to an almost uniform distribution: we require that, for every set S , $\sum_{b \in S} \text{among}(S, b) = 1$, $\text{among}(S, b) > 0$ for all $b \in S$, and $\sum_{b \in S} \left| \text{among}(S, b) - \frac{1}{|S|} \right| \leq \epsilon$ for some $\epsilon > 0$. Indeed, probabilistic Turing machines can choose random elements uniformly only in sets of cardinal a power of 2. For other sets, they can choose random elements with a probability distribution as close as we wish to uniform, that is, we can make ϵ as small as we wish in the formula above. In case the same element a_1^0, \dots, a_k^0 occurs several times in the list l , the probability of choosing that element is the sum of the probabilities of all its occurrences. The probability of choosing a_1^0, \dots, a_k^0 is then close to $m/|l|$, where m is the number of times this element appears in l .

The rule (Call) implements the oracle call `let` $(x_1[\tilde{a}] : T_1, \dots, x_{k'}[\tilde{a}] : T_{k'}) = O[M_1, \dots, M_l](N_1, \dots, N_k)$ in P else P' . It evaluates the indices M_1, \dots, M_l of the oracle to call into \tilde{a}' and its arguments N_1, \dots, N_k into b_1, \dots, b_k ; after evaluation, we want to call the oracle $O[\tilde{a}'](b_1, \dots, b_k)$. Then, it looks for the definition Q_0 of the oracle $O[\tilde{a}']$ in the callable oracles \mathcal{Q} . It calls Q_0 by removing it from the callable oracles, storing b_1, \dots, b_k in the arguments of Q_0 , and running its body P'' . The element $(x_1[\tilde{a}], \dots, x_{k'}[\tilde{a}], P, P')$ is pushed on the stack \mathcal{R} : $x_1[\tilde{a}], \dots, x_{k'}[\tilde{a}]$ are the variables in which the return value of Q_0 should be stored, P is the process to execute when Q_0 returns, and P' is the process to execute when Q_0 terminates with `end`. The rule (Return) pops an element $((x_1[\tilde{a}], \dots, x_{k'}[\tilde{a}], P, P')$ from the stack, stores the return value in $x_1[\tilde{a}], \dots, x_{k'}[\tilde{a}]$, and executes P . It adds to the set of callable oracles \mathcal{Q} the

oracles \mathcal{Q}' defined in the oracle definition Q'' located after the return statement. The rule (End) also pops an element $((x_1[\tilde{a}], \dots, x_{k'}[\tilde{a}]), P, P')$ from the stack, but executes the process P' . The rule (Event) adds the executed event to the list of events \mathcal{E} .

The rules (Loop1) and (Loop2) implement the loop statement. The rule (Loop1) performs one iteration of the loop. To that effect, it creates two fresh variable names $r'_{a'_1, r}$ and $b_{a'_1, r}$, calls the oracle O and stores its return values in these variables. When the boolean $b_{a'_1, r}[\tilde{a}]$ returned by O is false, it ends the loop and continues by executing P with the result $r[\tilde{a}]$ bound to the value of $r'_{a'_1, r}[\tilde{a}]$. When $b_{a'_1, r}[\tilde{a}]$ is true, it reruns the loop. If the oracle O terminates with an end statement, it ends the loop and continues by executing P' . The rule (Loop2) handles the case in which the loop stops by reaching the bound N_1 of the loop index.

The initial configuration for running the oracle definition Q is $\mathfrak{C}_i(Q_0) \stackrel{\text{def}}{=} \emptyset, \text{let } x[] : \text{bitstring} = O_{\text{start}}() \text{ in return}(x) \text{ else end}, \mathcal{T}_0, \text{reduce}(Q_0), \emptyset, []$, where $\mathcal{T}_0(Tbl) = []$ for all tables Tbl . This configuration starts by calling oracle O_{start} . The oracle definition Q_0 typically contains a protocol in parallel with an adversary.

CryptoVerif verifies the following requirements on Q_0 :

Property 2.1 *Variables are renamed so that each variable has a single definition. The indices \tilde{i} of a variable $x[\tilde{i}]$ are always the indices of replications above the definition of x .*

Property 2.2 *The processes are well-typed. (In particular, functions and oracles receive arguments of their expected types. For brevity, we do not detail the type system; see [5] for a similar type system.)*

Property 2.3 *Oracles with the same name can be defined only in different branches of an if or get construct. In an oracle definition $O[\tilde{i}](x_1[\tilde{i}] : T_1, \dots, x_k[\tilde{i}] : T_k) := P$, the indices \tilde{i} are always the indices of replications above that oracle definition.*

Property 2.4 *We define types of oracles as follows. The type of a $\text{return}(M_1, \dots, M_k); Q$ statement consists of the types of M_1, \dots, M_k and the list of types of the oracle definitions at the beginning of Q , ordered from left to right. The type of an oracle definition consists of the oracle name, the bounds of the replications above that oracle definition, the types of the arguments of the oracle, and the common type of its return statements.*

An oracle may have several return statements, but they must be of the same type. When there are several definitions of an oracle with the same name O , they must be of the same type.

Property 2.1 makes sure that a distinct array cell is used in each copy of a process, so that all values of the variables during execution are kept in memory. (This helps in cryptographic proofs.) To lighten notations, we often omit the indices since they are determined by Property 2.1. Property 2.2 requires

the adversary to be well-typed. This requirement does not restrict its computing power, because it can always define type-cast functions $f : T \rightarrow T'$ to bypass the type system. Similarly, the type system does not restrict the class of protocols that we consider, since the protocol may contain type-cast functions. The type system just makes explicit which set of bitstrings may appear at each point of the protocol. Property 2.4 guarantees that the various definitions of an oracle are consistent, and can in fact be compiled into a single function in OCaml. Property 2.3 guarantees that there exists a single callable definition for each oracle. This property is formalized by the following lemma, proved in Appendix A.

Lemma 2.5 (Oracle name and indices unicity) *If the configuration $\mathcal{C} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$ is reachable from the initial configuration $\mathcal{C}_i(Q_0)$ by reductions \rightarrow_p , then the set of callable oracles \mathcal{Q} contains at most one oracle with a given name O and given replication indices \tilde{a} .*

This lemma proves that the rule (Call) is deterministic. Therefore, all rules are deterministic, except the rules (New) and (Get1) which may make probabilistic choices.

Another interesting point is that, if a configuration \mathcal{C} reduces into another configuration, then the sum of the probabilities of all the possible reductions from \mathcal{C} is 1:

$$\sum_{\{\mathcal{C}' | \mathcal{C} \rightarrow_p(\mathcal{C}') \mathcal{C}'\}} p(\mathcal{C}') = 1.$$

Definition 2.6 (Traces) *Let us denote traces with the symbol \mathcal{CT} . A trace is a sequence of reductions $\mathcal{CT} = \mathcal{C}_0 \rightarrow_{p_1} \dots \rightarrow_{p_n} \mathcal{C}_n$ where $\mathcal{C}_0, \dots, \mathcal{C}_n$ are semantic configurations such that $\mathcal{C}_i \rightarrow_{p_{i+1}} \mathcal{C}_{i+1}$ for $i = 0, \dots, n-1$.*

A complete trace is a trace such that there is no possible reduction from its last configuration.

The probability of the trace \mathcal{CT} is $\Pr[\mathcal{CT}] = p_1 \times \dots \times p_n$. When the traces in a set of traces $\mathcal{CT}\mathcal{S}$ are not prefix of one another, the probability of $\mathcal{CT}\mathcal{S}$ is the sum of the probabilities of its elements.

The notation $\mathcal{C} \rightarrow_p^ \mathcal{C}'$ means that there exists a trace beginning at \mathcal{C} and ending at \mathcal{C}' , and p is the probability of the set of all traces beginning at \mathcal{C} and stopping at their first occurrence of \mathcal{C}' .*

The notation $\mathcal{C} \rightarrow_p^+ \mathcal{C}'$ means that $\mathcal{C} \rightarrow_p^ \mathcal{C}'$ and $\mathcal{C} \neq \mathcal{C}'$, that is, all traces from \mathcal{C} to \mathcal{C}' have at least one step.*

The notation $\mathcal{C} \rightarrow^ \mathcal{C}'$ means $\mathcal{C} \rightarrow_1^* \mathcal{C}'$. We denote the number of steps in the trace \mathcal{CT} as $|\mathcal{CT}| = n$.*

Intuitively, when traces in $\mathcal{CT}\mathcal{S}$ are not prefix of one another, they correspond to disjoint cases, so the probability of $\mathcal{CT}\mathcal{S}$ is the sum of probabilities of the traces in $\mathcal{CT}\mathcal{S}$. (When \mathcal{CT} is a prefix of \mathcal{CT}' , the trace \mathcal{CT} is a particular case of \mathcal{CT}' .)

In CryptoVerif, since for every reduction with a probabilistic choice, the environment E is modified so that we can determine from E which reduction

was used, and one cannot remove elements from E , there will be at most one trace from one configuration to another. However, the notations of Definition 2.6 are also used for OCaml where there could be several configurations reducing to the same configuration, so they support this situation.

2.3 Annotations

In order to compile a CryptoVerif process into an implementation, we added annotations to the language, to specify implementation details.

First, we separate the parts of the process that correspond to different roles, such as client and server, which will be included in different OCaml programs in the generated implementation. We annotate processes to specify roles: the beginning of role is specified by oracle definitions $\text{role}\{Q$; the end of role is specified by a closing brace $\}$ between a $\text{return}(\dots)$ and its following oracle definition Q . We denote by $Q(\text{role})$ the part of the process corresponding to the role role .

The process for a role $Q(\text{role})$ may have free variables, but CryptoVerif requires that these free variables be defined under no replication, so that they can be passed from the process that defines them to the process $Q(\text{role})$, which uses them, simply by storing each variable in a file. (There must be a single value to store, not one for each value of the replication indices. Storing variables in files is useful for variables that are communicated across roles, for example long-term keys that are set in a key generation program and later used by client and server programs.) The user must also declare, for each free variable $x[]$ in a role, the file f in which the variable will be stored. Let files be the set of these pairs $(x[], f)$. Let also tables be the set of pairs (Tbl, f) such that the table Tbl will be stored in file f .

Example 2 Let us annotate the protocol of Example 1.

```

keygen[ $pk > pkfile, sk > skfile$ ]{ Ostart() :=
   $rk \stackrel{R}{\leftarrow} \text{keyseed}; pk \leftarrow \text{pkgen}(rk); sk \leftarrow \text{skgen}(rk);$ 
   $\text{return}(pk) \}$ ; (foreach  $i_1 \leq N_1$  do  $P_A$  | foreach  $i_2 \leq N_2$  do  $P_B$ )
 $P_A \stackrel{\text{def}}{=} \text{alice}\{ \text{OA}() :=$ 
   $m \stackrel{R}{\leftarrow} \text{nonce}; s \stackrel{R}{\leftarrow} \text{seed}; \text{return}(m, \text{sign}(m, sk, s))$ 
 $P_B \stackrel{\text{def}}{=} \text{bob}\{ \text{OB}(m' : \text{nonce}, s' : \text{signature}) :=$ 
  if  $\text{check}(m', pk, s')$  then  $\text{return}()$  else end

```

We divide this process into three parts. First, the key generation part is represented by the role keygen , which contains just the oracle Ostart . The annotation $pk > pkfile, sk > skfile$ means that we store the public key pk in the file $pkfile$ so that the oracle OB can access it, and analogously, we store the secret key sk in the file $skfile$ so that the oracle OA can access it. In other words, $\text{files} = \{(pk[], pkfile), (sk[], skfile)\}$.

The role `alice`, which contains the oracle OA, corresponds to the role of Alice and the role `bob`, which contains the oracle OB, corresponds to the role of Bob. For these two roles, there is no need to write the closing braces `}` because there is nothing after them.

Finally, the user annotations provide, for each CryptoVerif type T , the corresponding OCaml type $\mathbb{G}_T(T)$ as well as several OCaml functions:

- The function $\mathbb{G}_{\text{random}}(T) : \text{unit} \rightarrow \mathbb{G}_T(T)$ generates random numbers uniformly in T (when T is used in a random number generation).
- The serialization function $\mathbb{G}_{\text{ser}}(T) : \mathbb{G}_T(T) \rightarrow \text{string}$ converts an element of type $\mathbb{G}_T(T)$ to an OCaml string. The deserialization function $\mathbb{G}_{\text{deser}}(T) : \text{string} \rightarrow \mathbb{G}_T(T)$ performs the inverse operation. When deserialization fails, it must raise the exception `Bad_file`; this exception is raised only when a file has been corrupted. These functions are present when values of type T are written or read from tables and files.
- The predicate function $\mathbb{G}_{\text{pred}}(T) : \mathbb{G}_T(T) \rightarrow \text{bool}$ returns true if its argument corresponds to an element of type T and false otherwise (when T is present in the interface of the oracle definitions).

The user annotations also provide, for each CryptoVerif function $f : T_1 \times \dots \times T_m \rightarrow T$, a corresponding OCaml function $\mathbb{G}_f(f) : \mathbb{G}_T(T_1) \times \dots \times \mathbb{G}_T(T_m) \rightarrow \mathbb{G}_T(T)$. We assume that these functions are all provided in an OCaml module μ_{prim} .

CryptoVerif verifies the following properties:

Property 2.7 *There is a single occurrence of each role `role`. If an oracle O has a return (x_1, \dots, x_k) ; Q where Q contains a role definition, then there is only one return statement for the oracle O in the whole initial process.*

This property guarantees that we know which process to compile for a given role, and which roles start after the return from a given oracle.

Property 2.8 *There are no nested roles.*

Furthermore, for simplicity, we also assume the following points:

Assumption 2.9 *All oracle definitions are included in a role.*

Assumption 2.10 *No replication occurs above a parallel composition or a replication. When the definition of a role `role` is under replication `foreach $i \leq N$ do role{ Q , its contents Q consists of an oracle definition $O[\tilde{i}](\dots) := \dots$ or of a parallel composition of such oracle definitions (without replication).`*

A process can be transformed so that no replication occurs above a parallel composition by distributing the replications into the parallel compositions: `foreach $i \leq N$ do ($P_1 \mid P_2$)` can be transformed into `(foreach $i_1 \leq N_1$ do P_1) \mid (foreach $i_2 \leq N_2$ do P_2)`. We can encode nested replications by adding a dummy

oracle between the two replications: the process `foreach $i \leq N$ do foreach $j \leq N'$ do P` can be transformed into `foreach $i \leq N$ do $O() := \text{return}()$; foreach $j \leq N'$ do P` .

By Properties 2.3, 2.4, and 2.7, there cannot be, in the same process, a definition of an oracle O directly under replication and another definition of the same oracle O not directly under replication. Hence, we can use the phrase “ O is under replication” unambiguously. Moreover, by Property 2.4, the bound of the replication above a definition of an oracle O is the same for all definitions of O .

Assumption 2.11 *For each oracle O under replication, we let N_O be the bound of the replication above the definition of O . For each role `role` under replication, we let N_{role} be the bound of the replication above the definition of role. All these bounds N_O and N_{role} are pairwise distinct.*

This assumption allows us to be more precise when counting the number of times an oracle has been called, by using a distinct bound for each oracle.

These assumptions are relaxed in our implementation.

3 The OCaml Language

This section presents the OCaml language, the target language of our compiler, by giving its syntax and semantics. We omit some constructs, such as loops and type constructors, which are not used by our compiler. The subset that we consider is still Turing complete, so we do not lose expressivity by removing these constructs. To define the formal semantics, we adapted the semantics by Scott Owens et al. [15]. This semantics is a small step operational semantics of the core part of the OCaml language. We modified it in several ways, as detailed below.

3.1 Syntax and Informal Semantics

Figure 5 summarizes the syntax of our subset of OCaml. For brevity, we ignore types in this syntax.

Pattern-matching is a central feature of OCaml. A pattern pat describes the form of a value to be matched. When we match a value v with a pattern pat , if the value is of the correct form, then we bind each variable x occurring in the pattern pat to the corresponding part of v . Patterns must be linear, that is, no variable can occur more than once inside a pattern. When we match a value v with the pattern matching $pat_1 \rightarrow e_1 \mid \dots \mid pat_n \rightarrow e_n$, we match v sequentially to the patterns pat_1, \dots, pat_n . If the first pattern that matches v is pat_i , then we evaluate e_i . If no pattern matches v , then we raise the exception `Match_failure`.

The basic operations of the language are implemented by primitives *prim*. We write binary primitives in infix notation: for example, we write $v_1 = v_2$ rather than $(=) v_1 v_2$. We consider the following primitives: `not` is the boolean

$pat ::=$	pattern
x	variable
$\bar{}$	universal pattern
(pat_1, \dots, pat_n)	tuple
$pat_1 :: pat_2$	list constructor
$pm ::=$	pattern matching
$pat_1 \rightarrow e_1 \mid \dots \mid pat_n \rightarrow e_n$	pattern matching
$e ::=$	expression
$prim$	primitive
x	variable
l	location
c	constant (<code>[]</code> , <code>()</code> , <code>0</code> , <code>false</code> , ...)
(e_1, \dots, e_n)	tuple
$e_1 :: e_2$	list constructor
function pm	function
$e_1 e_2$	application
$e_1; e_2$	sequence
if e_1 then e_2 else e_3	if
match e with pm	pattern matching
try e with pm	try
let $pat = e_1$ in e_2	let
let rec $x_1 = \text{function } pm_1 \text{ and } \dots \text{ and } x_n = \text{function } pm_n \text{ in } e$	let rec
function $[env, pm]$	closure
letrec $[env, \{x_1 \mapsto \text{function } pm_1, \dots, x_n \mapsto \text{function } pm_n\} \text{ in } x_i]$	let rec closure, $1 \leq i \leq n$
addthread($program$)	addition of a thread
schedule(e)	schedule
$d ::=$	definition
let $pat = e$	let
let rec $x_1 = \text{function } pm_1 \text{ and } \dots \text{ and } x_n = \text{function } pm_n$	letrec
$definitions ::=$	definitions
ε	empty definition list
$d;; definitions$	definition list
$program ::=$	program
$definitions$	list of definitions
raise e	exception

Figure 5: OCaml syntax

$v ::=$	value
$prim\ v_1 \dots v_j$	partially applied primitives ($prim$ is n -ary and $0 \leq j < n$)
c	constant ($[], (), 0, false, \dots$)
l	location
(v_1, \dots, v_n)	tuple
$v_1 :: v_2$	list constructor
$function[env, pm]$	closure
$letrec[env, \{x_1 \mapsto function\ pm_1, \dots, x_n \mapsto function\ pm_n\} \text{ in } x_i]$	let rec closure

Figure 6: OCaml values

negation, $(=)$ is the equality test, $raise\ e$ raises the exception e . We use primitives to manage references, which are mutable memory cells. We represent memory cells by locations l ; we also use special locations to represent files. The reference creation $ref\ v$ creates a new location l , store the value v in l , and returns the location l . The assignment $l := v$ replaces the contents of the location l with the value v . The dereference $!l$ returns the contents of the location l . We also introduce a primitive for random number generation: $random\ ()$ returns a random boolean, true or false, with equal probability. This primitive was not present in [15]. It makes the semantics probabilistic. The language also includes primitives to manage other native types such as integers (e.g., addition and multiplication) and strings (e.g., concatenation, extraction of substrings, and conversion between integers in $\{0, \dots, 255\}$ and one-character strings). Strings are immutable values in our semantics. In contrast, in OCaml, values of type `string` are mutable. Our strings could be implemented in OCaml as an abstract type, on which only operations that do not mutate strings are implemented.

Most expressions are standard. Constants c can be integers, strings, boolean values true or false, the empty list $[]$, the unit constant $()$, and exceptions. The expression $function\ pm$ defines a function. When this function is applied to a value v , it matches that value using the pattern matching pm . The application $e_1\ e_2$ applies the function e_1 to the argument e_2 . The sequence operation $e_1; e_2$ evaluates e_1 , ignoring its result (but obviously keeping its side effects), then evaluates e_2 . The matching operation $match\ e\ \text{with}\ pm$ evaluates e and matches the result of e using the pattern matching pm . The try construct $try\ e\ \text{with}\ pm$ returns the result of e if e does not raise exceptions; if e raises an exception v matched by a pattern in pm , it returns the result of $match\ v\ \text{with}\ pm$; if e raises an exception v that is not matched by a pattern in pm , it also raises the exception v . The let binding $let\ pat = e_1\ \text{in}\ e_2$ evaluates e_1 , matches the result with the pattern pat , which binds the variables in pat , and finally evaluates e_2 . When the pattern matching fails, it raises the exception `Match_failure`. This construct is equivalent to $match\ e_1\ \text{with}\ pat \rightarrow e_2$. The let rec binding $let\ rec\ x_1 = function\ pm_1\ \text{and}\ \dots\ \text{and}\ x_n = function\ pm_n\ \text{in}\ e$ defines n mutually recursive

functions x_1, \dots, x_n , and evaluates the expression e using these functions.

Closures are not present in the initial program, but they serve to represent functional values internally. The closure `function[env, pm]` comes from the function `function pm`. It contains the code of the function (pm), and an environment env that maps the free variables of pm to their values. Closures allow one to evaluate functions using the values that the free variables of the function had at the definition of the function. (In other words, OCaml uses static variable binding.) The `let rec` closure `letrec[env, {x1 ↦ function pm1, ..., xn ↦ function pmn} in xi]` is similar, but for mutually recursive functions. It records several mutually recursive bindings together.

A security protocol typically involves several programs running in parallel on different machines. We model this situation by considering several threads. To manage threads, we introduce two new expressions, `addthread(program)` and `schedule(e)`. The expression `addthread(program)` creates a new thread that runs the program $program$. The expression `schedule(e)` stops execution of the current thread and continues execution of the thread number e . (Threads are designated by integer numbers. The initial thread, started at the beginning of the program, has number 1. The threads created by subsequent calls to `addthread` have numbers starting at 2 and increasing by one each time a new thread is created.)

We define the list expression $[e_1; e_2; \dots; e_n]$ as syntactic sugar for $e_1 :: (e_2 :: \dots :: (e_n :: [])) \dots$. The expression $e \&\& e'$ is syntactic sugar for `if e then e' else false`, and $e || e'$ is syntactic sugar for `if e then true else e'`.

A program is a list of top level definitions d , or the raising of an exception. We omit the final ε in a sequence of definitions when it is not empty.

Expressions reduce into values or exceptional values. As summarized in Figure 6, the values v are either functional values like closures, or tuples and lists of constants c or locations l . An exceptional value is `raise v`, where v is an exception value (a constant).

3.2 Formal Semantics

We define step by step the semantics of the various constructs of the language.

3.2.1 Pattern matching

We define the predicate `matches` in Figure 7: we have $v \text{ matches } pat \triangleright env$ when the value v matches the pattern pat , and the environment env is a mapping from the variables of pat to their values, computed by the pattern matching. The operation $env \oplus env' \stackrel{\text{def}}{=} env \upharpoonright_{\text{Dom}(env')} \cup env'$ adds the bindings of env' to those of env ; when a variable is bound in both environments, the binding of env' is kept. Since patterns are linear, in Figure 7, the operation $env \oplus env'$ is always used with environments env and env' that have disjoint domains; the general case is used below. We also define $v \text{ matches } pat$ as $\exists env, v \text{ matches } pat \triangleright env$.

$$\begin{array}{r}
v \text{ matches } x \triangleright \{x \mapsto v\} \quad \text{(Variable)} \\
v \text{ matches } _ \triangleright \emptyset \quad \text{(Any)} \\
\frac{\forall 1 \leq i \leq n, v_i \text{ matches } pat_i \triangleright env_i}{(v_1, \dots, v_n) \text{ matches } (pat_1, \dots, pat_n) \triangleright \bigoplus_{i=1}^n env_i} \quad \text{(Tuple)} \\
\frac{v_1 \text{ matches } pat_1 \triangleright env_1 \quad v_2 \text{ matches } pat_2 \triangleright env_2}{v_1 :: v_2 \text{ matches } pat_1 :: pat_2 \triangleright env_1 \oplus env_2} \quad \text{(List)}
\end{array}$$

Figure 7: Matches predicate

3.2.2 Primitives

The semantics of primitives is defined in Figure 8. This semantics is defined by rules of the form $prim\ v_1 \dots v_n \xrightarrow{L}_p e$ where $prim$ is an n -ary primitive. Such a rule means that $prim\ v_1 \dots v_n$ reduces to e with probability p . In contrast to [15], the semantics is probabilistic, because of the presence of the primitive `random`. The probability p is omitted when it is 1. The label L is used to reflect the operations on locations. It is empty when locations are unaffected. The label $ref\ v = l$ means that a new location l is created, with contents v . The label $!l = v$ means that the current contents of location l is v . The label $l := v$ means that the contents of the location l is changed into v . The rules are straightforward; they reflect the semantics defined informally in Section 3.1. One is not allowed to test equality between functional values, so we use the predicate `funval`, also defined in Figure 8, to test whether a value is functional, and raise the exception `Invalid_argument` when we try to test equality between functional values. There is no rule for the primitive `raise`: `raise v` is an exceptional value, it does not reduce.

3.2.3 Expressions and Programs

The semantics of [15] substitutes variables with their values. Instead, we define an environment env that maps variables to their values. This way, it is easier to relate the OCaml state to the CryptoVerif state which also contains an environment. Because of this change, we also need to add an explicit call stack $stack$. The stack is a list of pairs (env, C) , where C is a minimal *evaluation context*, that is, an expression with a hole $[\cdot]$, such that the expression inside the hole can be immediately evaluated. We define a minimal evaluation context as:

$C_m ::=$	minimal expression evaluation context
$e\ [\cdot]$	apply
$[\cdot]\ v$	apply function
$\text{let } pat = [\cdot] \text{ in } e$	let
$[\cdot]; e$	sequence
$\text{if } [\cdot] \text{ then } e_1 \text{ else } e_2$	if

Functional values:

$$\frac{\text{prim is an } n\text{-ary primitive} \quad 0 \leq j < n}{\text{funval}(\text{prim } v_1 \dots v_j)} \quad (\text{Primitive})$$

$$\text{funval}(\text{function}[env, pm]) \quad (\text{Function})$$

$$\text{funval}(\text{letrec}[env, \{x_1 \mapsto \text{function } pm_1, \dots, x_n \mapsto \text{function } pm_n\} \text{ in } x_i]) \quad (\text{Let rec})$$

Primitives:

$$\text{not} \left\{ \begin{array}{ll} \text{not false} \rightarrow \text{true} & (\text{Not1}) \\ \text{not true} \rightarrow \text{false} & (\text{Not2}) \end{array} \right.$$

$$\text{(=)} \left\{ \begin{array}{ll} \frac{\text{funval}(v) \text{ or } \text{funval}(v')}{v = v' \rightarrow \text{raise Invalid_argument}} & (\text{Funval}) \\ c = c \rightarrow \text{true} & (\text{Constant1}) \\ \frac{c \neq c'}{c = c' \rightarrow \text{false}} & (\text{Constant2}) \\ v_1 :: v_2 = v'_1 :: v'_2 \rightarrow v_1 = v'_1 \ \&\& \ v_2 = v'_2 & (\text{List1}) \\ v_1 :: v_2 = [] \rightarrow \text{false} & (\text{List2}) \\ [] = v'_1 :: v'_2 \rightarrow \text{false} & (\text{List3}) \\ (v_1, \dots, v_n) = (v'_1, \dots, v'_n) \rightarrow v_1 = v'_1 \ \&\& \ \dots \ \&\& \ v_n = v'_n & (\text{Tuples}) \end{array} \right.$$

$$\text{ref} \left\{ \begin{array}{ll} \text{ref } v \xrightarrow{\text{ref } v=l} l & (\text{New ref}) \end{array} \right.$$

$$\text{(:=)} \left\{ \begin{array}{ll} l := v \xrightarrow{l:=v} () & (\text{Assign}) \end{array} \right.$$

$$! \left\{ \begin{array}{ll} !l \xrightarrow{!l=v} v & (\text{Dereference}) \end{array} \right.$$

$$\text{random} \left\{ \begin{array}{ll} \frac{a \in \{\text{true}, \text{false}\}}{\text{random } () \rightarrow_{1/2} a} & (\text{Random}) \end{array} \right.$$

Figure 8: Rules for OCaml primitives

$\text{match } [\cdot] \text{ with } pm$	match
$\text{try } [\cdot] \text{ with } pm$	try
$(e_1, \dots, e_{i-1}, [\cdot], v_{i+1}, \dots, v_n)$	tuple
$e :: [\cdot]$	cons1
$[\cdot] :: v$	cons2
$\text{schedule}([\cdot])$	schedule
$C_{pm} ::=$	$\text{minimal program evaluation context}$
$\text{let } pat = [\cdot];; definitions$	let

For instance, $e [\cdot]$ and $[\cdot] v$ are evaluation contexts, so we evaluate the argument of applications first, and when it becomes a value, we evaluate the function. We denote by $C_m[e]$ the context C_m with the hole $[\cdot]$ replaced by e , and similarly for C_{pm} . The stack contains a minimal program evaluation context C_{pm} in the last element of the list and expression evaluation contexts C_m in the other elements if it is non-empty.

Hence, we evaluate expressions and programs by reducing triples $env, pe, stack$, where pe means program *program* or expression *e*. The reduction rules $env, pe, stack \xrightarrow{L}_p env', pe', stack'$ are defined in Figures 9 and 10 for expressions and Figure 11 for programs. The label L is defined above in Section 3.2.2. These reductions are probabilistic; the probability p is omitted when it is 1. Most rules are straightforward. In order to evaluate an expression $C_m[e]$, we need to reduce e under the context C_m . To do that, we push the context C_m on the stack with the current environment by rule (Context in), evaluate the expression e until it becomes a value v , and finally pop the context C_m from the stack by rule (Context out), inserting the obtained value v in C_m , yielding $C_m[v]$. In case the expression e raises an exception v , we use rules (Context raise1) and (Context raise2). If the context C_m is not a try context, the result of $C_m[e]$ is also $\text{raise } v$ by (Context raise1). If C_m is a try context, we evaluate that try by (Context raise2), followed by (Try2). The rules (Let ctx in), (Let ctx out), and (Let ctx raise) play the same role as (Context in), (Context out), and (Context raise1) respectively, for programs instead of expressions: they allow reducing under the minimal program evaluation context $\text{let } pat = [\cdot];; definitions$. There is no rule corresponding to (Context raise2) for programs because there is no try program context.

Example 3 Let us present as an example the reduction of a simple program in an empty environment and an empty stack:

$$\emptyset, \text{let } x = \text{if random } () \text{ then } 0 \text{ else } 1;;, [] .$$

We first reduce the expression part of the let, by keeping in the stack the fact that the expression is under the context $\text{let } x = [\cdot]$. This expression reduces eventually to a value, and at this point we plug this value back into the context. So we first reduce the previous configuration by (Let ctx in) into:

$$\emptyset, \text{if random } () \text{ then } 0 \text{ else } 1, [\emptyset, \text{let } x = [\cdot];;]$$

By (Context in), we prepare to reduce the condition of the if:

$$\emptyset, \text{random } (), [\emptyset, \text{if } [\cdot] \text{ then } 0 \text{ else } 1; \emptyset, \text{let } x = [\cdot];;]$$

By (Random), `random ()` reduces to true with probability 1/2 and false with probability 1/2. For the purpose of the example, let us consider the case where `random ()` reduces to true. By (Primitives), the configuration reduces with probability 1/2 into

$$\emptyset, \text{true}, [\emptyset, \text{if } [\cdot] \text{ then } 0 \text{ else } 1; \emptyset, \text{let } x = [\cdot];;]$$

By (Context out), we insert the value of the condition back into the if:

$$\emptyset, \text{if true then } 0 \text{ else } 1, [\text{let } x = [\cdot];;]$$

By (If1), we evaluate the if:

$$\emptyset, 0, [\text{let } x = [\cdot];;]$$

By (Let ctx out), we insert the obtained value back into the context `let x = [];`:

$$\emptyset, \text{let } x = 0;;, []$$

By (Variable), we have that 0 matches $x \triangleright \{x \mapsto 0\}$. So, by (Let match1), the configuration reduces into the following last configuration:

$$\{x \mapsto 0\}, \varepsilon, []$$

The expressions `addthread(program)` and `schedule(e)` are treated specially because they alter parts of the semantic configuration other than *env*, *pe*, *stack*. Their treatment is detailed in Section 3.2.5.

3.2.4 Store

As usual, the contents of locations are stored in a *store*, which maps locations to their current values. Figure 12 defines the relation $store \xrightarrow{L} store'$. If a program or an expression reduces by $env, pe, stack \xrightarrow{L}_p env', pe', stack'$, then the store *store* will be updated into *store'* such that $store \xrightarrow{L} store'$. When *L* is empty, the store is unchanged by rule (Store empty). When *L* is $!l = v$, the store is also unchanged, but the reduction succeeds only when the contents of *l* is *v*, by rule (Store lookup). When *L* is $l := v$, the store is updated so that *l* contains *v*, by rule (Store assign). When *L* is $\text{ref } v = l$, a new location *l* is created with contents *v*, so the contents of *l* must not be defined in the initial store, by rule (Store alloc).

$$\begin{array}{c}
\text{env}, x, \text{stack} \rightarrow \text{env}, \text{env}(x), \text{stack} \quad (\text{Env}) \\
\\
\frac{\text{prim } v_1 \dots v_n \xrightarrow{L}_p e}{\text{env}, \text{prim } v_1 \dots v_n, \text{stack} \xrightarrow{L}_p \text{env}, e, \text{stack}} \quad (\text{Primitives}) \\
\\
\text{e is not a value} \\
\text{and, when } C_m \text{ is a try context, } e \text{ is not an exceptional value} \\
\frac{}{\text{env}, C_m[e], \text{stack} \rightarrow \text{env}, e, (\text{env}, C_m) :: \text{stack}} \quad (\text{Context in}) \\
\\
\text{env}', v, (\text{env}, C_m) :: \text{stack} \rightarrow \text{env}, C_m[v], \text{stack} \quad (\text{Context out}) \\
\\
\frac{C_m \text{ is not a try context}}{\text{env}', \text{raise } v, (\text{env}, C_m) :: \text{stack} \rightarrow \text{env}, \text{raise } v, \text{stack}} \quad (\text{Context raise1}) \\
\\
\frac{C_m \text{ is a try context}}{\text{env}', \text{raise } v, (\text{env}, C_m) :: \text{stack} \rightarrow \text{env}, C_m[\text{raise } v], \text{stack}} \quad (\text{Context raise2}) \\
\\
\text{env}, \text{function } pm, \text{stack} \rightarrow \text{env}, \text{function}[\text{env}, pm], \text{stack} \quad (\text{Closure}) \\
\\
\text{env}, \text{function}[\text{env}', pm] v_0, \text{stack} \rightarrow \text{env}', \text{match } v_0 \text{ with } pm, \text{stack} \quad (\text{Expr apply}) \\
\\
\text{env}, v; e, \text{stack} \rightarrow \text{env}, e, \text{stack} \quad (\text{Sequence}) \\
\\
\text{env}, \text{if true then } e_1 \text{ else } e_2, \text{stack} \rightarrow \text{env}, e_1, \text{stack} \quad (\text{If1}) \\
\\
\text{env}, \text{if false then } e_1 \text{ else } e_2, \text{stack} \rightarrow \text{env}, e_2, \text{stack} \quad (\text{If2}) \\
\\
\frac{v \text{ matches } pat \triangleright \text{env}'}{\text{env}, \text{match } v \text{ with } pat \rightarrow e \mid pat_1 \rightarrow e_1 \mid \dots \mid pat_n \rightarrow e_n, \text{stack} \rightarrow \text{env} \oplus \text{env}', e, \text{stack}} \quad (\text{Match1}) \\
\\
\frac{\neg(v \text{ matches } pat)}{\text{env}, \text{match } v \text{ with } pat \rightarrow e \mid pat_1 \rightarrow e_1 \mid \dots \mid pat_n \rightarrow e_n, \text{stack} \rightarrow \text{env}, \text{match } v \text{ with } pat_1 \rightarrow e_1 \mid \dots \mid pat_n \rightarrow e_n, \text{stack}} \quad (\text{Match2}) \\
\\
\frac{\neg(v \text{ matches } pat)}{\text{env}, \text{match } v \text{ with } pat \rightarrow e, \text{stack} \rightarrow \text{env}, \text{raise Match_failure}, \text{stack}} \quad (\text{Match fail}) \\
\\
\text{env}, \text{try } v \text{ with } pm, \text{stack} \rightarrow \text{env}, v, \text{stack} \quad (\text{Try1}) \\
\\
\text{env}, \text{try raise } v \text{ with } pm, \text{stack} \rightarrow \text{env}, \text{match } v \text{ with } pm \mid _ \rightarrow \text{raise } v, \text{stack} \quad (\text{Try2})
\end{array}$$

Figure 9: Rules for expressions

$$\begin{array}{c}
\frac{v \text{ matches } pat \triangleright env'}{env, \text{ let } pat = v \text{ in } e, stack \rightarrow env \oplus env', e, stack} \quad (\text{Let1}) \\
\frac{\neg(v \text{ matches } pat)}{env, \text{ let } pat = v \text{ in } e, stack \rightarrow env, \text{ raise Match_failure}, stack} \quad (\text{Let2}) \\
\frac{letrecenv = \{x_1 \mapsto \text{function } pm_1, \dots, x_n \mapsto \text{function } pm_n\}}{env, \text{ let rec } x_1 = \text{function } pm_1 \text{ and } \dots \text{ and } x_n = \text{function } pm_n \text{ in } e, stack \\ \rightarrow env[x_1 \mapsto \text{letrec}[env, letrecenv \text{ in } x_1], \\ \dots, \\ x_n \mapsto \text{letrec}[env, letrecenv \text{ in } x_n]], e, stack} \quad (\text{Closure let rec}) \\
\frac{letrecenv = \{x_1 \mapsto \text{function } pm_1, \dots, x_n \mapsto \text{function } pm_n\}}{env, \text{ letrec}[env', letrecenv \text{ in } x_i] v_0, stack \rightarrow \\ env'[x_1 \mapsto \text{letrec}[env', letrecenv \text{ in } x_1], \\ \dots, \\ x_n \mapsto \text{letrec}[env', letrecenv \text{ in } x_n]], \\ \text{match } v_0 \text{ with } pm_i, stack} \quad (\text{Expr letrec apply})
\end{array}$$

Figure 10: Rules for expressions (continued)

$$\begin{array}{c}
\frac{e \text{ is not a value}}{env, \text{ let } pat = e;; definitions, [] \rightarrow env, e, [env, \text{ let } pat = [::]; definitions]} \quad (\text{Let ctx in}) \\
env', v, [env, \text{ let } pat = [::]; definitions] \rightarrow env, \text{ let } pat = v;; definitions, [] \quad (\text{Let ctx out}) \\
env', \text{ raise } v, [env, \text{ let } pat = [::]; definitions] \rightarrow env, \text{ raise } v, [] \quad (\text{Let ctx raise}) \\
\frac{v \text{ matches } pat \triangleright env'}{env, \text{ let } pat = v;; definitions, [] \rightarrow env \oplus env', definitions, []} \quad (\text{Let match1}) \\
\frac{\neg(v \text{ matches } pat)}{env, \text{ let } pat = v;; definitions, [] \rightarrow env, \text{ raise Match_failure}, []} \quad (\text{Let match2}) \\
\frac{letrecenv = \{x_1 \mapsto \text{function } pm_1, \dots, x_n \mapsto \text{function } pm_n\}}{env, \text{ let rec } x_1 = \text{function } pm_1 \text{ and } \dots \text{ and } x_n = \text{function } pm_n;; definitions, [] \\ \rightarrow env[x_1 \mapsto \text{letrec}[env, letrecenv \text{ in } x_1], \\ \dots, \\ x_n \mapsto \text{letrec}[env, letrecenv \text{ in } x_n]], definitions, []} \quad (\text{Closure let rec})
\end{array}$$

Figure 11: Rules for programs

$store \rightarrow store$	(Store empty)
$\frac{store(l) = v}{store \xrightarrow{!l=v} store}$	(Store lookup)
$\frac{l \in \text{Dom}(store)}{store \xrightarrow{l:=v} store[l \mapsto v]}$	(Store assign)
$\frac{l \notin \text{Dom}(store)}{store \xrightarrow{\text{ref } v=l} store[l \mapsto v]}$	(Store alloc)

Figure 12: Store rules

3.2.5 Toplevel Reduction

As mentioned in Section 3.1, and in contrast to [15], we consider several threads running in parallel. Each thread has a configuration $Th_i = \langle env_i, pe_i, stack_i, store_i \rangle$ that contains the current $env_i, pe_i, stack_i$ as explained in Section 3.2.3, as well as the contents of locations local to this thread, in a store $store_i$, as explained in Section 3.2.4. The complete semantic configuration is then

$$\mathcal{C} = [Th_1, \dots, Th_n], globalstore, tj$$

where tj is the number of the thread currently being executed, and $globalstore$ is a store for locations shared between threads. We use it to model the communication between threads by storing messages in global locations, and to store the files containing private data from the CryptoVerif process (free variables of roles and tables). In practice, these files may be copied from one machine to another by the user, so they are actually shared between several threads. The values in the global store contain no closure and no reference. (In OCaml, closures and references can be written to a file only by marshalling, but marshalling is ruled out by Assumption A4, since it may bypass the type system.) The global store contains locations in a set S_g , while the local stores contain locations in an infinite set S_l , with $S_g \cap S_l = \emptyset$.

The reduction rules for semantic configurations \mathcal{C} are defined in Figure 13. Actually, this figure defines three relations. The relation $Th \rightarrow_p Th'$, defined by rule (Thread), handles all operations that deal with the current thread only. It updates the store using the same label L as the one used for evaluating the program or the expression, and it checks that this label concerns the local store of the thread. (The location l , if any, must be in S_l .)

Second, the relation $Th, globalstore \rightarrow_p Th', globalstore'$, defined by rules (Globalstore1) and (Globalstore2), handles all operations local to one thread and the global store operations. By rule (Globalstore1), it uses the relation $Th \rightarrow_p Th'$ to handle the operations local to one thread. By rule (Globalstore2), it handles the global store operations. It updates the global store using the

same label L as the one used for evaluating the program or the expression, and it checks that this label concerns the global store. The location l must be in S_g , and the creation of a location in the global store is forbidden. (Otherwise, one would need a way to tell the system whether a new location should be created in the local or in the global store, and to communicate the global locations to the other threads.) We assume that all locations of the global store are initialized at the beginning of the program.

Finally, the relation $\mathcal{C} \rightarrow_p \mathcal{C}'$, defined by the last four rules of Figure 13, gives the semantics of the full language. Rule (Toplevel) runs the current thread t_j , using the relation $Th, globalstore \rightarrow_p Th', globalstore'$. Rule (Toplevel add thread) defines the semantics of `addthread(program)`: it creates a new thread that runs the program *program*, with empty environment, stack, and store. Rules (Toplevel schedule1) and (Toplevel schedule2) define the semantics of `schedule`: `schedule(tj')` schedules thread number t_j' when this thread exists, and otherwise it raises the exception `Invalid_argument`.

Splitting the definition of the semantics into three relations allows us to lighten notations in proofs: we can use the reduction on a thread, or on a thread and the global store, without mentioning the other components when they are not affected.

The construct `addthread` does not allow using the same local store in several threads, which corresponds to forbidding fork in the middle of a role. Moreover, we reduce only the active thread, and we change threads only with `schedule`. So we can only change threads in code defined by the adversary, because neither the primitives nor the generated modules use `schedule`. So a call to an oracle cannot be interleaved with other threads. This property corresponds to Assumption A6: if several oracles cannot interleave reads and writes in the same table file, one can reconstruct a well-defined call order for these oracles in the `CryptoVerif` process, which processes one oracle call after another, so that the calls can be simulated in our semantics.

3.2.6 Modules

OCaml programs typically contain several modules. We adopt a very simple model of modules. A module named μ consists of an OCaml program $program(\mu)$ and its interface $interface(\mu)$ that is the set of OCaml identifiers defined in μ and usable in other modules. When needed to distinguish identifiers coming from different modules, we use identifiers of the form $\mu.x$ for variables defined in module μ . A correct OCaml program is then of the form $program = program(\mu_1); \dots; program(\mu_n);$, where, for all $i \leq n$, the free variables of μ_i are defined in the interfaces of μ_j with $j < i$.

Such a program is run by using the previous reduction rules from the initial configuration

$$\mathcal{C}_0(program) = [\langle \emptyset, program, [], \emptyset \rangle, globalstore_0, 1]$$

where $globalstore_0 = \{l \mapsto initval_l \mid l \in S_g\}$ is the initial value of the global store, and $initval_l$ is the default value for location l : the empty list `[]` for lists,

$$\begin{array}{c}
\text{store} \xrightarrow{L} \text{store}' \\
L \text{ is empty or } L \text{ is } !l = v, l := v, \text{ or } \text{ref } v = l \text{ with } l \in S_l \\
\text{env}, \text{pe}, \text{stack} \xrightarrow{L}_p \text{env}', \text{pe}', \text{stack}' \\
\hline
\langle \text{env}, \text{pe}, \text{stack}, \text{store} \rangle \rightarrow_p \langle \text{env}', \text{pe}', \text{stack}', \text{store}' \rangle \quad (\text{Thread})
\end{array}$$

$$\begin{array}{c}
\text{Th} \rightarrow_p \text{Th}' \\
\hline
\text{Th}, \text{globalstore} \rightarrow_p \text{Th}', \text{globalstore} \quad (\text{Globalstore1})
\end{array}$$

$$\begin{array}{c}
\text{globalstore} \xrightarrow{L} \text{globalstore}' \\
L \text{ is } !l = v \text{ or } l := v \text{ with } l \in S_g \\
\text{env}, \text{pe}, \text{stack} \xrightarrow{L}_p \text{env}', \text{pe}', \text{stack}' \\
\hline
\langle \text{env}, \text{pe}, \text{stack}, \text{store} \rangle, \text{globalstore} \rightarrow_p \langle \text{env}', \text{pe}', \text{stack}', \text{store}' \rangle, \text{globalstore}' \quad (\text{Globalstore2})
\end{array}$$

$$\begin{array}{c}
\text{Th}, \text{globalstore} \rightarrow_p \text{Th}', \text{globalstore}' \\
\hline
[\text{Th}_1, \dots, \text{Th}_{t_j-1}, \text{Th}, \text{Th}_{t_j+1}, \dots, \text{Th}_n], \text{globalstore}, t_j \rightarrow_p \\
[\text{Th}_1, \dots, \text{Th}_{t_j-1}, \text{Th}', \text{Th}_{t_j+1}, \dots, \text{Th}_n], \text{globalstore}', t_j \quad (\text{Toplevel})
\end{array}$$

$$\begin{array}{c}
[\text{Th}_1, \dots, \text{Th}_{t_j-1}, \langle \text{env}, \text{addthread}(\text{program}), \text{stack}, \text{store} \rangle, \text{Th}_{t_j+1}, \dots, \text{Th}_n], \\
\text{globalstore}, t_j \rightarrow \\
[\text{Th}_1, \dots, \text{Th}_{t_j-1}, \langle \text{env}, () \rangle, \text{stack}, \text{store}], \text{Th}_{t_j+1}, \dots, \text{Th}_n, \langle \emptyset, \text{program}, [], \emptyset \rangle, \\
\text{globalstore}, t_j \quad (\text{Toplevel add thread})
\end{array}$$

$$\begin{array}{c}
1 \leq t_j' \leq n \\
\hline
[\text{Th}_1, \dots, \text{Th}_{t_j-1}, \langle \text{env}, \text{schedule}(t_j'), \text{stack}, \text{store} \rangle, \text{Th}_{t_j+1}, \dots, \text{Th}_n], \\
\text{globalstore}, t_j \rightarrow \\
[\text{Th}_1, \dots, \text{Th}_{t_j-1}, \langle \text{env}, () \rangle, \text{stack}, \text{store}], \text{Th}_{t_j+1}, \dots, \text{Th}_n, \text{globalstore}, t_j' \quad (\text{Toplevel schedule1})
\end{array}$$

$$\begin{array}{c}
t_j' < 1 \text{ or } t_j' > n \\
\hline
[\text{Th}_1, \dots, \text{Th}_{t_j-1}, \langle \text{env}, \text{schedule}(t_j'), \text{stack}, \text{store} \rangle, \text{Th}_{t_j+1}, \dots, \text{Th}_n], \\
\text{globalstore}, t_j \rightarrow \\
[\text{Th}_1, \dots, \text{Th}_{t_j-1}, \langle \text{env}, \text{raise Invalid_argument}, \text{stack}, \text{store} \rangle, \text{Th}_{t_j+1}, \dots, \text{Th}_n], \\
\text{globalstore}, t_j \quad (\text{Toplevel schedule2})
\end{array}$$

Figure 13: Top level rules

the empty string "" for strings, 0 for integers, false for booleans. (Each location implicitly has a type. Values in the global store cannot contain locations and closures, so we do not define a default value for them.) The program *program* does not contain closures nor locations in S_l , but may contain locations in S_g . (Closures are created by `function` and `letrec`; locations in S_l are created by `ref`.)

Although we ignore types in our syntax, we suppose that our OCaml programs are well-typed, which is checked by the OCaml compiler, and we use the guarantee that well-typed programs do not go wrong: a program stops only when the current thread has been reduced into the empty definition list or an exception `raise v` (with the empty stack).

3.2.7 Equivalence Modulo Renaming of Locations

The rule (Store alloc) is non-deterministic, since the new location l can be any unused location in S_l . To remove this non-determinism, we consider equivalence classes of OCaml semantic configurations modulo renaming of locations in S_l . We still denote these equivalence classes as OCaml configurations \mathcal{C} , and denote an equivalence class by one of its members. On these equivalence classes, the semantics is purely probabilistic. (There is no non-deterministic choice.) If a configuration \mathcal{C} can reduce, then the sum of the probabilities of all possible reductions is 1:

$$\sum_{\{\mathcal{C}' | \mathcal{C} \rightarrow_{p(\mathcal{C}')} \mathcal{C}'\}} p(\mathcal{C}') = 1$$

Moreover, for each reduction $\mathcal{C} \rightarrow_p \mathcal{C}'$, we have $p > 0$.

We will also use notations similar to Definition 2.6 for the OCaml semantics. We denote by \mathcal{CT} an OCaml trace, \mathcal{CTS} a set of OCaml traces, and we also use the notation \rightarrow^* for reductions with several steps.

4 Instrumentation of the OCaml Semantics

In order to prove the correctness of our compiler, we instrument OCaml code in three ways; this section details this instrumentation and proves that it does not alter the semantics of OCaml.

First, we add a new kind of functions and closures `tagfunction` that behave exactly in the same way as regular functions and closures, but are labeled with additional tags. We use these tagged functions to differentiate functions coming from our generated code and functions coming from the adversary. Hence, we add two new expressions `tagfunctiont pm` for tagged functions and `tagfunctiont,τ[env, pm]` for the corresponding closures. We also add `tagfunctiont,τ[env, pm]` to the values. The tag t indicates the origin of the function or closure; it will be an oracle name or a role name, indicating that the function implements this oracle or role. The tag τ is a fresh tag generated when the function is reduced into a closure: each new closure gets a different tag, so that two closures are the same if and only if they have the same tag. This property will be used in Section 6 to count the number of calls to the same

$$\begin{array}{c}
\text{funval}(\text{tagfunction}^{t,\tau}[\text{env}, pm]) \qquad \text{(Tagged funval)} \\
\hline
\tau \text{ fresh} \\
\hline
\text{env}, \text{tagfunction}^t pm, \text{stack} \rightarrow \text{env}, \text{tagfunction}^{t,\tau}[\text{env}, pm], \text{stack} \\
\text{(Tagged closure)} \\
\hline
\text{env}, \text{tagfunction}^{t,\tau}[\text{env}', pm] v_0, \text{stack} \rightarrow \text{env}', \text{match } v_0 \text{ with } pm, \text{stack} \\
\text{(Tagged expr apply)}
\end{array}$$

Figure 14: Semantics of tagged functions

closure. The semantic rules for tagged functions are given in Figure 14. They are the same as those for ordinary functions, except for the addition of tags. Much like for locations, we consider traces modulo renaming of tags τ , so that the choice of a fresh tag τ in (Tagged closure) does not lead to non-determinism. The condition that τ is fresh in this rule means that τ is distinct from all tags previously used in the considered trace.

Second, we need to be able to match CryptoVerif events, so we add to the semantic configuration an element *events* that contains the list of the events executed until now. We add the expression `event $ev(e_1, \dots, e_k)$` that adds the event $ev(v_1, \dots, v_k)$ to *events* when e_1, \dots, e_k evaluate to the values v_1, \dots, v_k respectively. We consider a new minimal expression evaluation context `event($ev(e_1, \dots, e_{i-1}, [], v_{i+1}, \dots, v_n)$)`, for evaluating the arguments of events. Events serve in specifying security properties of protocols, so they appear in generated code, but cannot be used by the adversary.

Third, the roles of a CryptoVerif process cannot be executed in any order: if a role is defined after the return from an oracle, it can be executed only after the previous oracle has returned. For instance, we can run a server only after generating its keys. We need to enforce this constraint also in the OCaml program. Each CryptoVerif role *role* is translated by our compiler into an OCaml module μ_{role} . We add to the OCaml configuration the multiset of callable modules \mathcal{MI} that contains pairs (μ_{role}, b) of a module μ_{role} and a boolean b , indicating, if true, that the module can be called any number of times and if false that the module can be called only once. Hence, the instrumented semantic configuration is

$$\mathcal{CI} = [Th_1, \dots, Th_n], \text{globalstore}, j, \mathcal{MI}, \text{events}$$

We adapt the toplevel semantic rules to this configuration as shown in Figure 15. The instrumented semantic rules (New toplevel), (New toplevel schedule1), and (New toplevel schedule2) are straightforwardly adapted from the corresponding rules in the non-instrumented semantics by adding the components $\mathcal{MI}, \text{events}$. The rule (Toplevel event) gives the semantics of `event`: it adds its argument $e(v_1, \dots, v_n)$ to the list *events* in the configuration and returns (v_1, \dots, v_n) . The rule (New toplevel add thread) gives the instrumented semantics of `addthread`: the `addthread` construct is modified to reject new programs

$$\begin{array}{c}
\frac{Th, globalstore \longrightarrow_p Th', globalstore'}{[Th_1, \dots, Th_{t_j-1}, Th, Th_{t_j+1}, \dots, Th_n], globalstore, t_j, \mathcal{MI}, events \longrightarrow_p [Th_1, \dots, Th_{t_j-1}, Th', Th_{t_j+1}, \dots, Th_n], globalstore', t_j, \mathcal{MI}, events} \\
\text{(New toplevel)} \\
\\
\text{program} = \text{program}(\mu_{\text{prim}});; \text{program}(\mu_1);; \dots;; \text{program}(\mu_l);; \text{program}' \\
\text{program}' \text{ does not contain } \text{program}(\mu_{\text{prim}}) \text{ nor any } \text{program}(\mu) \text{ for } \mu \in \mathcal{M}_g \\
\mathcal{M} = \{\mu_1, \dots, \mu_l\} \subseteq \mathcal{M}_g \\
\forall \mu \in \mathcal{M}, \exists b, (\mu, b) \in \mathcal{MI} \\
\mathcal{MI}' = \{(\mu, \text{false}) \mid \mu \in \mathcal{M} \wedge (\mu, \text{false}) \in \mathcal{MI}\} \\
\text{or} \\
\text{program} \text{ does not contain } \text{program}(\mu_{\text{prim}}) \text{ nor any } \text{program}(\mu) \text{ for } \mu \in \mathcal{M}_g \\
\mathcal{M} = \emptyset, \mathcal{MI}' = \emptyset \\
\\
\frac{[Th_1, \dots, Th_{t_j-1}, \langle env, \text{addthread}(\text{program}), stack, store \rangle, Th_{t_j+1}, \dots, Th_n], globalstore, t_j, \mathcal{MI}, events \longrightarrow [Th_1, \dots, Th_{t_j-1}, \langle env, () \rangle, stack, store], Th_{t_j+1}, \dots, Th_n, \langle \emptyset, \text{program}, [], \emptyset \rangle, globalstore, t_j, \mathcal{MI} \setminus \mathcal{MI}', events}{} \\
\text{(New toplevel add thread)} \\
\\
\frac{1 \leq t_j' \leq n}{[Th_1, \dots, Th_{t_j-1}, \langle env, \text{schedule}(t_j'), stack, store \rangle, Th_{t_j+1}, \dots, Th_n], globalstore, t_j, \mathcal{MI}, events \longrightarrow [Th_1, \dots, Th_{t_j-1}, \langle env, () \rangle, stack, store], Th_{t_j+1}, \dots, Th_n, globalstore, t_j', \mathcal{MI}, events} \\
\text{(New toplevel schedule1)} \\
\\
\frac{t_j' < 1 \text{ or } t_j' > n}{[Th_1, \dots, Th_{t_j-1}, \langle env, \text{schedule}(t_j'), stack, store \rangle, Th_{t_j+1}, \dots, Th_n], globalstore, t_j, \mathcal{MI}, events \longrightarrow [Th_1, \dots, Th_{t_j-1}, \langle env, \text{raise Invalid_argument}, stack, store \rangle, Th_{t_j+1}, \dots, Th_n], globalstore, t_j, \mathcal{MI}, events} \\
\text{(New toplevel schedule2)} \\
\\
\frac{[Th_1, \dots, Th_{t_j-1}, \langle env, \text{return}(\mathcal{MI}', v), stack, store \rangle, Th_{t_j+1}, \dots, Th_n], globalstore, t_j, \mathcal{MI}, events \longrightarrow [Th_1, \dots, Th_{t_j-1}, \langle env, v, stack, store \rangle, Th_{t_j+1}, \dots, Th_n], globalstore, t_j, \mathcal{MI} \cup \mathcal{MI}', events}{} \\
\text{(Toplevel return)} \\
\\
\frac{[Th_1, \dots, Th_{t_j-1}, \langle env, \text{event } ev(v_1, \dots, v_n), stack, store \rangle, Th_{t_j+1}, \dots, Th_n], globalstore, t_j, \mathcal{MI}, events \longrightarrow [Th_1, \dots, Th_{t_j-1}, \langle env, (v_1, \dots, v_n), stack, store \rangle, Th_{t_j+1}, \dots, Th_n], globalstore, t_j, \mathcal{MI}, e(v_1, \dots, v_n) :: events}{} \\
\text{(Toplevel event)}
\end{array}$$

Figure 15: Updated toplevel rules for the instrumented semantics

that contain a module that cannot be called. We let \mathcal{M}_g be the set of generated modules. The programs spawned by `addthread` can be of two forms. Either they are *attacker programs* that contain neither the module corresponding to the primitives μ_{prim} nor any generated module in \mathcal{M}_g , or they are *protocol programs* that first contain the module corresponding to the primitives μ_{prim} , then the necessary generated modules μ_1, \dots, μ_l in \mathcal{M}_g , and finally any non-generated program *program'*. (We require this order on the modules for simplicity.) The generated modules μ_1, \dots, μ_l must be callable according to the value of \mathcal{MI} . The modules μ_1, \dots, μ_l that can be called only once are removed from the callable modules by removing the multiset \mathcal{MI}' from \mathcal{MI} .

We also add the expression `return(\mathcal{MI}' , e)` that adds to the multiset \mathcal{MI} the generated modules present in \mathcal{MI}' , and returns the result of e , as defined by rule (Toplevel return). This expression is useful to add modules newly defined at the return from an oracle. We also add the minimal expression evaluation context `return(\mathcal{MI} , $[\cdot]$)` to be able to evaluate the second argument of `return`.

Let us now show that this instrumentation does not alter the semantics of OCaml: an instrumented program behaves exactly in the same way as that program with the instrumentation deleted, provided only allowed roles are executed, as assumed by Assumption A2. This assumption is formalized as follows:

Assumption 4.1 (Only allowed roles) *The instrumented addthread rule (New toplevel add thread) never fails.*

We first show that, when a program or expression is a value v or an exceptional value `raise v` , the environment does not matter. To prove this property, we define the following equivalence.

Definition 4.2 *We define the equivalence \approx_{vTh} on threads by*

$$\langle env, pe, stack, store \rangle \approx_{vTh} \langle env', pe', stack', store' \rangle$$

if and only if $pe, stack, store = pe', stack', store'$, and if pe is not a value v or an exceptional value `raise v` , then $env = env'$.

We extend this equivalence to non instrumented configurations \mathcal{C} and \mathcal{C}' by $\mathcal{C} \approx_v \mathcal{C}'$ if and only if

- $\mathcal{C} = [Th_1, \dots, Th_n], globalstore, tj$,
- $\mathcal{C}' = [Th'_1, \dots, Th'_n], globalstore, tj$,
- $\forall tj' \leq n, Th_{tj'} \approx_{vTh} Th'_{tj'}$.

We show that configurations equivalent by \approx_v reduce in the same way.

Lemma 4.3 *If $\mathcal{C} \approx_v \mathcal{C}'$ and $\mathcal{C} \rightarrow_p \mathcal{C}''$, then $\mathcal{C}' \rightarrow_p \mathcal{C}'''$ and $\mathcal{C}'' \approx_v \mathcal{C}'''$.*

Proof No semantic rule uses the environment when the program or expression is a value or an exceptional value.

Indeed, the only semantic rules that apply when the program or expression is a value or an exceptional value are (Context out), (Context raise1), (Context raise2), (Let ctx out), and (Let ctx raise). All these rules replace the current environment with the one stored at the top of the stack. \square

By reviewing the changes to the semantics, we can see that the total probability of all reductions is still 1 for the instrumented semantics: If an instrumented semantic configuration \mathcal{CI} can reduce, then

$$\sum_{\{\mathcal{CI}' \mid \mathcal{CI} \rightarrow_p(\mathcal{CI}') \mathcal{CI}'\}} p(\mathcal{CI}') = 1.$$

Moreover, for each reduction $\mathcal{CI} \rightarrow_p \mathcal{CI}'$, we have $p > 0$.

Let us now define the function $\text{noinstr}_{\mathcal{CI}}$ that takes a configuration in the instrumented semantics and returns the corresponding configuration in the non-instrumented semantics.

Definition 4.4 *The function noinstr_{Th_1} applied to a thread replaces*

1. every $\text{return}(\mathcal{MI}, e)$ with e ,
2. every event $ev(e_1, \dots, e_n)$ with (e_1, \dots, e_n) ,
3. and all tagfunction functions and closures with regular ones

in this thread.

The function noinstr_{Th_2} modifies the stack of the thread by

- removing any pair of the form $(env, \text{return}(\mathcal{MI}, [\cdot]))$,
- and transforming each pair of the form $(env, \text{event}(ev(e_1, \dots, e_{i-1}, [\cdot], v_{i+1}, \dots, v_n)))$ into the pair $(env, (e_1, \dots, e_{i-1}, [\cdot], v_{i+1}, \dots, v_n))$.

Let $\text{noinstr}_{Th} \stackrel{\text{def}}{=} \text{noinstr}_{Th_1} \circ \text{noinstr}_{Th_2}$.

Finally, let us define

$$\text{noinstr}_{\mathcal{CI}}([Th_1, \dots, Th_n], \text{globalstore}, tj, \mathcal{MI}, \text{events}) \stackrel{\text{def}}{=} [\text{noinstr}_{Th}(Th_1), \dots, \text{noinstr}_{Th}(Th_n)], \text{globalstore}, tj$$

We do not need to replace elements of the global store, as they cannot contain closures: event, return, and tagged functions cannot appear in them.

The next proposition shows that, with Assumption 4.1, there is a weak bisimulation between the non-instrumented semantics and the instrumented semantics, that is, the reductions match in the two semantics, but the number of steps may differ. Indeed, the return and event expressions introduce an additional transition in the instrumented semantics. All other constructs reduce in the same number of steps in both semantics. Hence, the instrumentation does not alter the semantics of the language. This result is proved in Appendix B.

Proposition 4.5 1. If $\mathcal{C} \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI})$ and $\mathcal{C}_1, \dots, \mathcal{C}_n$ are pairwise distinct configurations such that for all $i \leq n$, we have $\mathcal{C} \rightarrow_{p_i} \mathcal{C}_i$ with $\sum_{i \leq n} p_i = 1$, then there exist pairwise distinct instrumented configurations $\mathcal{CI}_1, \dots, \mathcal{CI}_n$ such that for all $i \leq n$, we have $\mathcal{CI} \rightarrow_{p_i}^* \mathcal{CI}_i$ and $\mathcal{C}_i \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI}_i)$.

2. If $\mathcal{C} \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI})$ and $\mathcal{CI}_1, \dots, \mathcal{CI}_n$ are pairwise distinct instrumented configurations such that for all $i \leq n$, we have $\mathcal{CI} \rightarrow_{p_i} \mathcal{CI}_i$ with $\sum_{i \leq n} p_i = 1$, then there exist pairwise distinct configurations $\mathcal{C}_1, \dots, \mathcal{C}_n$ such that for all $i \leq n$, we have $\mathcal{C} \rightarrow_{p_i}^* \mathcal{C}_i$ and $\mathcal{C}_i \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI}_i)$.

In the rest of the paper, we use only the instrumented semantics. Furthermore, we denote instrumented configurations by \mathcal{C} to lighten notations.

5 Translation

In this section, we describe how our compiler translates an annotated CryptoVerif process. It translates each CryptoVerif role `role` into an OCaml module μ_{role} and each CryptoVerif oracle into a function. Let \mathbb{G}_{var} be an injective function that takes a CryptoVerif variable name and returns an OCaml variable name.

The function $\mathbb{G}_{\mathbb{M}}$ transforms a CryptoVerif term M into an OCaml term. It is defined as follows:

$$\begin{aligned} \mathbb{G}_{\mathbb{M}}(x[\tilde{i}]) &\stackrel{\text{def}}{=} \mathbb{G}_{\text{var}}(x) && \text{(Variable)} \\ \mathbb{G}_{\mathbb{M}}(f(M_1, \dots, M_m)) &\stackrel{\text{def}}{=} \mathbb{G}_f(f) (\mathbb{G}_{\mathbb{M}}(M_1), \dots, (\mathbb{G}_{\mathbb{M}}(M_m))) && \text{(Function call)} \end{aligned}$$

Before defining the translation of an oracle, let us first introduce some notations. For each CryptoVerif variable x , we denote by T_x the type of x , and by extension, for each CryptoVerif term M , we denote by T_M the type of M . More precisely, if M is the variable x , then $T_M \stackrel{\text{def}}{=} T_x$, and if M is a function application with a function of type $T_1 \times \dots \times T_n \rightarrow T$, then $T_M \stackrel{\text{def}}{=} T$.

We say that an oracle or role definition occurs *at the beginning of* Q when it occurs in Q just under replication or parallel composition. We define the function `reduce'` that returns a description of the oracles made available by an oracle definition Q . In more detail, `reduce'(Q)` is a list $[(Q_1, b_1), \dots, (Q_l, b_l)]$ such that Q_1, \dots, Q_l are the oracle definitions at the beginning of Q , from left to right, and the boolean b_l is true when Q_l is under replication, and false otherwise. In this function, the replication indices \tilde{i} can be partially instantiated into integer values. In contrast to the function `reduce`, `reduce'(foreach $i \leq n$ do Q)` does not instantiate the replication index i .

$$\begin{aligned} \text{reduce}'(0) &\stackrel{\text{def}}{=} [] && \text{(Nil)} \\ \text{reduce}'(Q_1 \mid Q_2) &\stackrel{\text{def}}{=} \text{reduce}'(Q_1) @ \text{reduce}'(Q_2) && \text{(Par)} \end{aligned}$$

$$\begin{aligned}
& \text{reduce}'(\text{foreach } i \leq n \text{ do } Q) \stackrel{\text{def}}{=} [(Q_1, \text{true}), \dots, (Q_l, \text{true})] \\
& \quad \text{when } \text{reduce}'(Q) = [(Q_1, b_1), \dots, (Q_l, b_l)] \text{ for some } b_1, \dots, b_l \\
& \hspace{20em} (\text{Repl}) \\
& \text{reduce}'(O[\tilde{i}](x_1[\tilde{i}], \dots, x_k[\tilde{i}]) := P) \stackrel{\text{def}}{=} [(O[\tilde{i}](x_1[\tilde{i}], \dots, x_k[\tilde{i}]) := P, \text{false})] \\
& \hspace{20em} (\text{Oracle}) \\
& \text{reduce}'(\text{role } \{Q\}) \stackrel{\text{def}}{=} [] \hspace{10em} (\text{Role})
\end{aligned}$$

The function reduce' takes processes Q that follow **return** statements that do not end a role. By Assumption 2.9, we are inside a role, so by Property 2.8, the construct $\text{role } \{Q'\}$ cannot appear in Q before a **return** statement that ends the current oracle. So, the function reduce' will never be called on $\text{role } \{Q'\}$.

We also define the function $\mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}$ that returns a description of the modules that correspond to roles defined at the beginning of an oracle definition Q . The function $\mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}$ is similar to the function reduce' above: it returns pairs containing the module generated for the role and a boolean indicating whether the role is under replication or not. In contrast to reduce' , it returns a set and not a list.

$$\begin{aligned}
& \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(0) \stackrel{\text{def}}{=} \emptyset \hspace{10em} (\text{Nil}) \\
& \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q_1 \mid Q_2) \stackrel{\text{def}}{=} \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q_1) \cup \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q_2) \hspace{2em} (\text{Par}) \\
& \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(\text{foreach } i \leq n \text{ do } Q) \stackrel{\text{def}}{=} \{(\mu, \text{true}) \mid \exists b, (\mu, b) \in \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q)\} \hspace{2em} (\text{Repl}) \\
& \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(O[\tilde{i}](x_1[\tilde{i}], \dots, x_k[\tilde{i}]) := P) \stackrel{\text{def}}{=} \emptyset \hspace{10em} (\text{Oracle}) \\
& \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(\text{role } \{Q\}) \stackrel{\text{def}}{=} \{(\mu_{\text{role}}, \text{false})\} \hspace{2em} (\text{Role})
\end{aligned}$$

The function $\mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}$ takes processes Q that follow **return** statements that end the current role. By Assumption 2.9, there cannot be an oracle definition before a $\text{role } \{Q'\}$ in Q . So the function $\mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}$ will never be called on oracle definitions.

To translate an oracle, we translate the body of the oracle using the function \mathbb{G} defined in Figure 16. Most cases are straightforward. After defining a variable, we store it in a file if needed, using $\mathbb{G}_{\text{file}}(x[\tilde{i}])$, defined by $\mathbb{G}_{\text{file}}(x[\tilde{i}]) \stackrel{\text{def}}{=} (f := \mathbb{G}_{\text{ser}}(T_x) \mathbb{G}_{\text{var}}(x))$ if $(x[\tilde{i}], f) \in \text{files}$ and $\mathbb{G}_{\text{file}}(x[\tilde{i}]) \stackrel{\text{def}}{=} ()$ otherwise. A file is modeled by a global store location.

For the **return** case, if the **return** is not at the end of a role, we return the closures corresponding to the oracles defined after the **return**, as defined in (Return1). Otherwise, we update the set of available roles using the primitive **return** introduced in Section 4, as defined in (Return2).

In the **insert** case, we add the inserted element to the considered table Tbl contained in the global store location f . In the **get** case, we read the table by read_table , keeping only the elements that satisfy the required condition (which is tested by \mathbb{G}_{test}). These elements are stored in the list l . If l is empty, we run P' ; otherwise, we choose a random element in l and run P with that element. To choose that element, we use a function random_l such that $\text{random}_l l$ returns a random element of the list l , such that the probability of returning

$$\begin{array}{l}
\mathbb{G}(x[\tilde{i}] \stackrel{R}{\leftarrow} T; P) \stackrel{\text{def}}{=} \text{let } \mathbb{G}_{\text{var}}(x) = \mathbb{G}_{\text{random}}(T) () \text{ in } \mathbb{G}_{\text{file}}(x[\tilde{i}]); \mathbb{G}(P) \quad (\text{New}) \\
\mathbb{G}(x[\tilde{i}] \leftarrow M; P) \stackrel{\text{def}}{=} \text{let } \mathbb{G}_{\text{var}}(x) = \mathbb{G}_{\text{M}}(M) \text{ in } \mathbb{G}_{\text{file}}(x[\tilde{i}]); \mathbb{G}(P) \quad (\text{Let}) \\
\mathbb{G}(\text{if } M \text{ then } P \text{ else } P') \stackrel{\text{def}}{=} \text{if } \mathbb{G}_{\text{M}}(M) \text{ then } \mathbb{G}(P) \text{ else } \mathbb{G}(P') \quad (\text{If}) \\
\frac{[(Q_1, b_1), \dots, (Q_l, b_l)] \stackrel{\text{def}}{=} \text{reduce}'(Q)}{\mathbb{G}(\text{return}(N_1, \dots, N_k); Q) \stackrel{\text{def}}{=} (\mathbb{G}_{\text{O}}(Q_1, b_1), \dots, \mathbb{G}_{\text{O}}(Q_l, b_l), \mathbb{G}_{\text{M}}(N_1), \dots, \mathbb{G}_{\text{M}}(N_k))} \quad (\text{Return1}) \\
\mathbb{G}(\text{return}(N_1, \dots, N_k)); Q) \stackrel{\text{def}}{=} (\text{return}(\mathbb{G}_{\text{get}_{\text{Mx}}}(Q), (\mathbb{G}_{\text{M}}(N_1), \dots, \mathbb{G}_{\text{M}}(N_k)))) \quad (\text{Return2}) \\
\mathbb{G}(\text{end}) \stackrel{\text{def}}{=} (\text{raise Match_failure}) \quad (\text{End}) \\
\mathbb{G}(\text{event } ev(M_1, \dots, M_k); P) \stackrel{\text{def}}{=} \text{event } ev(\mathbb{G}_{\text{M}}(M_1), \dots, \mathbb{G}_{\text{M}}(M_k)); \mathbb{G}(P) \quad (\text{Event}) \\
\frac{(Tbl, f) \in \text{tables}}{\mathbb{G}(\text{insert } Tbl(M_1, \dots, M_k); P) \stackrel{\text{def}}{=} (f := (\mathbb{G}_{\text{ser}}(T_{M_1}) \mathbb{G}_{\text{M}}(M_1), \dots, \mathbb{G}_{\text{ser}}(T_{M_k}) \mathbb{G}_{\text{M}}(M_k)) :: (!f); \mathbb{G}(P))} \quad (\text{Insert}) \\
\mathbb{G}_{\text{test}}((x_1, \dots, x_k), M) \stackrel{\text{def}}{=} \\
\begin{array}{l}
(\text{function } (\mathbb{G}_{\text{var}}(x_1), \dots, \mathbb{G}_{\text{var}}(x_k)) \rightarrow \\
\quad \text{let } \mathbb{G}_{\text{var}}(x_1) = \mathbb{G}_{\text{deser}}(T_{x_1}) \mathbb{G}_{\text{var}}(x_1) \text{ in } \dots \\
\quad \text{let } \mathbb{G}_{\text{var}}(x_k) = \mathbb{G}_{\text{deser}}(T_{x_k}) \mathbb{G}_{\text{var}}(x_k) \text{ in} \quad (\text{Test}) \\
\quad \text{if } (\mathbb{G}_{\text{M}}(M)) \text{ then } (\mathbb{G}_{\text{var}}(x_1), \dots, \mathbb{G}_{\text{var}}(x_k)) \\
\quad \quad \text{else raise Match_failure} \\
| _ \rightarrow \text{raise Bad_file}) \\
\mathbb{G}_{\text{fold}} \stackrel{\text{def}}{=} f \rightarrow \text{function } a \rightarrow \text{function } [] \rightarrow a \mid x :: l \rightarrow f (\text{fold } f a l) x \quad (\text{Fold}) \\
\text{read_table}(f, c) \stackrel{\text{def}}{=} \text{let rec fold} = \text{function } \mathbb{G}_{\text{fold}} \text{ in} \\
\quad \text{fold} (\text{function } a \rightarrow \text{function } x \rightarrow \quad (\text{Read table}) \\
\quad \quad (\text{try } (c x) :: a \text{ with} \\
\quad \quad \quad \text{Match_failure} \rightarrow a)) [] !f \\
\frac{(Tbl, f) \in \text{tables}}{\mathbb{G}(\text{get } Tbl(x_1[\tilde{i}], \dots, x_k[\tilde{i}]) \text{ suchthat } M \text{ in } P \text{ else } P') \stackrel{\text{def}}{=} \\
\quad \text{let } l = \text{read_table}(f, \mathbb{G}_{\text{test}}((x_1, \dots, x_k), M)) \text{ in} \\
\quad \text{if } l = [] \text{ then } \mathbb{G}(P') \\
\quad \quad \text{else let } (\mathbb{G}_{\text{var}}(x_1), \dots, \mathbb{G}_{\text{var}}(x_k)) = \text{random}_l l \text{ in} \\
\quad \quad \quad (\mathbb{G}_{\text{file}}(x_1[\tilde{i}]); \dots; \mathbb{G}_{\text{file}}(x_k[\tilde{i}]); \mathbb{G}(P))} \quad (\text{Get})
\end{array}
\end{array}$$

Figure 16: Translation function \mathbb{G} of an oracle body in OCaml

$$\begin{aligned} \mathbb{G}_O(Q, \text{false}) &\stackrel{\text{def}}{=} \text{let } token = \text{ref true in tagfunction}^O pm_{\text{false}}(Q) \\ \text{where } pm_{\text{false}}(O[\tilde{i}](x_1[\tilde{i}] : T_1, \dots, x_k[\tilde{i}] : T_k) := P) &\stackrel{\text{def}}{=} \\ &(\mathbb{G}_{\text{var}}(x_1), \dots, \mathbb{G}_{\text{var}}(x_k)) \rightarrow \\ &\text{if } (!token) \&\& (\mathbb{G}_{\text{pred}}(T_1) \mathbb{G}_{\text{var}}(x_1)) \&\& \dots \&\& (\mathbb{G}_{\text{pred}}(T_k) \mathbb{G}_{\text{var}}(x_k)) \\ &\text{then } (token := \text{false}; \mathbb{G}_{\text{file}}(x_1[\tilde{i}]); \dots; \mathbb{G}_{\text{file}}(x_k[\tilde{i}]); \mathbb{G}(P)) \\ &\text{else raise Bad_Call} \end{aligned} \tag{Oracle1}$$

$$\begin{aligned} \mathbb{G}_O(Q, \text{true}) &\stackrel{\text{def}}{=} \text{tagfunction}^O pm_{\text{true}}(Q) \\ \text{where } pm_{\text{true}}(O[\tilde{i}](x_1[\tilde{i}] : T_1, \dots, x_k[\tilde{i}] : T_k) := P) &\stackrel{\text{def}}{=} \\ &(\mathbb{G}_{\text{var}}(x_1), \dots, \mathbb{G}_{\text{var}}(x_k)) \rightarrow \\ &\text{if } (\mathbb{G}_{\text{pred}}(T_1) \mathbb{G}_{\text{var}}(x_1)) \&\& \dots \&\& (\mathbb{G}_{\text{pred}}(T_k) \mathbb{G}_{\text{var}}(x_k)) \\ &\text{then } (\mathbb{G}_{\text{file}}(x_1[\tilde{i}]); \dots; \mathbb{G}_{\text{file}}(x_k[\tilde{i}]); \mathbb{G}(P)) \\ &\text{else raise Bad_Call} \end{aligned} \tag{Oracle2}$$

Figure 17: Translation of an oracle

the j -th element of l is among $(\{1, \dots, |l|\}, j)$. We suppose that this function is programmed from the OCaml primitive `random`, and is present in the module for cryptographic primitives μ_{prim} .

An oracle $O(x_1, \dots, x_n) := P$ is transformed into a closure by the function \mathbb{G}_O as shown in Figure 17. When the oracle O is not under replication (second argument of \mathbb{G}_O `false`, in (Oracle1)), we use a boolean token $token$ to make sure that it can be called only once: $token$ is initially true, it is set to false in the first call. In subsequent calls, the exception `Bad_Call` will be raised. The translation of an oracle always checks that the arguments are correct values for their CryptoVerif types, and stores them in files if necessary by calling \mathbb{G}_{file} .

Finally, we generate an OCaml module μ_{role} for each role $role$ in the CryptoVerif process. This module provides a single function `init`, which returns the functions implementing the oracles defined at the beginning of $Q(\text{role})$, so its interface is $\text{interface}(\mu_{\text{role}}) \stackrel{\text{def}}{=} \{\mu_{\text{role}}.\text{init}\}$ and its program is

$$\begin{aligned} \text{program}(\mu_{\text{role}}) &\stackrel{\text{def}}{=} \text{let } \mu_{\text{role}}.\text{init} = \text{let } token = \text{ref true in tagfunction}^{\text{role}} pm_{\text{role}} \\ \text{where } pm_{\text{role}} &\stackrel{\text{def}}{=} () \rightarrow \\ &\text{if } (!token) \text{ then} \\ &\quad (token := \text{false}; \\ &\quad \mathbb{G}_{\text{read}}(x_1[]) \text{ in } \dots \text{ in } \mathbb{G}_{\text{read}}(x_m[]) \text{ in} \\ &\quad (\mathbb{G}_O(Q_1, b_1), \dots, \mathbb{G}_O(Q_k, b_k))) \\ &\text{else raise Bad_Call} \end{aligned}$$

where $[(Q_1, b_1), \dots, (Q_k, b_k)] = \text{reduce}'(Q(\text{role}))$ and $x_1[], \dots, x_m[]$ are the free variables of $Q(\text{role})$, which are the variables we need to retrieve from the files. The function $\mathbb{G}_{\text{read}}(x[])$, which reads the contents of the file associated to $x[]$, is defined by $\mathbb{G}_{\text{read}}(x[]) \stackrel{\text{def}}{=} \text{let } \mathbb{G}_{\text{var}}(x) = \mathbb{G}_{\text{deser}}(T_x) !(f) \text{ if } (x[], f) \in \text{files.}$

Example 4 Let us explain the translation of the role `alice` described in Examples 1 and 2. This role contains only the following oracle

$$\text{OA}() := m \stackrel{R}{\leftarrow} \text{nonce}; s \stackrel{R}{\leftarrow} \text{seed}; \text{return}(m, \text{sign}(m, sk, s))$$

This role is translated into the module μ_{alice} . Its program $\text{program}(\mu_{\text{alice}})$ is:

```

let  $\mu_{\text{alice}}.\text{init} = \text{let } token = \text{ref true in tagfunction}^{\text{alice}} () \rightarrow$ 
  if (! $token$ ) then
    ( $token := \text{false};$ 
      $\mathbb{G}_{\text{read}}(sk[])$  in
     let  $token = \text{ref true in tagfunction}^{\text{OA}} () \rightarrow$ 
       if (! $token$ ) then
         ( $token := \text{false};$ 
          let  $\mathbb{G}_{\text{var}}(m) = \mathbb{G}_{\text{random}}(\text{nonce}) ()$  in
          let  $\mathbb{G}_{\text{var}}(s) = \mathbb{G}_{\text{random}}(\text{seed}) ()$  in
          ( $\mathbb{G}_{\text{var}}(m), \mathbb{G}_{\text{f}}(\text{sign}) (\mathbb{G}_{\text{var}}(m), \mathbb{G}_{\text{var}}(sk), \mathbb{G}_{\text{var}}(s))$ ))
        else raise Bad_Call)
     else raise Bad_Call

```

This program defines the function $\mu_{\text{alice}}.\text{init}$, which expects $()$ as argument and returns the function that implements oracle OA. This function itself expects $()$ as argument and returns the OCaml representation of the message $m, \text{sign}(m, sk, s)$ returned by OA.

The function $\mu_{\text{alice}}.\text{init}$ can be called only once, which is guaranteed using a boolean reference $token$. If $token$ is already false, then this function has already been called, so we raise the exception `Bad_Call`. Otherwise, we set $token$ to false, we read the value of sk from the file $skfile$ and we continue by the translation of the CryptoVerif oracle OA. The oracle OA can also be called only once (but the program $\text{program}(\mu_{\text{alice}})$ can be launched N_1 times in different threads). So we define a new boolean reference $token$, to guarantee this property, and define the function that implements OA. When this function is called for the first time, it sets this second $token$ to false, creates a new nonce $\mathbb{G}_{\text{var}}(m)$ and a new seed $\mathbb{G}_{\text{var}}(s)$, and finally returns the nonce $\mathbb{G}_{\text{var}}(m)$ with the signature. When this function is called again, it raises the exceptional value `Bad_Call`.

To call the translation of oracle OA, one can execute:

$$\mu_{\text{alice}}.\text{init} () ()$$

This code initializes the role and calls the closure returned by $\mu_{\text{alice}}.\text{init} ()$, which corresponds to the translation of OA.

The generated modules \mathcal{M}_g (μ_{role} for each role in the CryptoVerif process) are included in manually-written programs that represent the full implementation

of the protocol, for instance a client and a server. In particular, these programs are responsible for sending the result of oracles to the network and receiving messages to be passed as arguments to oracles. These programs interact with an adversary that we model as an OCaml program $program_0$. We consider that the programs of the protocol are launched by the adversary $program_0$ using the `addthread` construct. The generated modules depend only on the module containing the cryptographic primitives μ_{prim} , so when the program of a thread uses the primitives or the generated modules, we can order the programs of the modules in the argument of `addthread` in the order $program(\mu_{\text{prim}}); program(\mu_{\text{role}_1}); \dots; program(\mu_{\text{role}_k}); program'$ where $program'$ contains no generated module, as required by the instrumented semantics of `addthread` (New toplevel add thread). We assume that $program_0$ uses the generated modules only inside `addthread`, and that $program_0$ is a well-typed OCaml program. (The network code is well-typed by Assumption A4. The adversary itself is any probabilistic Turing machine, which can be implemented by a well-typed OCaml program.) We assume that only the generated modules use events, tagged functions, and `return`. The adversary must not use events, which serve for specifying security properties of the protocol, nor `return`, which serves for updating the set of callable generated modules. He uses regular functions rather than tagged functions. Moreover, as mentioned in Assumption A3, we suppose that only the generated modules access files that contain private CryptoVerif data (free variables of roles and tables). So we let $S_{\text{priv}} \stackrel{\text{def}}{=} \{f \mid (x[], f) \in \text{files} \text{ or } (Tbl, f) \in \text{tables}\} \subseteq S_g$ be the set of global locations reserved for private CryptoVerif data, and we have the following assumption:

Assumption 5.1 *The locations in S_{priv} occur only in the programs of generated modules; they do not occur elsewhere in $program_0$.*

The program $program_0$ is run in the initial (instrumented) OCaml configuration $\mathcal{C}_0(Q_0, program_0)$ defined as follows:

$$\mathcal{C}_0(Q_0, program_0) \stackrel{\text{def}}{=} [(\emptyset, program_0, [], \emptyset), globalstore_0, 1, \mathbb{G}_{\text{get}, \mathcal{M}\mathcal{I}}(Q_0), []]$$

where $\mathbb{G}_{\text{get}, \mathcal{M}\mathcal{I}}(Q_0)$ is the set of modules available at the beginning of the execution and $globalstore_0 \stackrel{\text{def}}{=} \{l \mapsto \text{initval}_l \mid l \in S_g\}$ is the initial value of global store, as defined in Section 3.2.6. Tables are represented by lists, and their initial value initval_l is the empty list `[]`, representing that the tables are initially empty. Files that contain free variables of roles are represented by strings, and their initial value initval_l is the empty string `""`. For other elements, the initial value initval_l is the default value for the type of location l .

6 Proof of Security

This section presents the proof of correctness of our compiler. We give ourselves a CryptoVerif process Q_0 that corresponds to a cryptographic protocol. By using our compiler, we first generate modules \mathcal{M}_g that correspond to the




Figure 18: Overview of our proof

roles present inside Q_0 , as explained in the previous section. We consider an adversary interacting with the protocol implementation, modeled as an OCaml program $program_0$ that uses the generated modules in \mathcal{M}_g . Informally, when CryptoVerif shows that Q_0 satisfies a certain security property, it shows that for any CryptoVerif adversary Q , the probability that $Q \mid Q_0$ breaks the security property is bounded by a certain bound, which CryptoVerif computes. Our goal is to show that the same probability bound also applies to the generated implementation, that is, the probability that $program_0$ breaks the security property is bounded by the same bound. As illustrated in Figure 18, to prove this property, we build from the OCaml adversary $program_0$ a CryptoVerif adversary $Q_{adv}(Q_0, program_0)$ that simulates $program_0$. Basically, we run the OCaml program $program_0$ inside a CryptoVerif primitive *simulate*; when $program_0$ would call the translation of an oracle in \mathcal{M}_g , $Q_{adv}(Q_0, program_0)$ calls the corresponding oracle in Q_0 . We prove that $Q_{adv}(Q_0, program_0) \mid Q_0$ and $program_0$ using \mathcal{M}_g behave similarly, hence they have the same probability of breaking the security property. To achieve this goal, we need to prove, firstly, that the translations of the oracles behave in the same way as the CryptoVerif oracles, and secondly, that our simulation is sound.

In Section 6.1, we state our assumptions on the cryptographic primitives, and show that the primitives behave correctly independently of the rest of the program. In Section 6.2, we prove that the OCaml translation of a CryptoVerif oracle behaves like the oracle. In Section 6.3, we define the CryptoVerif adversary that simulates the OCaml adversary $program_0$. Finally, in Section 6.4, we prove that the CryptoVerif adversary interacting with Q_0 behaves like the OCaml adversary interacting with the generated implementation. This result shows the desired correctness of our compiler.

6.1 Correctness of Cryptographic Primitives

Let us first formalize the assumptions we make about the implementation of cryptographic primitives. Let $program_{prim} \stackrel{\text{def}}{=} program(\mu_{prim})$ be the program of the module that defines the primitives and $interface_{prim} \stackrel{\text{def}}{=} interface(\mu_{prim})$ be its interface. The interface $interface_{prim}$ consists of the function $random_l$, the functions $\mathbb{G}_f(f)$ for each CryptoVerif function f , and the functions $\mathbb{G}_{random}(T)$, $\mathbb{G}_{ser}(T)$, $\mathbb{G}_{deser}(T)$, and $\mathbb{G}_{pred}(T)$ for each CryptoVerif type T for which these

functions are used in the translation, as described in Section 2.3. (The functions $\mathbb{G}_{\text{ser}}(T)$ and $\mathbb{G}_{\text{deser}}(T)$ are either both present or both absent in $\text{interface}_{\text{prim}}$.) We rely on the following assumptions.

Assumption 6.1 *There are no `schedule`, `addthread`, `return`, nor event operations and no global store locations in $\text{program}_{\text{prim}}$.*

An OCaml semantic configuration in which the current thread does not use `addthread`, `return`, `event`, `schedule` operations, nor global store locations reduces by using the rule (Thread) of Figure 13, so we can reduce it by considering as configuration only a thread Th . We denote by \mathcal{TT} the traces over threads.

Let $Th_0^s \stackrel{\text{def}}{=} \langle \emptyset, \text{program}_{\text{prim}};;, [], \emptyset \rangle$ be a thread configuration that evaluates only the implementation of the cryptographic primitives module.

Assumption 6.2 *There exists a unique complete thread trace \mathcal{TT} beginning at the configuration Th_0^s and there exists env_{prim} such that the last configuration of the trace \mathcal{TT} is:*

$$Th = \langle \text{env}_{\text{prim}}, \varepsilon, [], \emptyset \rangle$$

This assumption means that there are no uncaught exceptions, no access to the store, and no `random` operations in the initialization of the module μ_{prim} , so that the environment env_{prim} is always the same. Typically, the initialization just defines functions, so this assumption is not restrictive. Random choices and a limited access to the store explained below are allowed during calls to primitives. By definition of a module, we have $\text{interface}_{\text{prim}} \subseteq \text{Dom}(\text{env}_{\text{prim}})$.

Assumption 6.3 *For each `CryptoVerif` type T , OCaml values of the corresponding type $\mathbb{G}_T(T)$ does not contain closures nor store or global store locations.*

In particular, bitstrings received or returned by cryptographic primitives are typically represented by strings. Strings are values in our semantics, and do not contain locations. So one cannot alter the contents of a string after creation: string values are not mutable. In OCaml, the type `string` is mutable, so either one should use an immutable abstract type instead of `string` to represent bitstrings, or one needs to assume that the network code does not modify the contents of strings that go through our generated closures, as mentioned in Assumption A5.

To establish the correspondence between `CryptoVerif` values and OCaml values, we define a function $\mathbb{G}_{\text{val}T}$, which maps each `CryptoVerif` bitstring a to its associated value v in OCaml. For a given type T , $\mathbb{G}_{\text{val}T}$ must be a bijection between T and the set of OCaml values of type $\mathbb{G}_T(T)$ satisfying the predicate function $\mathbb{G}_{\text{pred}}(T)$. Furthermore, the OCaml value `true` and the `CryptoVerif` value `true` are such that $\mathbb{G}_{\text{val}bool}(\text{true}) = \text{true}$, and the same goes for `false`. We extend this function to events by $\mathbb{G}_{\text{ev}}(e(a_1, \dots, a_j)) = e(\mathbb{G}_{\text{val}T_1}(a_1), \dots, \mathbb{G}_{\text{val}T_j}(a_j))$ if e is of type $T_1 \times \dots \times T_j$. This function is naturally extended to lists of events.

The next assumption states that the primitives have been correctly implemented, following Assumption A1: the implementation of the cryptographic

primitives in $interface_{\text{prim}}$ correctly emulates the corresponding behavior of CryptoVerif.

Assumption 6.4 (Correct primitives) *For each CryptoVerif function f of type $T_1 \times \dots \times T_n \rightarrow T$, for each CryptoVerif values a_1, \dots, a_n of types T_1, \dots, T_n , there exist env and store such that*

$$\langle \emptyset, env_{\text{prim}}(\mathbb{G}_f(f)) (\mathbb{G}_{\text{val}T_1}(a_1), \dots, \mathbb{G}_{\text{val}T_n}(a_n)), [], \emptyset \rangle \rightarrow^* \langle env, \mathbb{G}_{\text{val}T}(f(a_1, \dots, a_n)), [], store \rangle.$$

For each CryptoVerif type T such that the function $\mathbb{G}_{\text{random}}(T)$ is in $interface_{\text{prim}}$, for each CryptoVerif value $a \in T$, there exist env and store such that

$$\langle \emptyset, env_{\text{prim}}(\mathbb{G}_{\text{random}}(T)) (), [], \emptyset \rangle \rightarrow_{1/|T|}^* \langle env, \mathbb{G}_{\text{val}T}(a), [], store \rangle.$$

For each CryptoVerif type T such that the function $\mathbb{G}_{\text{pred}}(T)$ is in $interface_{\text{prim}}$, for each value v of the OCaml type $\mathbb{G}_T(T)$, there exist env and store such that

$$\langle \emptyset, env_{\text{prim}}(\mathbb{G}_{\text{pred}}(T)) v, [], \emptyset \rangle \rightarrow^* \langle env, v', [], store \rangle$$

where $v' = \text{true}$ when $\mathbb{G}_{\text{val}T}^{-1}(v)$ exists, and $v' = \text{false}$ otherwise.

For each CryptoVerif type T such that the functions $\mathbb{G}_{\text{ser}}(T)$ and $\mathbb{G}_{\text{deser}}(T)$ are in $interface_{\text{prim}}$, for each CryptoVerif value $a \in T$, there exists an OCaml string value $\text{ser}(T, a)$, such that there exist env and store such that

$$\langle \emptyset, env_{\text{prim}}(\mathbb{G}_{\text{ser}}(T)) \mathbb{G}_{\text{val}T}(a), [], \emptyset \rangle \rightarrow^* \langle env, \text{ser}(T, a), [], store \rangle$$

and there exist env and store such that

$$\langle \emptyset, env_{\text{prim}}(\mathbb{G}_{\text{deser}}(T)) \text{ser}(T, a), [], \emptyset \rangle \rightarrow^* \langle env, \mathbb{G}_{\text{val}T}(a), [], store \rangle.$$

If v is a non-empty list, then for each $a \in v$, there exist env and store such that

$$\langle \emptyset, env_{\text{prim}}(\text{random}_l) v, [], \emptyset \rangle \rightarrow_{\sum_{j \in S} \text{among}(\{1, \dots, |v|\}, j)}^* \langle env, a, [], store \rangle$$

where $S \stackrel{\text{def}}{=} \{1 \leq j \leq |v| \mid \text{nth}(v, j) = a\}$.

In contrast to the conference version of this paper [9], in this paper, we allow the cryptographic primitives to use the store for their internal computations (which often happens in practice). However, the primitives are not allowed to communicate across calls or to communicate data to the adversary or to the rest of the code using the store. They can create new locations and use these locations internally; these locations become unreachable as soon as the call to the primitive ends. This assumption corresponds to the intuition that the cryptographic primitives are pure functions: their usage of the store should not have any visible side effect. This assumption is modeled above by considering that the primitives are initially called in an empty store. Since their return value does not contain locations, the store at the end of the call will be unreachable.

The last statement of Assumption 6.4 guarantees that random_l is programmed correctly: $\text{random}_l v$ returns a random element of the list v , such that the probability of returning the j -th element of v is $\text{among}(\{1, \dots, |v|\}, j)$. In case the same element occurs several times in v , the probability of that element is then the sum of the probabilities of all its occurrences.

In general, when primitives make probabilistic choices, they might return the same result in several traces with a different environment and store. To simplify notations, Assumption 6.4 states that this does not happen, so that we have the same environment and store in all final configurations that yield the same result. Our proof could easily be extended to the general case if desired.

The next proposition shows that the primitives always return correct results, when they are called inside an OCaml program, so possibly with a non-empty store and a non-empty stack. We prove it in Appendix C. It is a consequence of Assumption 6.4.

Proposition 6.5 (Correct behavior of the primitives) *Let us consider a thread $Th \stackrel{\text{def}}{=} \langle env, env_{\text{prim}}(s) v, stack, store \rangle$.*

- *If $s = \mathbb{G}_f(f)$, f is a *CryptoVerif* function of type $T_1 \times \dots \times T_n \rightarrow T$, and $v = (\mathbb{G}_{\text{val}T_1}(a_1), \dots, \mathbb{G}_{\text{val}T_n}(a_n))$ for some *CryptoVerif* values a_1, \dots, a_n of types T_1, \dots, T_n , then there exist env' and $store'$ such that*

$$Th \rightarrow^* \langle env', \mathbb{G}_{\text{val}T}(f(a_1, \dots, a_n)), stack, store' \rangle.$$

- *If $s = \mathbb{G}_{\text{random}}(T)$ and $v = ()$, then for each *CryptoVerif* value $a \in T$, there exist env' and $store'$ such that*

$$Th \rightarrow_{1/|T|}^* \langle env', \mathbb{G}_{\text{val}T}(a), stack, store' \rangle.$$

- *If $s = \mathbb{G}_{\text{pred}}(T)$, then there exist env' and $store'$ such that*

$$Th \rightarrow^* \langle env', v', stack, store' \rangle$$

where $v' = \text{true}$ when $\mathbb{G}_{\text{val}T}^{-1}(v)$ exists, and $v' = \text{false}$ otherwise.

- *If $s = \mathbb{G}_{\text{ser}}(T)$ and $v = \mathbb{G}_{\text{val}T}(a)$, then there exist env' and $store'$ such that*

$$Th \rightarrow^* \langle env', \text{ser}(T, a), stack, store' \rangle.$$

- *If $s = \mathbb{G}_{\text{deser}}(T)$ and $v = \text{ser}(T, a)$, then there exist env' and $store'$ such that*

$$Th \rightarrow^* \langle env', \mathbb{G}_{\text{val}T}(a), stack, store' \rangle.$$

- *If $s = \text{random}_l$ and v is a non-empty list, then for each $a \in v$, there exist env' and $store'$ such that*

$$Th \rightarrow_{\sum_{j \in S} \text{among}(\{1, \dots, |v|\}, j)}^* \langle env', a, stack, store' \rangle$$

where $S \stackrel{\text{def}}{=} \{1 \leq j \leq |v| \mid \text{nth}(v, j) = a\}$.

In all cases, we have $store' \supseteq store$.

6.2 Correctness of the Translation of Oracle Bodies

In this section, we show the correctness of the translation of oracle bodies in our compiler: we show a correspondence between the semantics of the oracle body in CryptoVerif and the semantics of its translation into OCaml.

Let $\text{fv}(M)$, $\text{fv}(P)$, $\text{fv}(Q)$ be the free variables of the CryptoVerif term M and processes P and Q , respectively, defined as follows:

$$\begin{aligned}
\text{fv}(x[\tilde{i}]) &\stackrel{\text{def}}{=} \{x[\tilde{i}]\} \\
\text{fv}(f(M_1, \dots, M_k)) &\stackrel{\text{def}}{=} \bigcup_{i=1}^k \text{fv}(M_i) \\
\text{fv}(0) &\stackrel{\text{def}}{=} \emptyset \\
\text{fv}(Q \mid Q') &\stackrel{\text{def}}{=} \text{fv}(Q) \cup \text{fv}(Q') \\
\text{fv}(\text{foreach } i \leq n \text{ do } Q) &\stackrel{\text{def}}{=} \text{fv}(Q) \\
\text{fv}(O[\tilde{i}](x_1[\tilde{i}] : T_1, \dots, x_k[\tilde{i}] : T_k) := P) &\stackrel{\text{def}}{=} \text{fv}(P) \setminus \{x_1[\tilde{i}], \dots, x_k[\tilde{i}]\} \\
\text{fv}(\text{return}(M_1, \dots, M_k); Q) &\stackrel{\text{def}}{=} \bigcup_{i=1}^k \text{fv}(M_i) \cup \text{fv}(Q) \\
\text{fv}(\text{end}) &\stackrel{\text{def}}{=} \emptyset \\
\text{fv}(x[\tilde{i}] \stackrel{R}{\leftarrow} T; P) &\stackrel{\text{def}}{=} \text{fv}(P) \setminus \{x[\tilde{i}]\} \\
\text{fv}(x[\tilde{i}] \leftarrow M; P) &\stackrel{\text{def}}{=} \text{fv}(P) \setminus \{x[\tilde{i}]\} \\
\text{fv}(\text{if } M \text{ then } P \text{ else } P') &\stackrel{\text{def}}{=} \text{fv}(M) \cup \text{fv}(P) \cup \text{fv}(P') \\
\text{fv}(\text{insert } Tbl(M_1, \dots, M_k); P) &\stackrel{\text{def}}{=} \bigcup_{i=1}^k \text{fv}(M_i) \cup \text{fv}(P) \\
\text{fv}(\text{get } Tbl(x_1[\tilde{i}], \dots, x_k[\tilde{i}]) \text{ suchthat } M \text{ in } P \text{ else } P') &\stackrel{\text{def}}{=} \\
&(\text{fv}(M) \cup \text{fv}(P) \setminus \{x_1[\tilde{i}], \dots, x_k[\tilde{i}]\}) \cup \text{fv}(P') \\
\text{fv}(\text{event } e(M_1, \dots, M_k); P) &\stackrel{\text{def}}{=} \bigcup_{i=1}^k \text{fv}(M_i) \cup \text{fv}(P) \\
\text{fv}(\text{let } (x_1[\tilde{i}] : T_1, \dots, x_{k'}[\tilde{i}] : T_{k'}) = O[M_1, \dots, M_l](M'_1, \dots, M'_l) \text{ in } P \text{ else } P') &\stackrel{\text{def}}{=} \\
&\bigcup_{i=1}^l \text{fv}(M_i) \cup \bigcup_{i=1}^{k'} \text{fv}(M'_i) \cup (\text{fv}(P) \setminus \{x_1[\tilde{i}], \dots, x_{k'}[\tilde{i}]\}) \cup \text{fv}(P') \\
\text{fv}(\text{let } x[\tilde{i}] : T = \text{loop } O[M_1, \dots, M_l](M') \text{ in } P \text{ else } P') &\stackrel{\text{def}}{=} \\
&\bigcup_{i=1}^l \text{fv}(M_i) \cup \text{fv}(M') \cup (\text{fv}(P) \setminus \{x[\tilde{i}]\}) \cup \text{fv}(P')
\end{aligned}$$

The only unusual point in this definition is that variables are arrays $x[\tilde{i}]$. We extend this definition to terms and processes in which the replication indices \tilde{i} have been instantiated to bitstrings: for example, $\text{fv}(x[\tilde{a}]) = \{x[\tilde{a}]\}$. We extend this definition to sets of processes by $\text{fv}(\mathcal{Q}) = \bigcup_{Q \in \mathcal{Q}} \text{fv}(Q)$ and to stacks by $\text{fv}(\mathcal{R}) = \bigcup_{((x_1[\tilde{a}], \dots, x_k[\tilde{a}]), P_1, P_2) \in \mathcal{R}} \text{fv}(P_1) \setminus \{x_1[\tilde{a}], \dots, x_k[\tilde{a}]\} \cup \text{fv}(P_2)$.

Next, we define the OCaml value corresponding to a CryptoVerif table, and we use this definition to define the OCaml environment and global store corresponding to a CryptoVerif environment and to CryptoVerif tables.

Definition 6.6 (CryptoVerif table to OCaml list) *Let us consider a table Tbl of type $T_1 \times \dots \times T_l$. The serialized OCaml value that corresponds to an element of this table is*

$$\mathbb{G}_{\text{tbl}}(Tbl, (b_1, \dots, b_l)) \stackrel{\text{def}}{=} (\text{ser}(T_1, \mathbb{G}_{\text{val}T_1}(b_1)), \dots, \text{ser}(T_l, \mathbb{G}_{\text{val}T_l}(b_l))).$$

Let $t = [a_1; \dots; a_k]$ be the contents of the table Tbl : each a_i is an element of the

table. Let us denote

$$\mathbb{G}_{\text{tbl}}(\text{Tbl}, t) \stackrel{\text{def}}{=} [\mathbb{G}_{\text{tbl}}(a_1); \dots; \mathbb{G}_{\text{tbl}}(a_k)]$$

the OCaml list corresponding to t .

Definition 6.7 (Minimal environment and global store)

$$\begin{aligned} \text{env}(E, P) &\stackrel{\text{def}}{=} \{\mathbb{G}_{\text{var}}(x) \mapsto \mathbb{G}_{\text{val}}(b) \mid x[\tilde{a}] \in \text{fv}(P), E(x[\tilde{a}]) = b\} \quad (\text{Environment}) \\ \text{globalstore}(E, \mathcal{T}) &\stackrel{\text{def}}{=} \{f \mapsto \mathbb{G}_{\text{tbl}}(\text{Tbl}, \mathcal{T}(\text{Tbl})) \mid (\text{Tbl}, f) \in \text{tables}\} \\ &\cup \{f \mapsto \text{ser}(T_x, a) \mid (x[], f) \in \text{files}, E(x[]) = a\} \\ &\cup \{f \mapsto "" \mid (x[], f) \in \text{files}, x \text{ not defined in } E\} \\ &\quad (\text{Globalstore}) \end{aligned}$$

We define $\text{env}(E, M)$ and $\text{env}(E, Q)$ in the same way.

The *globalstore* function defined above returns the global store in which the contents of the files and the tables is correct with respect to the CryptoVerif configuration elements E and \mathcal{T} . The *env* function returns the environment corresponding to E for the free variables in P (or M , or Q).

First, we show a correspondence between a CryptoVerif term and its OCaml translation.

Lemma 6.8 (Term reduction) *Let M be a CryptoVerif term of type T . If*

$$\text{Th} = \langle \text{env}, \mathbb{G}_{\text{M}}(M), \text{stack}, \text{store} \rangle \text{ with } \text{env} \supseteq \text{env}_{\text{prim}} \cup \text{env}(E, M),$$

and $E, M \Downarrow a$, then $\text{Th} \rightarrow^* \text{Th}'$ where

$$\text{Th}' \stackrel{\text{def}}{=} \langle \text{env}', \mathbb{G}_{\text{val}T}(a), \text{stack}, \text{store}' \rangle$$

for some env' and store' such that $\text{store}' \supseteq \text{store}$.

Proof We prove this result by induction on the syntax of terms.

- Case $M = x[\tilde{a}']$: Since $E, M \Downarrow a$ is derived by (Var), we have $E(x[\tilde{a}']) = a$. Since $\text{env}(E, M) \subseteq \text{env}$, we have $\text{env}(\mathbb{G}_{\text{var}}(x)) = \mathbb{G}_{\text{val}T}(a)$, so

$$\text{Th} = \langle \text{env}, \mathbb{G}_{\text{var}}(x), \text{stack}, \text{store} \rangle \rightarrow \text{Th}' = \langle \text{env}, \mathbb{G}_{\text{val}T}(a), \text{stack}, \text{store} \rangle$$

- Case $M = f(M_1, \dots, M_k)$, where f is of type $T_1 \times \dots \times T_k \rightarrow T$. Since $E, M \Downarrow a$ is derived by (Fun), we have $E, M_i \Downarrow a_i$ for all $i \leq k$, for some a_1, \dots, a_k such that $f(a_1, \dots, a_k) = a$.

$$\begin{aligned} \text{Th} &= \langle \text{env}, \mathbb{G}_f(f) (\mathbb{G}_{\text{M}}(M_1), \dots, \mathbb{G}_{\text{M}}(M_k)), \text{stack}, \text{store} \rangle \\ &\rightarrow \langle \text{env}, (\mathbb{G}_{\text{M}}(M_1), \dots, \mathbb{G}_{\text{M}}(M_k)), \text{stack}', \text{store} \rangle \\ &\quad \text{where } \text{stack}' \stackrel{\text{def}}{=} (\text{env}, \mathbb{G}_f(f) [\cdot]) :: \text{stack} \end{aligned}$$

$$\begin{aligned}
& \rightarrow Th_1 \stackrel{\text{def}}{=} \langle env, \mathbb{G}_M(M_k), stack'', store \rangle \\
& \quad \text{where } stack'' \stackrel{\text{def}}{=} (env, (\mathbb{G}_M(M_1), \dots, \mathbb{G}_M(M_{k-1}), [\cdot])) :: stack' \\
\rightarrow^* Th_2 & \stackrel{\text{def}}{=} \langle env', \mathbb{G}_{valT_k}(a_k), stack'', store' \rangle \\
& \quad \text{by induction hypothesis applied to } M_k \\
& \rightarrow \langle env, (\mathbb{G}_M(M_1), \dots, \mathbb{G}_M(M_{k-1}), \mathbb{G}_{valT_k}(a_k)), stack', store' \rangle \\
\rightarrow^* Th_3 & \stackrel{\text{def}}{=} \langle env, (\mathbb{G}_{valT_1}(a_1), \dots, \mathbb{G}_{valT_k}(a_k)), stack', store'' \rangle \\
& \quad \text{by an easy induction} \\
& \rightarrow \langle env, \mathbb{G}_f(f) (\mathbb{G}_{valT_1}(a_1), \dots, \mathbb{G}_{valT_k}(a_k)), stack, store'' \rangle \\
& \rightarrow \langle env, \mathbb{G}_f(f), stack''', store'' \rangle \\
& \quad \text{where } stack''' \stackrel{\text{def}}{=} (env, [\cdot] (\mathbb{G}_{valT_1}(a_1), \dots, \mathbb{G}_{valT_k}(a_k))) :: stack \\
& \rightarrow \langle env, env_{\text{prim}}(\mathbb{G}_f(f)), stack''', store'' \rangle \quad \text{since } env_{\text{prim}} \subseteq env \\
& \rightarrow \langle env, env_{\text{prim}}(\mathbb{G}_f(f)) (\mathbb{G}_{valT_1}(a_1), \dots, \mathbb{G}_{valT_k}(a_k)), stack, store'' \rangle \\
\rightarrow^* Th' & \stackrel{\text{def}}{=} \langle env'', \mathbb{G}_{valT}(a), stack, store''' \rangle \quad \text{by Proposition 6.5} \\
\end{aligned}$$

By Proposition 6.5 and induction hypothesis, we have $store''' \supseteq store'' \supseteq store' \supseteq store$. \square

Let us now introduce some notations.

Definition 6.9 (Helper functions) *For an OCaml configuration*

$$\mathcal{C} = [Th_1, \dots, Th_n], globalstore, tj', \mathcal{MI}, events$$

with $Th_{tj} = \langle env_{tj}, pe_{tj}, stack_{tj}, store_{tj} \rangle$ for all $tj \leq n$, let us define the following functions:

- $\mathcal{C}_{pe}(\mathcal{C}, tj) \stackrel{\text{def}}{=} pe_{tj}, \quad \mathcal{C}_{pe}(\mathcal{C}) \stackrel{\text{def}}{=} \mathcal{C}_{pe}(\mathcal{C}, tj')$,
- $\mathcal{C}_{Th}(\mathcal{C}, tj) \stackrel{\text{def}}{=} Th_{tj}, \quad \mathcal{C}_{Th}(\mathcal{C}) \stackrel{\text{def}}{=} \mathcal{C}_{Th}(\mathcal{C}, tj')$,
- $\mathcal{C}_t(\mathcal{C}) \stackrel{\text{def}}{=} tj'$,
- $\mathcal{C}_{tmax}(\mathcal{C}) \stackrel{\text{def}}{=} n$,
- $\mathcal{C}_{globalstore}(\mathcal{C}) \stackrel{\text{def}}{=} globalstore$,
- $\mathcal{C}_{events}(\mathcal{C}) \stackrel{\text{def}}{=} events$,

We also define these functions on traces, where they return the elements in the last configuration of the trace.

We also define

$$\begin{aligned}
\mathcal{C}[Th \mapsto Th', globalstore \mapsto globalstore', MI \mapsto MI', events \mapsto events'] & \stackrel{\text{def}}{=} \\
[Th_1, \dots, Th_{tj'-1}, Th', Th_{tj'+1}, \dots, Th_n], globalstore', tj', MI', events' & .
\end{aligned}$$

In this notation, one can omit $globalstore$, MI , or $events$. When omitted, we keep the corresponding element of the configuration \mathcal{C} .

Next, we prove that the CryptoVerif oracle bodies P are correctly translated into OCaml as $\mathbb{G}(P)$. We extend the translation $\mathbb{G}(P)$ to processes in which some replication indices have been instantiated into their values, using the formulas of Section 5 where replication indices i may be replaced with their value a . It is easy to see that $\mathbb{G}(P\{a/i\}) = \mathbb{G}(P)$.

Lemma 6.10 (Inner reduction) *Let \mathfrak{C} be a CryptoVerif configuration. Suppose that the program part P of \mathfrak{C} is not in a return, end, call, or loop form. Suppose that we have n possible reductions beginning at this configuration:*

$$\mathfrak{C} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_{p_i} \mathfrak{C}_i = E_i, P_i, \mathcal{Q}, \mathcal{T}_i, \mathcal{R}, \mathcal{E}_i$$

for $i \leq n$. Let \mathcal{C} be an OCaml configuration such that

$$\begin{aligned} \mathcal{C}_{Th}(\mathcal{C}) &= \langle env, \mathbb{G}(P), stack, store \rangle \text{ with } env \supseteq env_{\text{prim}} \cup env(E, P), \\ \mathcal{C}_{globalstore}(\mathcal{C}) &\supseteq globalstore(E, \mathcal{T}), \\ \mathcal{C}_{events}(\mathcal{C}) &= \mathbb{G}_{\text{ev}}(\mathcal{E}). \end{aligned}$$

Then there exist n disjoint sets of OCaml traces $\mathcal{CTS}_1, \dots, \mathcal{CTS}_n$ all starting at \mathcal{C} such that none of these traces is a prefix of another of these traces, $\Pr[\mathcal{CTS}_i] = p_i$ for all $i \leq n$, and if \mathcal{C}' is the last configuration of a trace in \mathcal{CTS}_i , then we have $\mathcal{C}' = \mathcal{C}[\text{Th} \mapsto \text{Th}', globalstore \mapsto globalstore', events \mapsto events']$ where

$$\begin{aligned} \text{Th}' &= \langle env', \mathbb{G}(P_i), stack, store' \rangle \\ &\text{with } env' \supseteq env_{\text{prim}} \cup env(E_i, P_i) \text{ and } store' \supseteq store, \\ globalstore' &\supseteq globalstore(E_i, \mathcal{T}_i), \\ globalstore'(l) &= \mathcal{C}_{globalstore}(\mathcal{C})(l) \text{ for all } l \notin S_{\text{priv}}, \\ events' &= \mathbb{G}_{\text{ev}}(\mathcal{E}_i). \end{aligned}$$

The proof of this lemma can be found in Appendix D. This lemma is proved by cases on the process P . We use Lemma 6.8 when we need to evaluate a term. The cases `end` and `return` will be handled when we prove the invariant for the whole system. The oracle bodies that we translate into OCaml do not contain calls nor loops. This lemma shows that the following invariants are preserved during the evaluation of oracle bodies: the OCaml environment and global store contain the minimal environment and global store corresponding to the CryptoVerif configuration; the public part of the global store does not change; the OCaml and CryptoVerif events match. Locations may be added in the store, but the contents of existing locations does not change.

6.3 Simulation of the OCaml Adversary

In this section, we show how to simulate in CryptoVerif any OCaml program $program_0$ that corresponds to an adversary interacting with the protocol implementation generated from the CryptoVerif process Q_0 . Basically, we run the OCaml program $program_0$ inside the CryptoVerif primitive `simulate` (which is

possible since these primitives can represent any deterministic Turing machine). When $program_0$ needs to call an oracle of Q_0 , the primitive returns and the call is made by CryptoVerif. When $program_0$ needs to generate a random number, this generation is performed by CryptoVerif.

We assume that the OCaml program $program_0$ runs in bounded time, so makes a bounded number of oracle calls. By Assumption 2.11, when an oracle O (resp. role $role$) is under replication, this replication has bound N_O (resp. N_{role}). When oracle O is under replication, we let N_O be the maximum number of calls to the same closure $tagfunction^{O,\tau}[env, pm]$ corresponding to oracle O . When a role $role$ is under replication, we let N_{role} be the maximum number of executions of $addthread(program)$ for some $program$ that contains μ_{role} . These replication bounds are chosen such that the OCaml program $program_0$ never exhausts the number of oracle calls allowed by the CryptoVerif process. We let $N_{rand+calls}$ be the maximum number of oracle calls and random number generations that the OCaml program $program_0$ can make plus one. We let N_{steps} be the maximum number of reduction steps of the program $program_0$ in the semantics of OCaml. Formally, we use the following definition:

Definition 6.11 *The number of calls to the closure with tag O, τ in a trace \mathcal{CT} , denoted $N_{calls}(O, \tau, \mathcal{CT})$, is the number of configurations \mathcal{C} such that $\mathcal{C}_{pe}(\mathcal{C}) = tagfunction^{O,\tau}[env, pm]$ v in \mathcal{CT} excluding its last configuration.*

The number of executions of role $role$ in a trace \mathcal{CT} , denoted $N_{exec}(role, \mathcal{CT})$, is the number of configurations \mathcal{C} such that $\mathcal{C}_{pe}(\mathcal{C}) = addthread(program)$ where $program$ contains $program(\mu_{role})$ in \mathcal{CT} excluding its last configuration.

The number of random number generations in a trace \mathcal{CT} , denoted $N_{rand}(\mathcal{CT})$, is the number of transitions derived by rule (Random) in \mathcal{CT} .

We define

$$\begin{aligned}
N_O &\stackrel{\text{def}}{=} \max_{\mathcal{CT}, \tau} N_{calls}(O, \tau, \mathcal{CT}) \\
N_{role} &\stackrel{\text{def}}{=} \max_{\mathcal{CT}} N_{exec}(role, \mathcal{CT}) \\
N_{rand+calls} &\stackrel{\text{def}}{=} \max_{\mathcal{CT}} \left(N_{rand}(\mathcal{CT}) + \sum_{O, \tau} N_{calls}(O, \tau, \mathcal{CT}) \right) + 1 \\
N_{steps} &\stackrel{\text{def}}{=} \max_{\mathcal{CT}} |\mathcal{CT}|
\end{aligned}$$

where \mathcal{CT} ranges over traces that begin with the configuration $\mathcal{C}_0(Q_0, program_0)$.

While N_O is an optimal bound, N_{role} is not optimal. Consider for instance a process of the form

foreach $i \leq N_O$ do $O() := \dots$ } foreach $j \leq N_{role}$ do role { \dots

By distributing the instantiations of role on every available index i , the optimal bound of the replication j is the maximum during all executions of $program_0$ of the number of instantiations of role divided by the number of calls to O made


```

1  $Q_{\text{adv}}(Q_0, \text{program}_0) = Q_{\text{start}}(Q_0, \text{program}_0) \mid Q_c(Q_0, \text{program}_0)$ 
2  $Q_{\text{start}}(Q_0, \text{program}_0) = O_{\text{start}}() :=$ 
3    $s_0 : T_{CS} \leftarrow s_0(Q_0, \text{program}_0);$ 
4   let  $r : T_{CS} = \text{loop } O_{\text{loop}}(s_0)$  in end else end
5  $Q_c(Q_0, \text{program}_0) = \text{foreach } i' \leq N_{\text{rand+calls}}$  do
6    $O_{\text{loop}}[i'](s : T_{CS}) :=$ 
7     let  $(s' : T_{CS}, o : T_o, i : \text{bitstring}, \text{args} : \text{bitstring}) = \text{simulate}(s)$  in
8     if  $o = o_S$  then
9       return( $s'$ , false)
10    else if  $o = o_1$  then
11      let  $(a_{1,1} : T_{1,1}, \dots, a_{1,m_1} : T_{1,m_1}) = \text{args}$  in
12      let  $(i_{1,1} : [1, N_{1,1}], \dots, i_{1,n_1} : [1, N_{1,n_1}]) = i$  in
13      let  $(r_{1,1} : T'_{1,1}, \dots, r_{1,m'_1} : T'_{1,m'_1}) = O_1[i_{1,1}, \dots, i_{1,n_1}](a_{1,1}, \dots, a_{1,m_1})$  in
14        return( $\text{simulate}'_{O_1}(s', (r_{1,1}, \dots, r_{1,m'_1}))$ ), true)
15      else return( $\text{simulate}''_{O_1}(s')$ , true)
16    else if  $o = o_2$  then
17       $\vdots$ 
18    else if  $o = o_R$  then
19       $b_R \stackrel{R}{\leftarrow} \text{bool};$ 
20      return( $\text{simulate}_R(s', b_R)$ ), true)

```

Figure 19: The program $Q_{\text{adv}}(Q_0, \text{program}_0)$

before these instantiations of `role`. To get this optimal bound, we would need to associate each new instantiation of `role` to the index i with the least number of associated instantiations of `role`. Since a role is often under at most one replication, we decided not to complicate the proof with details needed to get the optimal bound.

From the OCaml program program_0 , we define a CryptoVerif adversary $Q_{\text{adv}}(Q_0, \text{program}_0)$ given in Figure 19 and explained below. We will prove that this process, when executed in parallel with Q_0 , has the same behavior as the OCaml program program_0 . The initial CryptoVerif configuration is then $\mathfrak{C}_0(Q_0, \text{program}_0) = \mathfrak{C}_i(Q_0 \mid Q_{\text{adv}}(Q_0, \text{program}_0))$.

In Figure 19, we use a `let` construct with pattern matching, which can be defined as follows. We define the function $\text{tuple}_{T_1, \dots, T_j} : T_1 \times \dots \times T_j \rightarrow \text{bitstring}$ that creates a tuple with j elements (for instance by concatenating the j bitstrings with information on their length, so that they can be unambiguously recovered), and the associated projections $\pi_{k, T_1, \dots, T_j} : \text{bitstring} \rightarrow T_k$ with $k \leq j$ (which may return any value when their argument is not a tuple with j elements). The construct `let $(x_1 : T_1, \dots, x_j : T_j) = M$ in P` is an abbreviation

for:

$$\begin{aligned} x &\leftarrow M; x_1 \leftarrow \pi_{1, T_1, \dots, T_j}(x); \dots; x_j \leftarrow \pi_{j, T_1, \dots, T_j}(x); \\ &\text{if } x = \text{tuple}_{T_1, \dots, T_j}(x_1, \dots, x_j) \text{ then } P \end{aligned}$$

where x is a fresh variable. The CryptoVerif term (M_1, \dots, M_j) is an abbreviation for $\text{tuple}_{T_1, \dots, T_j}(M_1, \dots, M_j)$, where T_1, \dots, T_j are the types of M_1, \dots, M_j , respectively.

Let O_1, \dots, O_n be the oracle names in Q_0 . We define n constants o_1, \dots, o_n which are used to designate the oracles O_1, \dots, O_n respectively, o_R which corresponds to a random choice, and o_S which corresponds to the end of the OCaml program. We define the CryptoVerif type $T_o \stackrel{\text{def}}{=} \{o_R, o_S, o_1, \dots, o_n\}$, which contains all these bitstring constants.

The adversary is mainly encoded by the function *simulate*. This function takes as argument the bitstring representation $s = \text{repr}(\mathcal{CS})$ of a simulator configuration \mathcal{CS} . The configuration \mathcal{CS} consists of a non-instrumented OCaml configuration \mathcal{C} (with some extensions to the syntax described later) and sets \mathcal{RI} and \mathcal{I} that finitely represent the callable oracles \mathcal{Q} of the CryptoVerif configuration:

$$\mathcal{CS} = \underbrace{([Th_1, \dots, Th_n], \text{globalstore}, i)}_{\mathcal{C}}, \mathcal{RI}, \mathcal{I}.$$

The function *repr* is injective. We denote its inverse by repr^{-1} . We also define a CryptoVerif type $T_{\mathcal{CS}}$ that consists of all bitstrings in the image of *repr*, that is, all bitstrings that correspond to simulator configurations \mathcal{CS} . We also use the notations of Definition 6.9 for simulator configurations.

When we call an oracle or instantiate a role under replication, we must choose an unused replication index for this replication, and call the oracle or instantiate the role with that replication index. In this simulation, we will always choose the smallest replication index that has not been used yet, so that the used indices form an interval $[1, a - 1]$ and the unused indices are in $[a, N]$ where N is the bound of the considered replication. The sets \mathcal{RI} and \mathcal{I} represent the sets of callable roles and oracles, by storing the smallest index a that is not used yet.

More precisely, the set \mathcal{RI} represents the set of callable roles with their replication indices. Elements of \mathcal{RI} are either:

- of the form $\text{role}[[a, +\infty[, \tilde{a}']]$, which means the role *role* is under replication, the roles $\text{role}[1, \tilde{a}']$ to $\text{role}[a - 1, \tilde{a}']$ have been used, and the roles $\text{role}[a, \tilde{a}']$ to $\text{role}[N_{\text{role}}, \tilde{a}']$ are usable,
- or of the form $\text{role}[\tilde{a}]$, which means that *role* is not under replication and the role *role* is callable with the replication indices \tilde{a} .

The set \mathcal{RI} never contains simultaneously $\text{role}[[a, +\infty[, \tilde{a}']]$ and $\text{role}[a'', \tilde{a}']$ for the same *role* and \tilde{a}' , and it never contains simultaneously $\text{role}[[a, +\infty[, \tilde{a}']]$ and $\text{role}[[a'', +\infty[, \tilde{a}']]$ with $a \neq a''$ for the same *role* and \tilde{a}' .

The set \mathcal{I} represents the set of callable oracles with their replication indices. Elements of \mathcal{I} are either:

- of the form $O[[a, +\infty[, \tilde{a}']]$, which means that the oracle O is under replication and the oracles $O[1, \tilde{a}']$ to $O[a-1, \tilde{a}']$ have been used, and the oracles $O[a, \tilde{a}']$ to $O[N_O, \tilde{a}']$ are usable,
- or of the form $O[\tilde{a}]$ which means that O is an oracle not under replication that can be called with the replication indices \tilde{a} .

The set \mathcal{I} never contains simultaneously $O[[a, +\infty[, \tilde{a}']]$ and $O[a'', \tilde{a}']$ for the same O and \tilde{a}' , and it never contains simultaneously $O[[a, +\infty[, \tilde{a}']]$ and $O[[a'', +\infty[, \tilde{a}']]$ with $a \neq a''$ for the same O and \tilde{a}' .

Next, we define functions that manipulate these sets of oracles and roles. We define the subtraction operation $\mathcal{I} - O[\tilde{a}]$ on sets of oracles.

- If $O[[a, +\infty[, \tilde{a}']]$ is in \mathcal{I} , then

$$\mathcal{I} - (O[a, \tilde{a}']) \stackrel{\text{def}}{=} \mathcal{I} \setminus \{O[[a, +\infty[, \tilde{a}']]\} \cup \{O[[a+1, +\infty[, \tilde{a}']]\}.$$

- If $O[\tilde{a}]$ is in \mathcal{I} , then

$$\mathcal{I} - (O[\tilde{a}]) \stackrel{\text{def}}{=} \mathcal{I} \setminus \{O[\tilde{a}]\}.$$

We define similarly the subtraction on sets of roles $\mathcal{RI} - \text{role}[\tilde{a}]$. We also generalize this operator to sets:

$$\mathcal{RI} - \{\text{role}_1[\tilde{a}_1], \dots, \text{role}_k[\tilde{a}_k]\} \stackrel{\text{def}}{=} (\dots (\mathcal{RI} - \text{role}_1[\tilde{a}_1]) - \dots) - \text{role}_k[\tilde{a}_k].$$

We let $\text{smallest}(\mathcal{RI}, \text{role})$ be the smallest indices present for the role role in \mathcal{RI} : when $\tilde{a} = \text{smallest}(\mathcal{RI}, \text{role})$, we have $\text{role}[\tilde{a}] \in \mathcal{RI}$ or there exist a' and \tilde{a}' such that $\tilde{a} = a'$, \tilde{a}' and $\text{role}[[a', +\infty[, \tilde{a}']] \in \mathcal{RI}$.

Let us define the function $\text{oracles}'$, which is similar to the function reduce' , but just returns the oracle name and its replication indices \tilde{i} (which can be partly instantiated to values), instead of returning the entire oracle definition:

$$\text{oracles}'(0) \stackrel{\text{def}}{=} [] \quad (\text{Nil})$$

$$\text{oracles}'(Q_1 \mid Q_2) \stackrel{\text{def}}{=} \text{oracles}'(Q_1) @ \text{oracles}'(Q_2) \quad (\text{Par})$$

$$\begin{aligned} \text{oracles}'(\text{foreach } i' \leq n \text{ do } Q) &\stackrel{\text{def}}{=} [O_1[_, \tilde{i}], \dots, O_l[_, \tilde{i}]] \text{ when} \\ \text{oracles}'(Q) &= [O_1[i', \tilde{i}], \dots, O_l[i', \tilde{i}]] \end{aligned} \quad (\text{Repl})$$

$$\text{oracles}'(\text{role } \{Q\}) \stackrel{\text{def}}{=} [] \quad (\text{Role})$$

$$\text{oracles}'(O[\tilde{i}](x_1[\tilde{i}], \dots, x_k[\tilde{i}]) := P) \stackrel{\text{def}}{=} [O[\tilde{i}]] \quad (\text{Oracle})$$

This function returns elements of the form $O[\tilde{i}]$ for oracles that are not directly under replication and $O[_, \tilde{i}]$ for oracles directly under replication. Similarly to reduce' , this function returns an empty list when encountering a role definition.

Let us consider a process $Q' = \text{foreach } i' \leq n \text{ do } Q$. By Assumption 2.10, there is no replication in Q , and so all oracles in Q are under the same replications and have exactly the same replication indices i', \tilde{i} , where the indices \tilde{i} are the replication indices of replications above Q' . So, by rule (Repl), $\text{oracles}'(Q')$ produces the list of callable oracles in Q where we replace the replication index i' with $_$.

By Property 2.4, an oracle with a certain name O always takes arguments of the same types and always returns values of the same types. So we can say that the oracle O_i takes m_i arguments of types $T_{i,1}, \dots, T_{i,m_i}$, and returns m'_i bitstrings of types $T'_{i,1}, \dots, T'_{i,m'_i}$. We can also define $\text{returnoracles}(O[\tilde{i}]) \stackrel{\text{def}}{=} \text{oracles}'(Q)$ where Q is an oracle definition located after a return statement in a body of the oracle $O[\tilde{i}]$ in Q_0 . This definition is correct because, by Property 2.4, the structure of the processes Q after any return statement of a given oracle O is always the same, so the list $\text{oracles}'(Q)$ will be the same for each of these Q . The function returnoracles can take an oracle with its replication indices partly instantiated to values: $\text{returnoracles}(O[\tilde{a}]) \stackrel{\text{def}}{=} \text{returnoracles}(O[\tilde{i}])\{\tilde{a}/\tilde{i}\}$.

Let us denote by $Q(\text{role})$ the subprocess of Q_0 that corresponds to the role role . For a subprocess Q of Q_0 that is under replication indices \tilde{i} in Q_0 , we denote $Q[\tilde{a}]$ the process Q where we substituted elements of \tilde{i} by the respective elements of \tilde{a} .

Definition 6.12 (First oracle) *The first oracles of a role role are the oracles that can be called when we are at the beginning of the subprocess corresponding to the role, that is, $\text{oracles}'(Q(\text{role}))$.*

We define $\text{add}(\mathcal{I}, \mathcal{RI})$ as the addition of the first oracles present in \mathcal{RI} to \mathcal{I} :

$$\begin{aligned} \text{add}(\mathcal{I}, \mathcal{RI}) &\stackrel{\text{def}}{=} \mathcal{I} \cup \{O[\tilde{a}] \mid \text{role}[\tilde{a}] \in \mathcal{RI}, O[\tilde{a}] \in \text{oracles}'(Q(\text{role}))[\tilde{a}]\} \cup \\ &\quad \{O[[1, +\infty[\tilde{a}]] \mid \text{role}[\tilde{a}] \in \mathcal{RI}, O[_, \tilde{a}] \in \text{oracles}'(Q(\text{role}))[\tilde{a}]]\} \end{aligned}$$

The syntax of the language of the simulator is almost the same as the language we described in Section 3, with the addition of tagged functions introduced in Section 4. We add the functional values $\text{call}(O[\tilde{a}])$ and $\text{call}(O[_, \tilde{a}])$ that replace our generated closures for the oracle O . The value $\text{call}(O[\tilde{a}])$ is used when O is not directly under replication; $\text{call}(O[_, \tilde{a}])$ is used when O is directly under replication.

We present the semantics followed by our simulator in Figure 20. When we encounter a configuration containing a successful call to an oracle (by call) or a random operation, we cannot reduce. These operations are executed, but not inside the simulator: we stop the simulator in its current state, and in CryptoVerif, we call the requested oracle with the requested arguments, or generate a random bit. Otherwise, when the simulator configuration reduces into another configuration in the OCaml semantics, by rule (Simulator toplevel), we also reduce in the same way. By rules (FailedCall1) and (FailedCall2), we raise the exception `Bad_Call` when the call to the oracle is invalid, as our generated code does in this case. Notice that, in the OCaml implementation, the

$$\begin{array}{c}
O[\tilde{a}] \notin \mathcal{I} \\
\text{or } O \text{ does not have } k \text{ arguments} \\
\text{or } O \text{ has } k \text{ arguments of type } T_1, \dots, T_k \text{ and } \exists i, \exists a_i, v_i = \mathbb{G}_{\text{val}T_i}(a_i) \\
\hline
env, \text{call}(O[\tilde{a}]) (v_1, \dots, v_k), stack \rightarrow env, \text{raise Bad_Call}, stack \\
\text{(FailedCall1)} \\
\\
\forall a', O[[a', +\infty[\tilde{a}]] \notin \mathcal{I} \\
\text{or } O \text{ does not have } k \text{ arguments} \\
\text{or } O \text{ has } k \text{ arguments of type } T_1, \dots, T_k \text{ and } \exists i, \exists a_i, v_i = \mathbb{G}_{\text{val}T_i}(a_i) \\
\hline
env, \text{call}(O[_{\tilde{a}}]) (v_1, \dots, v_k), stack \rightarrow env, \text{raise Bad_Call}, stack \\
\text{(FailedCall2)} \\
\\
\mathcal{C} \rightarrow \mathcal{C}' \text{ using the rules of Figures 7–14, (FailedCall1), and (FailedCall2)} \\
\text{but not (Random) and (Toplevel add thread)} \\
\hline
\mathcal{C}, \mathcal{RI}, \mathcal{I} \rightarrow \mathcal{C}', \mathcal{RI}, \mathcal{I} \\
\text{(Simulator toplevel)} \\
\\
\text{program}^a = \text{program}_{\text{prim}};; \text{program}(\mu_{\text{role}_1});; \dots;; \text{program}(\mu_{\text{role}_l});; \text{program}' \\
\text{program}' \text{ does not contain } \text{program}(\mu_{\text{prim}}) \text{ nor any } \text{program}(\mu) \text{ for } \mu \in \mathcal{M}_g \\
\{\mu_{\text{role}_1}, \dots, \mu_{\text{role}_l}\} \subseteq \mathcal{M}_g \\
\tilde{a}_1 = \text{smallest}(\mathcal{RI}, \text{role}_1), \dots, \tilde{a}_l = \text{smallest}(\mathcal{RI}, \text{role}_l) \\
\mathcal{RI}'' = \{\text{role}_1[\tilde{a}_1], \dots, \text{role}_l[\tilde{a}_l]\} \quad \mathcal{RI}' = \mathcal{RI} - \mathcal{RI}'' \quad \mathcal{I}' = \text{add}(\mathcal{I}, \mathcal{RI}'') \\
\text{program}^b = \text{program}_{\text{prim}};; \text{program}'(\text{role}_1[\tilde{a}_1]);; \dots;; \text{program}'(\text{role}_l[\tilde{a}_l]);; \\
\text{program}' \\
\text{or} \\
\text{program}^a \text{ does not contain } \text{program}(\mu_{\text{prim}}) \text{ nor any } \text{program}(\mu) \text{ for } \mu \in \mathcal{M}_g \\
\mathcal{RI}'' = \emptyset \quad \mathcal{RI}' = \mathcal{RI} \quad \mathcal{I}' = \mathcal{I} \quad \text{program}^b = \text{program}^a \\
\hline
[Th_1, \dots, Th_{tj-1}, \langle env, \text{addthread}(\text{program}^a), stack, store \rangle, Th_{tj+1}, \dots, Th_n], \\
\text{globalstore}, tj, \mathcal{RI}, \mathcal{I} \longrightarrow \\
[Th_1, \dots, Th_{tj-1}, \langle env, (), stack, store \rangle, Th_{tj+1}, \dots, Th_n, \langle \emptyset, \text{program}^b, [], \emptyset \rangle], \\
\text{globalstore}, tj, \mathcal{RI}', \mathcal{I}' \\
\text{(Simulator add thread)}
\end{array}$$

Figure 20: Semantics followed by the simulator

adversary can test whether an oracle call succeeds or not, by catching the exception `Bad_Call`. In `CryptoVerif`, failed calls can happen only when the called oracle is not available, and in this case, the reduction blocks. This different behavior does not give additional power to the OCaml adversary, because the adversary can test before performing the call whether it will succeed or not. The rules (`FailedCall1`) and (`FailedCall2`) implement this test. By rule (`Simulator add thread`), we modify the behavior of the `addthread` construct to transform references to our generated modules $program(\mu_{\text{role}})$ into references to the corresponding role $program'(\text{role}[\tilde{a}])$ where \tilde{a} are the replication indices we chose for this particular reference and

$$\begin{aligned}
program'(\text{role}[\tilde{a}]) &\stackrel{\text{def}}{=} \text{let } \mu_{\text{role}}.\text{init} = \text{let } token = \text{ref true} \text{ in tagfunction}^{\text{role}} pm'_{\text{role}[\tilde{a}]} \\
\text{where } pm'_{\text{role}[\tilde{a}]} &\stackrel{\text{def}}{=} () \rightarrow \\
&\quad \text{if } (!token) \text{ then } (token := \text{false}; (\text{call}(O_1[\tilde{a}_1]), \dots, \text{call}(O_k[\tilde{a}_k]))) \\
&\quad \text{else raise Bad_Call}
\end{aligned}$$

where $oracles'(Q(\text{role})[\tilde{a}]) = [O_1[\tilde{a}_1], \dots, O_k[\tilde{a}_k]]$, and the \tilde{a}_j are either \tilde{a} or $_$, \tilde{a} . In particular, the initialization function defined in $program'(\text{role}[\tilde{a}])$ returns oracles represented by call values instead of closures.

The `CryptoVerif` function $simulate : T_{CS} \rightarrow \text{bitstring}$ follows the simulator semantics defined in Figure 20: we define $simulate(\text{repr}(\mathcal{CS})) \stackrel{\text{def}}{=} \text{simreturn}(\mathcal{CS}')$ where \mathcal{CS}' is the configuration such that either \mathcal{CS} reduces into \mathcal{CS}' in at most N_{steps} reductions and \mathcal{CS}' does not reduce, or \mathcal{CS} reduces into \mathcal{CS}' in exactly N_{steps} reductions, by the semantics of Figure 20, and $\text{simreturn}(\mathcal{CS}')$ is defined below. (We need to bound the number of reductions to make sure that $simulate$ is always defined. The proof of the simulation between OCaml and `CryptoVerif`, presented in the next section, shows that the simulator configuration always blocks after at most N_{steps} reductions, so that we are always in the first case.)

- If $\mathcal{C}_{pe}(\mathcal{CS}') = \text{call}(O[\tilde{a}]) (v_1, \dots, v_l)$, let T_1, \dots, T_l be the type of the arguments of the oracle O and let o be the constant associated to O . We define

$$\text{simreturn}(\mathcal{CS}') \stackrel{\text{def}}{=} (\text{repr}(\mathcal{CS}'), o, \tilde{a}, (\mathbb{G}_{\text{val}T_1}^{-1}(v_1), \dots, \mathbb{G}_{\text{val}T_l}^{-1}(v_l))).$$

- If $\mathcal{C}_{pe}(\mathcal{CS}') = \text{call}(O[_ , \tilde{a}]) (v_1, \dots, v_l)$, let T_1, \dots, T_l be the type of the arguments of the oracle O , let o be the constant associated to O , and let a' be the value such that $O[[a', +\infty, \tilde{a}]]$ is in the set \mathcal{I} where $\mathcal{CS}' = \mathcal{C}, \mathcal{RI}, \mathcal{I}$. We define

$$\text{simreturn}(\mathcal{CS}') \stackrel{\text{def}}{=} (\text{repr}(\mathcal{CS}'), o, (a', \tilde{a}), (\mathbb{G}_{\text{val}T_1}^{-1}(v_1), \dots, \mathbb{G}_{\text{val}T_l}^{-1}(v_l))).$$

- If $\mathcal{C}_{pe}(\mathcal{CS}') = \text{random } ()$, we define

$$\text{simreturn}(\mathcal{CS}') \stackrel{\text{def}}{=} (\text{repr}(\mathcal{CS}'), o_R, (), ()).$$

- Otherwise, we define

$$\text{simreturn}(\mathcal{CS}') \stackrel{\text{def}}{=} (\text{repr}(\mathcal{CS}'), o_S, (), ()).$$

The function *simulate* can be implemented by a deterministic Turing machine (since the random choices are handled outside *simulate*), so it can be used as a CryptoVerif primitive.

When *simulate* returns $(\text{repr}(\mathcal{CS}'), o, \tilde{a}, (a_1, \dots, a_l))$, the CryptoVerif process $Q_c(Q_0, \text{program}_0)$ performs the corresponding oracle call $O[\tilde{a}](a_1, \dots, a_l)$ (lines 10–17 of Figure 19). Similarly, when *simulate* returns $(\text{repr}(\mathcal{CS}'), o_R, (), ())$, the process $Q_c(Q_0, \text{program}_0)$ performs a random choice (lines 18–20), and when *simulate* returns $(\text{repr}(\mathcal{CS}'), o_S, (), ())$, the process $Q_c(Q_0, \text{program}_0)$ terminates (lines 8–9; the corresponding OCaml program also terminates).

The functions $\text{simulate}'_O$ and $\text{simulate}''_O$ replace, in the simulator configuration, the call expression with the result returned by the oracle, and raise the `Match_failure` exception, respectively. The function $\text{simulate}'_O$ handles the situation in which an oracle returns a result by `return`; the function $\text{simulate}''_O$ handles the situation in which the oracle terminates with `end`. Formally, these functions are defined as follows.

Definition 6.13 (Simulation of oracle return) *Let us consider a simulator configuration $\mathcal{CS} = \mathcal{C}, \mathcal{RI}, \mathcal{I}$, with*

$$\mathcal{C}_{pe}(\mathcal{CS}) = \text{call}(O[\tilde{a}]) (v_1, \dots, v_l) \text{ or } \text{call}(O[_], \tilde{a}') (v_1, \dots, v_l).$$

When $\mathcal{C}_{pe}(\mathcal{CS})$ is of the second form, we denote by \tilde{a} the indices a'', \tilde{a}' where a'' is such that $O[[a'', +\infty[\tilde{a}'] \in \mathcal{I}$. Let $\mathcal{I}' \stackrel{\text{def}}{=} \mathcal{I} - (O[\tilde{a}])$.

We define the CryptoVerif function $\text{simulate}'_O : T_{\mathcal{CS}} \times \text{bitstring} \rightarrow T_{\mathcal{CS}}$ as follows.

If the returns in oracle O end the current role, then by Property 2.7, there is only one `return` statement in O ; let Q be the oracle definition following this statement, and let

$$\begin{aligned} \mathcal{RI}' \stackrel{\text{def}}{=} & \{ \text{role}[\tilde{a}] \mid (\mu_{\text{role}}, \text{false}) \in \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q) \} \\ & \cup \{ \text{role}[[1, +\infty[\tilde{a}]] \mid (\mu_{\text{role}}, \text{true}) \in \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q) \}. \end{aligned}$$

Let T_1, \dots, T_n be the types of the return value of O . We define:

$$\text{simulate}'_O(\text{repr}(\mathcal{C}, \mathcal{RI}, \mathcal{I}), (r_1, \dots, r_n)) \stackrel{\text{def}}{=} \text{repr}(\mathcal{C}', \mathcal{RI} \cup \mathcal{RI}', \mathcal{I}')$$

where \mathcal{C}' is the configuration \mathcal{C} in which the current expression is replaced with the translated result: $(\mathbb{G}_{\text{val}_{T_1}}(r_1), \dots, \mathbb{G}_{\text{val}_{T_n}}(r_n))$.

If the returns in oracle O do not end the current role, then let us define $\mathcal{O} \stackrel{\text{def}}{=} \text{returnoracles}(O[\tilde{a}])$. Let \mathcal{I}'' be the set \mathcal{I}' to which we added the oracles present in \mathcal{O} :

$$\mathcal{I}'' \stackrel{\text{def}}{=} \mathcal{I}' \cup \{ O'[[1, +\infty[\tilde{a}]] \mid O'[_], \tilde{a}] \in \mathcal{O} \} \cup \{ O'[\tilde{a}] \mid O'[\tilde{a}] \in \mathcal{O} \}.$$

We define:

$$\text{simulate}'_{\mathcal{O}}(\text{repr}(\mathcal{C}, \mathcal{R}\mathcal{I}, \mathcal{I}), (r_1, \dots, r_n)) \stackrel{\text{def}}{=} \text{repr}(\mathcal{C}', \mathcal{R}\mathcal{I}, \mathcal{I}')$$

where \mathcal{C}' is the configuration \mathcal{C} in which the current expression is replaced with the translated result: $(\text{call}(O_1[\tilde{a}_1]), \dots, \text{call}(O_l[\tilde{a}_l]), \mathbb{G}_{\text{val}T_1}(r_1), \dots, \mathbb{G}_{\text{val}T_n}(r_n))$, with $\mathcal{O} = \{O_1[\tilde{a}_1], \dots, O_l[\tilde{a}_l]\}$ and the \tilde{a}_j are either \tilde{a} or $_ , \tilde{a}$.

In all other cases (that is, $\mathcal{C}\mathcal{S}$ is not of the form mentioned above or the bitstring a is not a tuple of n bitstrings of types T_1, \dots, T_n), $\text{simulate}'_{\mathcal{O}}(\text{repr}(\mathcal{C}\mathcal{S}), a)$ can take any value, since these cases are in fact not used.

Finally, we define the *CryptoVerif* function $\text{simulate}''_{\mathcal{O}} : T_{\mathcal{C}\mathcal{S}} \rightarrow T_{\mathcal{C}\mathcal{S}}$ by:

$$\text{simulate}''_{\mathcal{O}}(\text{repr}(\mathcal{C}, \mathcal{R}\mathcal{I}, \mathcal{I})) \stackrel{\text{def}}{=} \text{repr}(\mathcal{C}'', \mathcal{R}\mathcal{I}, \mathcal{I}')$$

where \mathcal{C}'' is the configuration \mathcal{C} in which the current expression is replaced with `raise Match_failure`. In all other cases (that is, $\mathcal{C}\mathcal{S}$ is not of the form mentioned above), $\text{simulate}''_{\mathcal{O}}(\text{repr}(\mathcal{C}\mathcal{S}))$ can take any value, since these cases are in fact not used.

When the returns in oracle \mathcal{O} end the current role, the function $\text{simulate}'_{\mathcal{O}}$ does not return the oracles following the current oracle, but adds the corresponding roles to the role set $\mathcal{R}\mathcal{I}$. The programs that contain these roles can then be launched by `addthread`.

Definition 6.14 (Random simulation) We define the *CryptoVerif* function $\text{simulate}_R : T_{\mathcal{C}\mathcal{S}} \times \text{bool} \rightarrow T_{\mathcal{C}\mathcal{S}}$ by

$$\text{simulate}_R(\text{repr}(\mathcal{C}, \mathcal{R}\mathcal{I}, \mathcal{I}), b) \stackrel{\text{def}}{=} \text{repr}(\mathcal{C}'(b), \mathcal{R}\mathcal{I}, \mathcal{I})$$

where $\mathcal{C}'(b)$ is the configuration \mathcal{C} in which the current expression is replaced with the OCaml boolean value $\mathbb{G}_{\text{valbool}}(b)$.

Let us finally define the initial state of the simulator. Let $\mathcal{R}\mathcal{I}_0$ be the set of initially callable roles of Q_0 with their replication indices: $\mathcal{R}\mathcal{I}_0 \stackrel{\text{def}}{=} \{\text{role}[] \mid (\mu_{\text{role}}, \text{false}) \in \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q_0)\} \cup \{\text{role}[1, +\infty[] \mid (\mu_{\text{role}}, \text{true}) \in \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q_0)\}$. We define:

$$s_0(Q_0, \text{program}_0) \stackrel{\text{def}}{=} \text{repr}(\langle [\emptyset, \text{program}_0, [], \emptyset], \text{globalstore}_0, 1, \mathcal{R}\mathcal{I}_0, \emptyset \rangle)$$

6.4 Correspondence between the *CryptoVerif* and OCaml Systems

In this section, we prove our main security theorem by relating the *CryptoVerif* and OCaml systems.

In *CryptoVerif*, the security properties are defined using distinguishers D which are functions that take a list of events \mathcal{E} and return true or false. We denote by $\text{Pr}[\mathfrak{C} : D]$ the probability of the set of complete *CryptoVerif* traces starting at \mathfrak{C} and such that the list of events \mathcal{E} in their last configuration satisfies $D(\mathcal{E}) = \text{true}$.

For instance, to show that a protocol Q_0 satisfies a correspondence c of the form “for all a , if $e_1(a)$ has been executed, then $e_2(a)$ has also been executed”, we define D_c by $D_c(\mathcal{E}) = \text{true}$ if and only if the correspondence does not hold, that is, \mathcal{E} contains $e_1(a)$ but not $e_2(a)$ for some a . We can represent the adversary for Q_0 by any CryptoVerif process Q_{adv} that does not contain events nor variables that occur in Q_0 . Then we bound the probability $\Pr[\mathfrak{C}_i(Q_0 \mid Q_{\text{adv}}) : D_c]$, that is, the probability that the adversary Q_{adv} breaks the correspondence in Q_0 , for any adversary Q_{adv} for Q_0 . We can also define secrecy using events and distinguishers [6].

We use a similar definition in OCaml: $\Pr[\mathcal{C} : D]$ is the probability of the set of complete OCaml traces starting at \mathcal{C} and such that the list of events *events* in their last configuration satisfies $D(\mathbb{G}_{\text{ev}}^{-1}(\text{events})) = \text{true}$.

Our goal is to prove that, for all protocols Q_0 , OCaml adversaries program_0 , and distinguishers D , we have $\Pr[\mathfrak{C}_0(Q_0, \text{program}_0) : D] = \Pr[\mathcal{C}_0(Q_0, \text{program}_0) : D]$. From this property, it is easy to see that, if CryptoVerif bounds the probability $\Pr[\mathfrak{C}_i(Q_0 \mid Q_{\text{adv}}) : D]$ for any adversary Q_{adv} for Q_0 , then the same bound also holds for the probability $\Pr[\mathcal{C}_0(Q_0, \text{program}_0) : D]$ corresponding to the generated implementation. Indeed, $\Pr[\mathcal{C}_0(Q_0, \text{program}_0) : D] = \Pr[\mathfrak{C}_0(Q_0, \text{program}_0) : D] = \Pr[\mathfrak{C}_i(Q_0 \mid Q_{\text{adv}}(Q_0, \text{program}_0)) : D]$ and $Q_{\text{adv}}(Q_0, \text{program}_0)$ is an adversary for Q_0 .

To that order, we first introduce an intermediate semantics for CryptoVerif that decomposes the evaluation of the function *simulate* into several small steps. We easily relate this semantics to the semantics of CryptoVerif. Next, in Section 6.4.2, we relate the intermediate semantics to the OCaml semantics. For this purpose, we introduce a relation between intermediate semantic configurations and OCaml traces, that, in particular, ensures that the events are the same on both sides and we prove that this relation is preserved by reduction. Finally, in Section 6.4.3, we use these results to prove our main theorem.

6.4.1 Intermediate Semantics

We introduce extended CryptoVerif configurations \mathfrak{C}^{cs} , which are configurations of the form \mathfrak{C} or $\mathfrak{C}, \text{steps}, \mathcal{CS}$, where \mathcal{CS} is a simulator configuration and *steps* is the maximum number of reductions of \mathcal{CS} that can still be performed. (We use the field *steps* to guarantee termination.) The configurations $\mathfrak{C}, \text{steps}, \mathcal{CS}$ serve to represent the state of the system during the evaluation of the function *simulate*. We define a reduction relation \rightsquigarrow on the extended configurations \mathfrak{C}^{cs} .

Definition 6.15 *Let us define the reduction relation \rightsquigarrow such that:*

$$\frac{E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_p \mathfrak{C}' \quad P \text{ is not of the form } x[a] \leftarrow \text{simulate}(s[a]); P' \text{ for any } x, a, P'}{E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightsquigarrow_p \mathfrak{C}'} \quad (\text{CryptoVerif})$$

$$\begin{array}{c}
\frac{E(s[a]) = \text{repr}(\mathcal{CS})}{E, x[a] \leftarrow \text{simulate}(s[a]); P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightsquigarrow} \quad (\text{Enter Simulator}) \\
E, x[a] \leftarrow \text{simulate}(s[a]); P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, N_{\text{steps}}, \mathcal{CS} \\
\hline
\frac{\mathcal{CS} \rightarrow \mathcal{CS}' \quad \text{steps} > 0}{E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, \text{steps}, \mathcal{CS} \rightsquigarrow E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, \text{steps} - 1, \mathcal{CS}'} \quad (\text{Simulator}) \\
\hline
\frac{\mathcal{CS} \text{ does not reduce or } \text{steps} = 0}{E, x[a] \leftarrow \text{simulate}(s[a]); P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, \text{steps}, \mathcal{CS} \rightsquigarrow} \quad (\text{Leave Simulator}) \\
E[x[a] \mapsto \text{simreturn}(\mathcal{CS})], P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}
\end{array}$$

When encountering a configuration $\mathfrak{C} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$ such that P is of the form $x[a] \leftarrow \text{simulate}(s[a]); P'$ and $E(s[a]) = \text{repr}(\mathcal{CS})$, we reduce \mathfrak{C} into an extended configuration $\mathfrak{C}, N_{\text{steps}}, \mathcal{CS}$ by (Enter Simulator). We reduce \mathcal{CS} by (Simulator) until it blocks or the number of allowed reductions N_{steps} is exhausted, and then we resume the CryptoVerif reductions by (Leave Simulator).

In the next lemma and proposition, we relate traces using \rightsquigarrow to traces using \rightarrow , to prove that all events have the same probability in these two semantics. These results are proved in Appendix E.

Lemma 6.16 *Let \mathfrak{C} be a CryptoVerif configuration.*

- If $\mathfrak{C} \rightarrow_p \mathfrak{C}'$, then there is a trace $\mathfrak{C} \rightsquigarrow_p^* \mathfrak{C}'$ and all intermediate configurations in this trace (if any) are of the form $\mathfrak{C}, \text{steps}, \mathcal{CS}$.
- If \mathfrak{C} does not reduce by \rightarrow , then it does not reduce by \rightsquigarrow either.

We denote by $\Pr[\mathfrak{C}^{\text{cs}}(\rightsquigarrow) : D]$ the probability of the set of complete CryptoVerif traces using \rightsquigarrow starting at \mathfrak{C}^{cs} and such that the list of events \mathcal{E} in their last configuration satisfies $D(\mathcal{E}) = \text{true}$. The next proposition shows that all events have the same probability in the intermediate semantics as in the CryptoVerif semantics.

Proposition 6.17 $\Pr[\mathfrak{C}(\rightsquigarrow) : D] = \Pr[\mathfrak{C} : D]$.

6.4.2 Relation between the Intermediate Semantics and the OCaml Semantics

Let us first give some preliminary definitions.

Definition 6.18 (Accessible trace, configuration) *An accessible trace \mathcal{CT} (for the adversary program₀) is a trace beginning with the initial configuration $\mathcal{C}_0(Q_0, \text{program}_0)$ defined in Section 5.*

A configuration \mathcal{C} is accessible if there exists an accessible trace such that its last configuration is \mathcal{C} .

Definition 6.19 (Concretization of \mathcal{I} and \mathcal{RI}) Let us define the sets of oracles $\mathcal{O}^\infty(\mathcal{I})$ and $\mathcal{O}^\infty(\mathcal{RI})$ represented by \mathcal{I} and \mathcal{RI} respectively:

$$\begin{aligned}\mathcal{O}^\infty(\mathcal{I}) &= \{O[b, \tilde{a}'] \mid O[[a, +\infty[, \tilde{a}'] \in \mathcal{I}, a \leq b\} \cup \{O[\tilde{a}] \mid O[\tilde{a}] \in \mathcal{I}\} \\ \mathcal{O}^\infty(\mathcal{RI}) &= \{O[b, \tilde{a}] \mid \text{role}[\tilde{a}] \in \mathcal{RI}, O[_-, \tilde{a}] \in \text{oracles}'(Q(\text{role})[\tilde{a}]), 1 \leq b\} \\ &\quad \cup \{O[\tilde{a}] \mid \text{role}[\tilde{a}] \in \mathcal{RI}, O[\tilde{a}] \in \text{oracles}'(Q(\text{role})[\tilde{a}])\} \\ &\quad \cup \{O[b, \tilde{a}'] \mid \text{role}[[a, +\infty[, \tilde{a}'] \in \mathcal{RI}, \\ &\quad \quad O[b, \tilde{a}'] \in \text{oracles}'(Q(\text{role})[b, \tilde{a}']), a \leq b\}\end{aligned}$$

The definition of $\mathcal{O}^\infty(\mathcal{I})$ and $\mathcal{O}^\infty(\mathcal{RI})$ ignores the replication bounds and allows the indices of oracles to go to infinity. Using unbounded indices is helpful in Properties 13 and 14 of Definition 6.27 below. By Assumption 2.10, when O is a first oracle of a role role under replication, O cannot be under replication in $Q(\text{role})$. So the last component of $\mathcal{O}^\infty(\mathcal{RI})$ cannot contain oracles under replication.

Let us describe how the sets \mathcal{I} and \mathcal{RI} represent the contents of the set of callable processes \mathcal{Q} .

Definition 6.20 (Relation between \mathcal{I} , \mathcal{RI} and \mathcal{Q}) Let us define the sets of oracles $\mathcal{O}(\mathcal{I})$ and $\mathcal{O}(\mathcal{RI})$ represented by \mathcal{I} and \mathcal{RI} respectively:

$$\begin{aligned}\mathcal{O}(\mathcal{I}) &= \{O[b, \tilde{a}'] \mid O[[a, +\infty[, \tilde{a}'] \in \mathcal{I}, a \leq b \leq N_O\} \cup \{O[\tilde{a}] \mid O[\tilde{a}] \in \mathcal{I}\} \\ \mathcal{O}(\mathcal{RI}) &= \{O[b, \tilde{a}] \mid \text{role}[\tilde{a}] \in \mathcal{RI}, \\ &\quad O[_-, \tilde{a}] \in \text{oracles}'(Q(\text{role})[\tilde{a}]), 1 \leq b \leq N_O\} \\ &\quad \cup \{O[\tilde{a}] \mid \text{role}[\tilde{a}] \in \mathcal{RI}, O[\tilde{a}] \in \text{oracles}'(Q(\text{role})[\tilde{a}])\} \\ &\quad \cup \{O[b, \tilde{a}'] \mid \text{role}[[a, +\infty[, \tilde{a}'] \in \mathcal{RI}, \\ &\quad \quad O[b, \tilde{a}'] \in \text{oracles}'(Q(\text{role})[b, \tilde{a}']), a \leq b \leq N_{\text{role}}\}\end{aligned}$$

We write $\mathcal{Q} \leftrightarrow \mathcal{RI}, \mathcal{I}$ when the following two properties hold:

- \mathcal{Q} consists of exactly one element $O[\tilde{a}](x_1[\tilde{a}] : T_1, \dots, x_k[\tilde{a}] : T_k) := P$ for each $O[\tilde{a}]$ present in the set $\mathcal{O}(\mathcal{I}) \cup \mathcal{O}(\mathcal{RI})$. We denote by $\mathcal{Q}(O[\tilde{a}])$ this element of \mathcal{Q} .
- If $O[[a, +\infty[, \tilde{a}'] \in \mathcal{I}$, then there exist a process Q and an index i such that i does not occur in $\text{fv}(Q)$ and for all $b \in \{a, \dots, N_O\}$, we have $\mathcal{Q}(O[b, \tilde{a}']) = Q\{b/i\}$.

In contrast to the sets we defined in Definition 6.19, the indices of oracles in \mathcal{Q} are bounded by the replication bounds. So we redefine sets of oracles $\mathcal{O}(\mathcal{RI})$ and $\mathcal{O}(\mathcal{I})$ that correspond to \mathcal{RI} and \mathcal{I} , but with indices bounded by N_O and N_{role} as appropriate. The sets $\mathcal{O}(\mathcal{RI})$ and $\mathcal{O}(\mathcal{I})$ are included in $\mathcal{O}^\infty(\mathcal{RI})$ and $\mathcal{O}^\infty(\mathcal{I})$, respectively. The set of processes \mathcal{Q} corresponds to $\mathcal{RI}, \mathcal{I}$ when it contains exactly one definition for each oracle in $\mathcal{O}(\mathcal{I}) \cup \mathcal{O}(\mathcal{RI})$. Furthermore, in case an oracle is under replication, the corresponding elements of \mathcal{Q} all have the same form; they differ only by the value of the replication index. We enforce this property in the last item of Definition 6.20.

Next, we define several sets of oracles and roles, which allow us to determine which oracles and roles are in which state (callable immediately, available later) in a simulator configuration.

Definition 6.21 (Oracle sets) Let $\mathcal{O}_{\text{call}}(Th)$ be the set of oracles $O[\tilde{a}]$ not under replication that occur in call constructs in the thread Th , without entering tagged functions and closures.

Let $\mathcal{O}_{\text{call-repl}}(Th)$ be the set of oracles $O[a, \tilde{a}]$ such that O is under replication, $a > N_O$, and $\text{call}(O[_{_}, \tilde{a}])$ occurs in the thread Th , without entering tagged functions and closures.

Let $\mathcal{R}_{\text{init-closure}}(Th)$ be the set of roles $\text{role}[\tilde{a}]$ such that there exists env such that a closure $\text{tagfunction}^{\text{role}, \tau}[\text{env}, \text{pm}'_{\text{role}[\tilde{a}]}]$ is present in the thread Th , and such that $\text{env}(\text{token})$ is bound in its store to true.

Let $\mathcal{R}_{\text{init-function}}(Th)$ be the set of roles $\text{role}[\tilde{a}]$ such that the initialization function program'($\text{role}[\tilde{a}]$) is present in the thread Th .

Let $\mathcal{O}_{\text{call}}(\mathcal{CS})$, $\mathcal{R}_{\text{init-closure}}(\mathcal{CS})$, and $\mathcal{R}_{\text{init-function}}(\mathcal{CS})$ be the unions of the corresponding sets for all threads of the configuration.

We have already defined the set $\text{returnoracles}(O[\tilde{a}])$, which is the set of oracles returned by $O[\tilde{a}]$. Let us define $\text{returnoracles}'$ such that:

$$\begin{aligned} \text{returnoracles}'(O[\tilde{a}]) &= \{O'[\tilde{a}'] \mid O'[\tilde{a}'] \in \text{returnoracles}(O[\tilde{a}])\} \\ &\quad \cup \{O'[b, \tilde{a}'] \mid O'[_{_}, \tilde{a}'] \in \text{returnoracles}(O[\tilde{a}]), 1 \leq b \leq N_{O'}\} \end{aligned}$$

Let $\mathcal{CS} = \mathcal{C}, \mathcal{RI}, \mathcal{I}$. We denote the callable set of oracles:

$$\text{callable}(\mathcal{CS}) = \mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}^\infty(\mathcal{RI}) \cup \mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(\mathcal{CS}) \cup \mathcal{R}_{\text{init-function}}(\mathcal{CS}))$$

Let $\text{willbeavailable}(\mathcal{CS})$ be the set of oracles that can eventually become available. This set is the least fixpoint of the function f defined by $f(C) = C \cup \text{returnoracles}'(C)$ that contains the set $\text{returnoracles}'(\text{callable}(\mathcal{CS}))$.

The definition of $\mathcal{O}_{\text{call-repl}}(Th)$ may be surprising, as it considers $O[a, \tilde{a}]$ with a greater than the replication bound N_O . We have made this choice to guarantee that $\mathcal{O}_{\text{call-repl}}(Th)$ is always included in $\mathcal{O}^\infty(\mathcal{I})$: the indices up to N_O may have been consumed by calls already made to the oracle, while the indices greater than N_O always remain, because we make at most N_O calls to this oracle by definition of N_O . This property is exploited in Property 14 of Definition 6.27 below.

The sets $\mathcal{R}_{\text{init-closure}}(\mathcal{CS})$ and $\mathcal{R}_{\text{init-function}}(\mathcal{CS})$ are sets of roles with their replication indices, which can be seen as a role set \mathcal{RI} . The set $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(\mathcal{CS}) \cup \mathcal{R}_{\text{init-function}}(\mathcal{CS}))$ is the set of the first oracles of roles present in $\mathcal{R}_{\text{init-closure}}(\mathcal{CS})$ and $\mathcal{R}_{\text{init-function}}(\mathcal{CS})$.

The following definitions present functions used to relate OCaml and simulator threads.

Definition 6.22 (Replace initialization) The function `replaceinitpm` replaces in its argument the pattern matchings corresponding to role initialization of the simulator by the OCaml module initialization: to be more precise,

$\text{replaceinitpm}(Th)$ replaces each occurrence of $\text{tagfunction}^{\text{role}} pm'_{\text{role}[\tilde{a}]}$ in Th with $\text{tagfunction}^{\text{role}} pm_{\mu_{\text{role}}}$ and each occurrence of $\text{tagfunction}^{\text{role},\tau}[env, pm'_{\text{role}[\tilde{a}]}]$ in Th with $\text{tagfunction}^{\text{role},\tau}[env, pm_{\mu_{\text{role}}}]$.

This function transforms every occurrence of the tagged closures corresponding to role initialization in the simulator, which are added by the `addthread` construct, into the corresponding tagged closures in OCaml.

Definition 6.23 (Correct closure) Suppose that $\mathcal{Q} \leftrightarrow \mathcal{RI}, \mathcal{I}$ for some \mathcal{RI} , l_{tok} is a function that maps each oracle $O[\tilde{a}]$ to the location of its token, and $\tau_{\mathcal{O}}$ is a function maps each oracle $O[_, \tilde{a}]$ to the tag τ of the corresponding closure. We define the set of closures that correspond to an oracle:

- for an oracle $O[\tilde{a}] \in \mathcal{I}$:

$$\begin{aligned} \text{correctclosure}(O[\tilde{a}], \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_{\mathcal{O}}) = \\ \{ \text{tagfunction}^{O,\tau}[env, pm_{\text{false}}(\mathcal{Q}(O[\tilde{a}]))] \mid \\ env \supseteq env_{\text{prim}} \cup env(E, \mathcal{Q}(O[\tilde{a}])), env(\text{token}) = l_{\text{tok}}(O[\tilde{a}]) \} \end{aligned}$$

- for an oracle $O[\tilde{a}] \notin \mathcal{I}$:

$$\begin{aligned} \text{correctclosure}(O[\tilde{a}], \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_{\mathcal{O}}) = \\ \{ \text{tagfunction}^{O,\tau}[env, pm_{\text{false}}(Q)] \mid \text{for any } Q, env(\text{token}) = l_{\text{tok}}(O[\tilde{a}]) \} \end{aligned}$$

- for an oracle $O[[a', +\infty[, \tilde{a}'']] \in \mathcal{I}$ with $a' \leq N_{\mathcal{O}}$,

$$\begin{aligned} \text{correctclosure}(O[_, \tilde{a}''], \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_{\mathcal{O}}) = \\ \{ \text{tagfunction}^{O,\tau}[env, pm_{\text{true}}(\mathcal{Q}(O[a', \tilde{a}'']))] \mid \\ \tau = \tau_{\mathcal{O}}(O[_, \tilde{a}'']), env \supseteq env_{\text{prim}} \cup env(E, \mathcal{Q}(O[a', \tilde{a}''])) \} \end{aligned}$$

- for an oracle $O[[a', +\infty[, \tilde{a}'']] \in \mathcal{I}$ with $a' > N_{\mathcal{O}}$,

$$\begin{aligned} \text{correctclosure}(O[_, \tilde{a}''], \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_{\mathcal{O}}) = \\ \{ \text{tagfunction}^{O,\tau}[env, pm_{\text{true}}(Q)] \mid \tau = \tau_{\mathcal{O}}(O[_, \tilde{a}'']), \text{for any } Q, env \} \end{aligned}$$

- for an oracle $O[[a', +\infty[, \tilde{a}'']] \notin \mathcal{I}$:

$$\text{correctclosure}(O[_, \tilde{a}''], \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_{\mathcal{O}}) = \emptyset$$

The function `correctclosure` serves to map calls `call(R)` in the simulator configuration into their corresponding closures in the OCaml configuration: `call(R)` will be mapped below to an element of `correctclosure(R, \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_{\mathcal{O}})`.

In the case $O[\tilde{a}] \in \mathcal{I}$, we map `call(O[\tilde{a}])` into the closure that translates the process $\mathcal{Q}(O[\tilde{a}])$.

The case $O[\tilde{a}] \notin \mathcal{I}$ may be used when the oracle $O[\tilde{a}]$ has been called but the thread still contains a call to this oracle. If the oracle is called again, the call will fail. The process $\mathcal{Q}(O[\tilde{a}])$ is removed from \mathcal{Q} after execution, so we do not know which process to translate to obtain the correct closure for $O[\tilde{a}]$, that is why the correct closures for a call to an already called oracle can contain the translation of any process Q . This translation will fail and raise the exception `Bad_Call` regardless of the translated process Q .

Oracles under replication cannot disappear from \mathcal{I} after having been added to it: when one calls the oracle $O[_ , \tilde{a}'']$, we just increment the counter a' of the element $O[[a', +\infty[, \tilde{a}'']]$ present in \mathcal{I} . We need to distinguish whether the adversary has exhausted all the N_O calls available for this oracle or not. If there remains available calls, the process $\mathcal{Q}(O[a', \tilde{a}''])$ is defined, and we require that $\text{call}(O[_ , \tilde{a}''])$ is mapped into a closure that translates this process. Otherwise, if all the calls are exhausted, $a' > N_O$, and $\mathcal{Q}(O[a', \tilde{a}''])$ is not defined, but we know that the adversary will not call the oracle again, so $\text{call}(O[_ , \tilde{a}''])$ can be mapped to closures that translate any process.

The case $O[[a', +\infty[, \tilde{a}'']] \notin \mathcal{I}$ never happens: it would mean that the oracle $O[_ , \tilde{a}'']$ can be called but there is no reference to it in the set \mathcal{I} .

Definition 6.24 (Replace call)

$\text{replacecalls}((env, pe, stack, store), \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_O) =$
 $\{ (env', \sigma(pe), \sigma(stack), \sigma(store)) \mid \text{if } pe \text{ is a value } v \text{ or an exceptional}$
 $\text{value raise } v, \text{ then } env' \text{ is any environment, else } env' = \sigma(env), \text{ where } \sigma$
 $\text{is a function that replaces, for each } R, \text{ each occurrence of } \text{call}(R) \text{ with an}$
 $\text{element of } \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_O) \}$

The function `replacecalls` replaces in its argument each call $\text{call}(O[\tilde{a}])$ with a closure that corresponds to the oracle $O[\tilde{a}]$, computed by `correctclosure`. It allows any environment when the current program or expression is a value or an exceptional value, because in these cases, the environment is not used.

Definition 6.25 (Token part of the store)

$$\begin{aligned} \text{gettokens}(\mathcal{I}, \mathcal{O}, l_{\text{tok}}) &= \{ l_{\text{tok}}(O[\tilde{a}]) \mapsto \text{true} \mid O[\tilde{a}] \in \mathcal{O} \cap \mathcal{I} \} \\ &\cup \{ l_{\text{tok}}(O[\tilde{a}]) \mapsto \text{false} \mid O[\tilde{a}] \in \mathcal{O} \setminus \mathcal{I} \} \end{aligned}$$

The function `gettokens` returns the part of the store corresponding to the tokens of the closures of oracles not under replication.

Definition 6.26 (Processes) *We use the following notations:*

P_{loop} is the process from line 7 to line 20 in Figure 19.

$Q_{loop} \stackrel{\text{def}}{=} O_{loop}[i'](s : T_{CS}) := P_{loop}$.

$P_{return-loop}(\alpha) \stackrel{\text{def}}{=} \text{if } b_{\alpha,r}[] \text{ then}$

let $r[] : T_{CS} = \text{loop } O_{loop}[\alpha + 1](r'_{\alpha,r}[]) \text{ in end else end}$

else $r[] \leftarrow r'_{\alpha,r}[]$; end.

$\mathcal{R}_{loop}(\alpha) \stackrel{\text{def}}{=} [(r'_{\alpha,r}[], b_{\alpha,r}[]), P_{return-loop}(\alpha), \text{end}], (x[], \text{return}(x[]), \text{end})]$.

Definition 6.27 (Relation between extended CryptoVerif configurations and OCaml traces) *Let \mathfrak{C}^{cs} be an extended CryptoVerif configuration and \mathcal{CT} be an accessible OCaml trace. We say that $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ when there exists an injective function $\tau_{\mathcal{O}}$ that maps oracles $O[_, \tilde{a}]$ such that $O[[a', +\infty], \tilde{a}] \in \mathcal{I}$ for some a' to tags τ , such that the following properties are all true:*

1. $\mathfrak{C}^{\text{cs}} = E, P_{loop}\{\alpha/i'\}, \mathcal{T}, \mathcal{Q}, \mathcal{R}_{loop}(\alpha), \mathcal{E}, \text{steps}, \mathcal{CS}$.
 $\mathcal{CS} = ([Th_1, \dots, Th_n], \text{globalstore}, tj), \mathcal{RI}, \mathcal{I}$.
For all $i \leq n$, $Th_i = \langle \text{env}_i, \text{pe}_i, \text{stack}_i, \text{store}_i \rangle$.
2. \mathcal{C} is the last configuration of \mathcal{CT} .
 $\mathcal{C} = [Th'_1, \dots, Th'_n], \text{globalstore}', tj, \mathcal{MI}, \text{events}$.
For all $i \leq n$, $Th'_i = \langle \text{env}'_i, \text{pe}'_i, \text{stack}'_i, \text{store}'_i \rangle$.
3. $\mathcal{Q} = \{Q_{loop}\{a/i'\} \mid \alpha < a \leq N_{\text{rand+calls}}\} \cup \mathcal{Q}_0$ and $\mathcal{Q}_0 \leftrightarrow \mathcal{RI}, \mathcal{I}$.
4. $\text{fv}(P_{loop}\{\alpha/i'\}) \cup \text{fv}(\mathcal{Q}) \cup \text{fv}(\mathcal{R}_{loop}(\alpha)) \subseteq \text{Dom}(E)$.
5. For $i \leq n$, all store locations in S_i present in Th_i are in $\text{Dom}(\text{store}_i)$.
6. For each thread $i \leq n$, one of the following two cases occurs:
 - (a) $Th'_i = \text{replaceinitpm}(Th_i)$.
 $Th_i = \langle \emptyset, \text{program}_{\text{prim}};; \text{program}'(\text{role}_1[\tilde{a}_1]);; \dots;; \text{program}'(\text{role}_l[\tilde{a}_l]);; \text{program}', [], \emptyset \rangle$.
There is no closure, no tagged function $\text{tagfunction}^t \text{ pm}$, no event, and no return in $\text{program}'$, except in $\text{program}(\mu_{\text{role}})$ in arguments of addthread .
 - (b) The following properties hold:
 - i. There exist store''_i and an injective function l_{tok} that associates to each $O[\tilde{a}]$ in $\mathcal{O}_{\text{call}}(Th_i)$ a store location that does not occur in \mathcal{CS} such that
$$\langle \text{env}'_i, \text{pe}'_i, \text{stack}'_i, \text{store}''_i \rangle$$

$$\in \text{replacecalls}(\text{replaceinitpm}(Th_i), \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_{\mathcal{O}}),$$

$$\text{store}''_i \cup \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th_i), l_{\text{tok}}) \subseteq \text{store}'_i.$$

- ii. There exists an injective function $l_{\text{init-tok}}$ that associates to each role $\text{role}[\tilde{a}]$ such that a closure $\text{tagfunction}^{\text{role}, \tau}[\text{env}, \text{pm}'_{\text{role}[\tilde{a}]}]$ occurs in the thread Th_i for some env and τ , a store location such that for all closures $\text{tagfunction}^{\text{role}, \tau}[\text{env}, \text{pm}'_{\text{role}[\tilde{a}]}]$ present in Th_i , we have $l_{\text{init-tok}}(\text{role}[\tilde{a}]) = \text{env}(\text{token})$.
 The locations $l_{\text{init-tok}}(\text{role}[\tilde{a}])$ and $l_{\text{tok}}(O[\tilde{a}'])$ are distinct for every role $\text{role}[\tilde{a}]$ and oracle $O[\tilde{a}']$.
 The locations $l_{\text{init-tok}}(\text{role}[\tilde{a}])$ occur only in $\text{Dom}(\text{store}_i)$ and in $\text{env}(\text{token})$ where env is the environment of a tagged closure $\text{tagfunction}^{\text{role}, \tau}[\text{env}, \text{pm}'_{\text{role}[\tilde{a}]}]$ in Th_i .
 - iii. For each tagged closure $\text{tagfunction}^{t, \tau}[\text{env}, \text{pm}]$ present in Th_i , the tag t is a role role , $\text{env}_{\text{prim}} \subseteq \text{env}$, and there exist indices \tilde{a} such that $\text{pm} = \text{pm}'_{\text{role}[\tilde{a}]}$.
 - iv. There is no tagged function $\text{tagfunction}^t \text{pm}$, no event, and no return in Th_i except in $\text{program}(\mu_{\text{role}})$ in arguments of addthread .
7. For all locations $l \in S_{\text{priv}}$, l does not occur in Th_1, \dots, Th_n except in $\text{program}(\mu_{\text{role}})$ in arguments of addthread .
 8. $\forall l \in S_{\text{priv}}, \text{globalstore}(l) = \text{initial}_l$.
 9. $\text{globalstore}(E, \mathcal{T}) \subseteq \text{globalstore}'$.
 10. $\forall l \notin S_{\text{priv}}, \text{globalstore}(l) = \text{globalstore}'(l)$.
 11. $\mathcal{MI} = \{(\mu_{\text{role}}, \text{false}) \mid \text{role}[\tilde{a}] \in \mathcal{RI}\} \cup \{(\mu_{\text{role}}, \text{true}) \mid \text{role}[[a', +\infty, \tilde{a}]] \in \mathcal{RI}\}$.
 12. $\text{events} = \mathbb{G}_{\text{ev}}(\mathcal{E})$.
 13. The sets $\mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{CS})$, $\mathcal{O}^\infty(\mathcal{RI})$, and $\text{willbeavailable}(\mathcal{CS})$ are pairwise disjoint.
 14. The $4n$ sets of oracles $\mathcal{O}_{\text{call}}(Th_i)$, $\mathcal{O}_{\text{call-repl}}(Th_i)$, $\mathcal{O}^\infty(\mathcal{R}_{\text{init-function}}(Th_i))$, and $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(Th_i))$ for $i \leq n$ are pairwise disjoint, and are all included in $\mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{CS})$.
 15. $|\mathcal{CT}| + \text{steps} \geq N_{\text{steps}}$.
 16. $\alpha \leq N_{\text{rand}}(\mathcal{CT}) + \sum_{O, \tau} N_{\text{calls}}(O, \tau, \mathcal{CT}) + 1$.
 17. If $O[[a', +\infty, \tilde{a}]] \in \mathcal{I}$, then $a' \leq N_{\text{calls}}(O, \tau_O(O[_, \tilde{a}]), \mathcal{CT}) + 1$.
 18. If $\text{role}[[a', +\infty, \tilde{a}]] \in \mathcal{RI}$, then $a' \leq N_{\text{exec}}(\text{role}, \mathcal{CT}) + 1$.

The relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ is our main tool to relate the CryptoVerif and OCaml systems. This relation holds only when the CryptoVerif adversary is evaluating the function *simulate* (line 7 of Figure 19), as shown by the form of the extended CryptoVerif configuration \mathfrak{C}^{cs} in Item 1. (The value α is the current value of

the index i' , that is, the number of iterations in the loop.) Items 1 and 2 also ensure that there is the same number of threads in the simulator configuration \mathcal{CS} and in the OCaml configuration \mathcal{C} .

Item 3 is an invariant on the CryptoVerif side: it relates the available oracles in \mathcal{Q} to elements of the simulator configuration. This item ensures basically that when the simulator calls an oracle present in \mathcal{I} , it is also present in \mathcal{Q} , and the oracle call in the CryptoVerif adversary (line 13 of Figure 19) can proceed. Item 4 is an invariant of the CryptoVerif semantics: the environment contains bindings for every free variable present in the current configuration. Item 5 is an invariant of the simulator: each store location that occurs in a thread is present in the domain of the store. (When a location is created, it is immediately added to the store.)

Item 6 relates the threads of the simulator and of the OCaml semantics. A thread can be in one of the following two states. If it satisfies Item 6(a), the thread is a protocol thread that was not scheduled yet. The simulator and OCaml threads correspond by transforming the program $program'(\text{role}[\tilde{a}])$ present in the simulator into the program of the module corresponding to the role, $program(\mu_{\text{role}})$. Otherwise, the thread satisfies Item 6(b). In this case, Item 6(b)i relates the contents of the simulator thread and the OCaml thread by replacing $program'(\text{role}[\tilde{a}])$ with $program(\mu_{\text{role}})$ as above, and by replacing calls to oracles using `call` with a corresponding tagged closure. The tokens that determine whether oracles can be called are absent from the simulator: the value of these tokens is determined from \mathcal{I} by the function `gettokens`, and we require that they are present in the OCaml store with their correct value. Item 6(b)ii ensures that all instances of a closure of a given role initialization $\text{role}[\tilde{a}]$ share the same store location for their tokens. This ensures that a role initialization closure is not called twice. Item 6(b)ii also ensures that all locations used for the tokens of role initialization are not accessible elsewhere. Item 6(b)iii ensures that every tagged closure present in the simulator is a correct closure for the initialization of a role. Item 6(b)iv is an invariant of the simulator that ensures that the adversary does not have access to our OCaml instrumentation features.

Items 7 to 10 relate the values of the global store in the simulator and in the OCaml semantics. The public part of the global store is the same on both sides (Item 10). The private part (files and tables) is empty in the simulator, since this part is handled by CryptoVerif itself (Item 8) and cannot be accessed by the adversary (Item 7). We require that the private part of the OCaml global store corresponds to the CryptoVerif configuration (Item 9).

Item 11 relates the OCaml multiset of callable modules \mathcal{MI} and the simulator set of callable roles \mathcal{RI} . Item 12 relates the OCaml and CryptoVerif events.

Items 13 and 14 present restrictions on sets of oracles. To understand how all these oracle sets interact, let us present the flow of an oracle not under replication $O[\tilde{a}]$ in these sets.

1. Initially, if the oracle occurs at the beginning of the process, it is in $O^\infty(\mathcal{RI})$; otherwise, it is in $\text{willbeavailable}(\mathcal{CS})$.

2. For an oracle occurring at the beginning of a role, when the role containing it is instantiated using `addthread`, the oracle moves from $\mathcal{O}^\infty(\mathcal{RI})$ to $\mathcal{O}^\infty(\mathcal{R}_{\text{init-function}}(Th))$. It is also added into $\mathcal{O}^\infty(\mathcal{I})$.
3. When the initialization function of the role is reduced into a closure, the oracle moves from $\mathcal{O}^\infty(\mathcal{R}_{\text{init-function}}(Th))$ to $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(Th))$.
4. When the initialization function of the role is called, the oracle moves from $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(Th))$ to $\mathcal{O}_{\text{call}}(Th)$.
5. When the oracle itself is called, it is removed from $\mathcal{O}^\infty(\mathcal{I})$, and when the call to the oracle disappears from the thread, it is removed from $\mathcal{O}_{\text{call}}(Th)$. The oracles made available after the call are removed from `willbeavailable(CS)` and added either to $\mathcal{O}^\infty(\mathcal{RI})$ if they start a role or to $\mathcal{O}^\infty(\mathcal{I})$ and $\mathcal{O}_{\text{call}}(Th)$ if they do not start a role.

The case of an oracle under replication is fairly similar, using $\mathcal{O}_{\text{call-repl}}(Th)$ instead of $\mathcal{O}_{\text{call}}(Th)$. Items 13 and 14 ensure that an oracle cannot be simultaneously in two different sets. These properties allow us to prove that the injections of Items 6(b)i and 6(b)ii are kept. (We use \mathcal{O}^∞ rather than \mathcal{O} to make sure that we cannot have simultaneously $O[[a', +\infty[, \tilde{a}] \in \mathcal{I}$ with $a' > N_O$ and $O[b, \tilde{a}] \in \text{willbeavailable}(\text{CS})$ for all b . Indeed, $\mathcal{O}(\{O[[a', +\infty[, \tilde{a}]\})$ is empty when $a' > N_O$, so this situation would not be prevented by Item 13 if it used \mathcal{O} . It is prevented using \mathcal{O}^∞ .)

Items 15 to 18 ensure that we never reach the limits on the number of simulator steps N_{steps} (Item 15), the number of calls to the oracles (Item 16 for the oracle O_{loop} and Item 17 for the other oracles), and the number of calls to roles (Item 18), by making sure that the number of calls on the `CryptoVerif` side is at most the number of calls on the `OCaml` side. The number of calls made to oracle $O[_ , \tilde{a}]$ in `CryptoVerif`, $a' - 1$ such that $O[[a', +\infty[, \tilde{a}] \in \mathcal{I}$, may be less than the number of calls to that oracle in the `OCaml` trace, $N_{\text{calls}}(O, \tau_O(O[_ , \tilde{a}]), \mathcal{CT})$, because failed calls are not counted on the `CryptoVerif` side.

The next two lemmas show that the relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ is preserved during execution. Lemma 6.28 shows that it holds at the beginning, as soon as the simulator reaches line 7 of Figure 19.

Lemma 6.28 *There exists a trace $\mathfrak{C}_0(Q_0, \text{program}_0) \rightsquigarrow^* \mathfrak{C}^{\text{cs}}$ where $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}_0$ and $\mathcal{CT}_0 = \mathcal{C}_0(Q_0, \text{program}_0)$.*

Lemma 6.29 shows that the relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ is preserved. More precisely, the relation does not hold at all steps (in particular because it holds only when the `CryptoVerif` adversary is executing `simulate`), but if it holds at some point, we can continue execution so that either it holds again at a later point, or execution ends with matching events.

Lemma 6.29 *Let \mathfrak{C}^{cs} such that there exists an accessible trace \mathcal{CT} satisfying $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$.*

- Either there exist n configurations $\mathfrak{C}_1^{\text{cs}}, \dots, \mathfrak{C}_n^{\text{cs}}$ and n traces $\mathfrak{C}^{\text{cs}} \rightsquigarrow_{p_1}^+ \mathfrak{C}_1^{\text{cs}}, \dots, \mathfrak{C}^{\text{cs}} \rightsquigarrow_{p_n}^+ \mathfrak{C}_n^{\text{cs}}$ such that none of these traces is a prefix of another, $\sum_{i \leq n} p_i = 1$, and for each accessible trace \mathcal{CT} such that $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$, there exist n pairwise disjoint trace sets $\mathcal{CTS}_1, \dots, \mathcal{CTS}_n$ such that all traces in these sets are extensions of \mathcal{CT} , none of these traces is a prefix of another, $\Pr[\mathcal{CTS}_i] = p_i \cdot \Pr[\mathcal{CT}]$, and for each trace $\mathcal{CT}' \in \mathcal{CTS}_i$, we have $\mathfrak{C}_i^{\text{cs}} \equiv \mathcal{CT}'$.
- Or for each accessible trace \mathcal{CT} such that $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$, the last configuration \mathcal{C} of \mathcal{CT} cannot reduce, $\mathfrak{C}^{\text{cs}} \rightarrow^+ \mathfrak{C}_1^{\text{cs}}$, the configuration $\mathfrak{C}_1^{\text{cs}}$ cannot reduce, and the event list \mathcal{E} of $\mathfrak{C}_1^{\text{cs}}$ and the event list events of \mathcal{C} satisfy $\text{events} = \mathbb{G}_{\text{ev}}(\mathcal{E})$.

We prove these lemmas in Appendix F. Let us present a proof sketch of Lemma 6.29.

Proof sketch Let us take an extended CryptoVerif configuration \mathfrak{C}^{cs} and an OCaml trace \mathcal{CT} such that $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$. Let \mathcal{C} be the last configuration of \mathcal{CT} . Let \mathcal{CS} be the configuration of the simulator in \mathfrak{C}^{cs} and Th be the current thread of \mathcal{CS} .

Case 1: the current thread of \mathcal{CS} verifies Item 6(a), we run the initialization of the module. The programs of the current threads of \mathcal{CS} and \mathcal{C} are the same except that $\text{program}'(\text{role}[\tilde{a}])$ present in \mathcal{CS} are transformed into $\text{program}(\mu_{\text{role}})$. We show that after having reduced the initialization of the primitives and the initialization of the roles in both sides, the current threads verify Item 6(b). The oracles in $\mathcal{O}^\infty(\mathcal{R}_{\text{init-function}}(Th))$ that correspond to the roles implemented in this initialization are moved to $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(Th))$. We prove that the relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ is preserved.

Case 2: the current thread of \mathcal{CS} verifies Item 6(b). We distinguish cases on the form of the simulator configuration \mathcal{CS} .

Let us first look at the cases in which the configuration \mathcal{CS} does not reduce. We use the rule (Leave Simulator), thus finishing the evaluation of the function *simulate*.

- If the current expression of \mathcal{CS} is $\text{call}(O_j[\tilde{a}]) v$, then the result of *simulate* is such that $o = o_j$, so the CryptoVerif adversary of Figure 19 calls the oracle O_j at line 13 in the branch $o = o_j$, ends one iteration of O_{loop} , and starts the next iteration until it reaches line 7. We use Lemma 6.10 and we exploit the definition of $\text{simulate}'_{O_j}$ and $\text{simulate}''_{O_j}$ to prove that the OCaml configuration reduces similarly, by calling the OCaml function generated for oracle O_j . The oracle $O_j[\tilde{a}]$ is removed from $\mathcal{O}^\infty(\mathcal{I})$, and from $\mathcal{O}_{\text{call}}(Th)$ if all occurrences of $\text{call}(O_j[\tilde{a}])$ have disappeared. The newly available oracles, added to sets $\mathcal{O}^\infty(\mathcal{RI})$ or $\mathcal{O}_{\text{call}}(Th)$ and $\mathcal{O}_{\text{call-repl}}(Th)$, are removed from the set $\text{willbeavailable}(\mathcal{CS})$. We prove that the relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ is preserved.
- If the current expression of \mathcal{CS} is $\text{random}()$, then the result of *simulate* is such that $o = o_R$, so the CryptoVerif adversary of Figure 19 samples a random boolean at line 19, ends one iteration of O_{loop} , and starts the next

iteration until it reaches line 7. The current expression of \mathcal{CS} is replaced with true with probability 1/2 and false with probability 1/2. The OCaml configuration reduces similarly: it samples a random boolean by using the rule (Random), and the relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ is preserved.

- Otherwise, the configuration \mathcal{CS} cannot reduce, and the corresponding configuration \mathcal{C} cannot reduce either. The result of *simulate* is such that $o = o_S$, so the CryptoVerif adversary of Figure 19 ends the current iteration of O_{loop} at line 9, and ends the loop at line 4, so it also stops. The events in the final CryptoVerif and OCaml configurations match, so the second case of the lemma holds.

If the current expression of \mathcal{CS} is `addthread(program)`, a new thread is created on both sides. If *program* is a protocol program, then this new thread satisfies Item 6(a) by definition of `addthread` in OCaml and in the simulator and by definition of `replaceinitpm`. The roles added in this new thread *Th* are removed from \mathcal{RI} and the corresponding oracles are added to $\mathcal{O}^\infty(\mathcal{R}_{\text{init-function}}(\textit{Th}))$ and to \mathcal{I} . Otherwise, the new thread satisfies Item 6(b). We prove that the relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ is preserved.

If the current expression of \mathcal{CS} is `call($O_j[\bar{a}]$) v` and \mathcal{CS} reduces by (FailedCall1) or (FailedCall2), then the simulator raises `Bad_Call`, and the corresponding tagged function in OCaml also raises `Bad_Call` (because the tokens in OCaml correspond to \mathcal{I} in the simulator by Item 6(b)i). We prove that the relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ is preserved.

If the current expression of \mathcal{CS} is `tagfunctionrole,τ[env, pm'role[\bar{a}]]` (\cdot), then we execute the initialization function of role *role*. We remove this role from $\mathcal{R}_{\text{init-closure}}(\textit{Th})$, and add the corresponding oracles to $\mathcal{O}_{\text{call}}(\textit{Th})$ and $\mathcal{O}_{\text{call-repl}}(\textit{Th})$. We prove that the relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$ is preserved.

The other cases are straightforward since the simulator mimics the OCaml semantics. They all preserve the relation $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$. \square

From Lemmas 6.28 and 6.29, we can prove the following proposition, by extending the traces using Lemma 6.29 until we get complete traces.

Proposition 6.30 *Let $\mathfrak{CT}_1, \dots, \mathfrak{CT}_n$ be complete CryptoVerif traces starting at $\mathfrak{C}_0(Q_0, pe_0)$.*

Then there exist disjoint sets of complete OCaml traces $\mathcal{CTS}_1, \dots, \mathcal{CTS}_n$ all starting at $\mathfrak{C}_0(Q_0, pe_0)$ such that for all $i \leq n$, $\Pr[\mathfrak{CT}_i] = \Pr[\mathcal{CTS}_i]$, and if \mathfrak{C} is the last configuration of \mathfrak{CT}_i and \mathcal{C} is the last configuration of a trace in \mathcal{CTS}_i , then the event list \mathcal{E} of \mathfrak{C} and the event list events of \mathcal{C} satisfy $\text{events} = \mathbb{G}_{\text{ev}}(\mathcal{E})$.

We prove this proposition in more detail in Appendix G. As an immediate consequence of this proposition, we obtain:

Proposition 6.31 $\Pr[\mathfrak{C}_0(Q_0, \textit{program}_0)(\rightsquigarrow) : D] = \Pr[\mathcal{C}_0(Q_0, \textit{program}_0) : D]$.

6.4.3 Security Result

By combining Propositions 6.17 and 6.31, we obtain the following theorem:

Theorem 6.32 (Security result)

$$\Pr[\mathfrak{C}_0(Q_0, program_0) : D] = \Pr[\mathcal{C}_0(Q_0, program_0) : D].$$

In other words, the adversary $program_0$ against our generated OCaml modules has the same probability of breaking the security property as the adversary $Q_{adv}(Q_0, program_0)$ against the CryptoVerif process.

CryptoVerif bounds the probability that an adversary Q breaks the security property D , that is, it finds a probability p that depends on the adversary such that, for all CryptoVerif adversaries Q for Q_0 ,

$$\Pr[\mathfrak{C}_i(Q_0 \mid Q) : D] \leq p.$$

The adversaries $Q_{adv}(Q_0, program_0)$ are CryptoVerif adversaries for Q_0 , so for all OCaml programs $program$ that obey our assumptions,

$$\Pr[\mathcal{C}_0(Q_0, program) : D] = \Pr[\mathfrak{C}_0(Q_0, program) : D] \leq p$$

Hence, all considered OCaml adversaries $program$ can break the security property D with probability at most p .

The probability bound p returned by CryptoVerif is a function that depends on many parameters, expressed on the CryptoVerif protocol specification. Let us relate these parameters to the OCaml implementation. These parameters are as follows:

- The maximum number of times the various oracles and roles have been called, N_O and N_{role} . As shown by our proof and by Definition 6.11, N_O can be set to the maximum number of calls to the same closure representing oracle O in any trace of the OCaml program, and N_{role} can be set to the maximum number of instantiations of the role $role$ in any trace of this program.
- The size of the CryptoVerif types T . The corresponding OCaml type $\mathbb{G}_T(T)$ is fixed by the annotations of the CryptoVerif specification. The size of T can be set to the size of $\mathbb{G}_T(T)$. Similarly, the size of the CryptoVerif values a (used when their type T has unbounded size) can be set to the size of the corresponding OCaml value $\mathbb{G}_{valT}(a)$.
- The execution time of the cryptographic primitives and of various CryptoVerif constructs. This time can be set to the execution time of the corresponding OCaml implementation.
- The execution time of the adversary. Our proof shows that the function *simulate* executes at most as many reduction steps as the OCaml adversary. However, the CryptoVerif adversary shown in Figure 19 also includes

additional steps and conversions between the OCaml semantic configuration and its CryptoVerif bitstring representation. By using the contents of the OCaml memory as bitstring representation of the semantic configuration in CryptoVerif, we can obtain an efficient implementation of the CryptoVerif adversary that does not take significantly more time than the OCaml adversary.

From the probability bound given by CryptoVerif, we can then obtain a bound on the probability of breaking the security properties in the generated OCaml implementation of the protocol.

As detailed in [8], CryptoVerif shows that our model of the SSH Transport Layer Protocol guarantees the authentication of the server to the client and the secrecy of the session keys. By Theorem 6.32, our generated implementation of this protocol satisfies the same properties, provided assumptions A1 to A6 hold.

7 Conclusion

We have proved that our compiler preserves security. Therefore, by using CryptoVerif, we can prove the desired security properties on the protocol specification, and then by using our compiler, we get a runnable implementation of the protocol, which satisfies the same security properties as the specification. Making such a proof is also useful because it clarifies the assumptions needed to ensure that the implementation is secure (Assumptions A1 to A6 in our case). The proof technique presented in this paper, simulating any adversary by a CryptoVerif process, is also useful to show that any Turing machine can be encoded as a CryptoVerif adversary, which is important for the validity of the verification by CryptoVerif. We have done the proof by hand. Formalizing it using a proof assistant (e.g. Coq) would be interesting future work.

Acknowledgments. This work was partly done while the authors were at École Normale Supérieure, Paris. It was partly supported by the ANR project ProSe (decision ANR 2010-VERS-004).

References

- [1] M. Aizatulin, A. D. Gordon, and J. Jürjens. Extracting and verifying cryptographic models from C protocol code by symbolic execution. In *CCS'11*, pages 331–340, New York, 2011. ACM.
- [2] M. Aizatulin, A. D. Gordon, and J. Jürjens. Computational verification of C protocol implementations by symbolic execution. In *CCS'12*, pages 712–723, New York, 2012. ACM.
- [3] J. Bengtson, K. Bhargavan, C. Fournet, A. Gordon, and S. Maffei. Refinement types for secure implementations. *ACM TOPLAS*, 33(2), 2011.

- [4] K. Bhargavan, C. Fournet, A. Gordon, and S. Tse. Verified interoperable implementations of security protocols. *ACM TOPLAS*, 31(1), 2008.
- [5] B. Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Transactions on Dependable and Secure Computing*, 5(4):193–207, 2008.
- [6] B. Blanchet. Automatically verified mechanized proof of one-encryption key exchange. In *CSF'12*, pages 325–339, Los Alamitos, 2012. IEEE.
- [7] B. Blanchet and D. Pointcheval. Automated security proofs with sequences of games. In *CRYPTO'06*, volume 4117 of *LNCS*, pages 537–554. Springer, 2006.
- [8] D. Cadé and B. Blanchet. From computationally-proved protocol specifications to implementations and application to SSH. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, 4(1):4–31, Mar. 2013.
- [9] D. Cadé and B. Blanchet. Proved generation of implementations from computationally-secure protocol specifications. In D. Basin and J. Mitchell, editors, *2nd Conference on Principles of Security and Trust (POST 2013)*, volume 7796 of *LNCS*, pages 63–82, Rome, Italy, Mar. 2013. Springer.
- [10] S. Chaki and A. Datta. ASPIER: An automated framework for verifying security protocol implementations. In *CSF'09*, pages 172–185, Los Alamitos, 2009. IEEE.
- [11] F. Dupressoir, A. D. Gordon, J. Jürjens, and D. A. Naumann. Guiding a general-purpose C verifier to prove cryptographic protocols. In *CSF'11*, pages 3–17, Los Alamitos, 2011. IEEE.
- [12] C. Fournet, M. Kohlweiss, and P.-Y. Strub. Modular code-based cryptographic verification. In *CCS'11*, pages 341–350, New York, 2011. ACM.
- [13] <http://msr-inria.inria.fr/projects/sec/fs2cv/>.
- [14] G. Milicia. χ -spaces: Programming security protocols. In *NWPT'02*, 2002.
- [15] S. Owens. A sound semantics for OCaml light. In S. Drossopoulou, editor, *ESOP'08*, volume 4960 of *LNCS*, pages 1–15, Heidelberg, 2008. Springer.
- [16] A. Pironti and R. Sisto. Provably correct Java implementations of spi calculus security protocols specifications. *Computers and Security*, 29(3):302–314, 2010.
- [17] N. Swamy, J. Chen, C. Fournet, P.-Y. Strub, K. Bhargavan, and J. Yang. Secure distributed programming with value-dependent types. In *ICFP'11*, pages 266–278, New York, 2011. ACM.

Appendices

Note to the reviewers: these appendices provide detailed proofs of all results of the paper. Since the paper is very long, we are open to publishing these proofs as a technical report instead of including them in the journal paper. They are provided here so that the reviewers can check our results.

A Proof of Lemma 2.5

Proof (of Lemma 2.5) As usual, a multiset S is defined as a function from elements to integers: $S(x)$ is the number of occurrences of x in the multiset S . Multiset union is defined as addition: $(S \uplus S')(x) = S(x) + S'(x)$. The maximum $\max(S, S')$ is the multiset such that $\max(S, S')(x) = \max(S(x), S'(x))$. The inclusion $S \subseteq S'$ is true when $\text{Dom}(S) \subseteq \text{Dom}(S')$ and $\forall x \in \text{Dom}(S), S(x) \leq S'(x)$.

We define the multiset of available oracles inductively as follows:

$$\text{Oracles}(0) = \emptyset$$

$$\text{Oracles}(Q_1 \mid Q_2) = \text{Oracles}(Q_1) \uplus \text{Oracles}(Q_2)$$

$$\text{Oracles}(\text{foreach } i \leq n \text{ do } Q) = \bigsqcup_{a \in [1, n]} \text{Oracles}(Q\{a/i\})$$

$$\text{Oracles}(O[\tilde{a}](x_1[\tilde{a}] : T_1, \dots, x_k[\tilde{a}] : T_k) := P) = \{O[\tilde{a}]\} \uplus \text{Oracles}(P)$$

$$\text{Oracles}(\text{return}(M_1, \dots, M_k); Q) = \text{Oracles}(Q)$$

$$\text{Oracles}(\text{end}) = \emptyset$$

$$\text{Oracles}(x[\tilde{a}] \stackrel{R}{\leftarrow} T; P) = \text{Oracles}(P)$$

$$\text{Oracles}(x[\tilde{a}] \leftarrow M; P) = \text{Oracles}(P)$$

$$\text{Oracles}(\text{insert } \text{Tbl}(M_1, \dots, M_l); P) = \text{Oracles}(P)$$

$$\text{Oracles}(\text{get } \text{Tbl}(x_1[\tilde{a}], \dots, x_l[\tilde{a}]) \text{ suchthat } M \text{ in } P \text{ else } P') = \max(\text{Oracles}(P), \text{Oracles}(P'))$$

$$\text{Oracles}(\text{event } e(M_1, \dots, M_l); P) = \text{Oracles}(P)$$

$$\text{Oracles}(\text{let } (x_1[\tilde{i}] : T_1, \dots, x_{k'}[\tilde{i}] : T_{k'}) = O[\tilde{M}](\tilde{M}') \text{ in } P \text{ else } P') = \max(\text{Oracles}(P), \text{Oracles}(P'))$$

$$\text{Oracles}(\text{let } x[\tilde{i}] : T = \text{loop } O[\tilde{M}](M') \text{ in } P \text{ else } P') = \max(\text{Oracles}(P), \text{Oracles}(P'))$$

$$\text{Oracles}(E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}) = \text{Oracles}(P) \uplus \bigsqcup_{Q \in \mathcal{Q}} \text{Oracles}(Q) \uplus$$

$$\bigsqcup_{((x_1[\tilde{a}], \dots, x_k[\tilde{a}]), P', P'') \in \mathcal{R}} \max(\text{Oracles}(P'), \text{Oracles}(P''))$$

We show that, for all configurations $\mathfrak{C} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$ reachable from the initial configuration $\mathfrak{C}_i(Q_0)$, $\text{Oracles}(\mathfrak{C})$ contains no duplicates.

Let us first show this property for the initial configuration. We show by an easy induction on Q that $\bigsqcup_{Q' \in \text{reduce}(Q)} \text{Oracles}(Q') \subseteq \text{Oracles}(Q)$. Therefore, by definition of \mathfrak{C}_i , we have $\text{Oracles}(\mathfrak{C}_i(Q_0)) = \bigsqcup_{Q' \in \text{reduce}(Q_0)} \text{Oracles}(Q') \subseteq \text{Oracles}(Q_0)$. Next, we show by induction on Q_0 that $\text{Oracles}(Q_0)$ contains no duplicates.

- In the case $Q \mid Q'$, the oracles defined in Q and Q' are not in different branches of if or get, so by Property 2.3, they have different names. Hence, $\text{Oracles}(Q)$ and $\text{Oracles}(Q')$ do not both contain $O[\tilde{a}]$ for the same O . We conclude that $\text{Oracles}(Q) \uplus \text{Oracles}(Q')$ contains no duplicates using the induction hypothesis.
- In the case **foreach** $i \leq n$ **do** Q , by Property 2.3, the replication index i occurs as index in all definitions of oracles in Q , and in the same position. So the multisets $\text{Oracles}(Q\{a/i\})$ are disjoint for different choices of a . We conclude that $\bigsqcup_{a \in [1, n]} \text{Oracles}(Q\{a/i\})$ contains no duplicates using the induction hypothesis.
- In the case $O[\tilde{a}](x_1[\tilde{a}] : T_1, \dots, x_k[\tilde{a}] : T_k) := P$, the definition of O is not in a branch of if or get different from P , so by Property 2.3, there is no definition of O in P . Hence $O[\tilde{a}] \notin \text{Oracles}(P)$. We conclude that $\{O[\tilde{a}]\} \uplus \text{Oracles}(P)$ contains no duplicates using the induction hypothesis.
- In all other cases, the result follows immediately from the induction hypothesis.

Furthermore, $\text{Oracles}(\mathfrak{C})$ decreases by reduction: if $\mathfrak{C} \rightarrow_p \mathfrak{C}'$, then we have $\text{Oracles}(\mathfrak{C}') \subseteq \text{Oracles}(\mathfrak{C})$. Indeed, the rules (New), (Let), (Insert), (Event), (Loop1) leave $\text{Oracles}(\mathfrak{C})$ unchanged. In the case of (Loop1), we use

$$\begin{aligned} \max(\max(\max(\text{Oracles}(P), \text{Oracles}(P')), \text{Oracles}(P)), \text{Oracles}(P')) = \\ \max(\text{Oracles}(P), \text{Oracles}(P')). \end{aligned}$$

The rules (If1), (If2), (Get1), (Get2), (Loop2), (Return), (End) decrease the multiset $\text{Oracles}(\mathfrak{C})$ by replacing $\max(\text{Oracles}(P), \text{Oracles}(P'))$ with either $\text{Oracles}(P)$ or $\text{Oracles}(P')$. In the case of (Return), we also use $\bigsqcup_{Q' \in \text{reduce}(Q'')} \text{Oracles}(Q') \subseteq \text{Oracles}(Q'')$. The rule (Call) removes the called oracle $O[\tilde{a}']$ from $\text{Oracles}(\mathfrak{C})$.

Therefore, for all configurations \mathfrak{C} reachable from the initial configuration $\mathfrak{C}_i(Q_0)$, $\text{Oracles}(\mathfrak{C})$ contains no duplicates.

By definition of Oracles , when $\mathfrak{C} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$, we have $\{O[\tilde{a}] \mid O[\tilde{a}](x_1[\tilde{a}] : T_1, \dots, x_k[\tilde{a}] : T_k) := P \in \mathcal{Q}\} \subseteq \text{Oracles}(\mathfrak{C})$. (Both sides of the inclusion are multisets.) Therefore, \mathcal{Q} contains at most one element $O[\tilde{a}](x_1[\tilde{a}] : T_1, \dots, x_k[\tilde{a}] : T_k) := P$ for each $O[\tilde{a}]$. \square

B Proof of Proposition 4.5

Proof (of Proposition 4.5) Let $n_{\text{ev,ret}}(\mathcal{CI})$ be the number of occurrences of event or return in the instrumented configuration \mathcal{CI} . Let us first prove the following property:

2'. If $\mathcal{C} \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI})$ and $\mathcal{CI}_1, \dots, \mathcal{CI}_n$ are pairwise distinct instrumented configurations such that for all $i \leq n$, $\mathcal{CI} \rightarrow_{p_i} \mathcal{CI}_i$ with $\sum_{i \leq n} p_i = 1$, then one of the following two properties holds:

P1. $n = 1$, $\mathcal{C} \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI}_1)$, and $n_{\text{ev,ret}}(\mathcal{CI}_1) < n_{\text{ev,ret}}(\mathcal{CI})$.

P2. there exist pairwise distinct configurations $\mathcal{C}_1, \dots, \mathcal{C}_n$ such that for all $i \leq n$, we have $\mathcal{C} \rightarrow_{p_i} \mathcal{C}_i$ and $\mathcal{C}_i \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI}_i)$.

We prove this property by case analysis on the possible reductions of \mathcal{CI} .

Let us first suppose that \mathcal{CI} reduces by (Globalstore1) and (Toplevel) from a reduction $Th \rightarrow_p Th'$ of the current thread, and let us distinguish cases depending on the latter reduction:

- The reduction comes from rules (Context in) or (Let ctx in): we have $Th = \langle env, C[e], stack, store \rangle \rightarrow Th' = \langle env, e, (env, C) :: stack, store \rangle$ where e is not a value and C is a minimal expression or program evaluation context. Let us distinguish cases on the form of C .
 - If $C = \text{return}(\mathcal{ML}, [\cdot])$, then we have $\text{noinstr}_{Th}(Th) = \text{noinstr}_{Th}(\langle env, e, stack, store \rangle) = \text{noinstr}_{Th}(Th')$ by Definition 4.4, so by expanding this property to the complete configuration, and noting that the reduction removes one return, Property P1 holds.
 - If $C = \text{event } ev(e_1, \dots, e_{i-1}, [\cdot], v_{i+1}, \dots, v_n)$, then we have by Definition 4.4,

$$\begin{aligned} \text{noinstr}_{Th}(Th) &= \\ &\text{noinstr}_{Th}(\langle env, (e_1, \dots, e_{i-1}, e, v_{i+1}, \dots, v_n), stack, store \rangle), \\ \text{noinstr}_{Th}(Th') &= \text{noinstr}_{Th}(\langle env, e, (env, (e_1, \dots, e_{i-1}, [\cdot], v_{i+1}, \dots, v_n)) :: stack, store \rangle), \end{aligned}$$

and we have $\text{noinstr}_{Th}(Th) \rightarrow \text{noinstr}_{Th}(Th')$, so Property P2 holds.

- If C is neither a return nor an event context, then the reduction $Th \rightarrow_p Th'$ implies $\text{noinstr}_{Th}(Th) \rightarrow_p \text{noinstr}_{Th}(Th')$, so Property P2 holds.
- The reduction comes from rules (Context out) or (Let ctx out): we have $Th = \langle env, v, (env', C) :: stack, store \rangle \rightarrow Th' = \langle env', C[v], stack, store \rangle$ where C is a minimal expression or program evaluation context. Let us distinguish cases on the form of C .

– If $C = \text{return}(\mathcal{M}\mathcal{I}, [\cdot])$, then we have by Definitions 4.4 and 4.2,

$$\begin{aligned} \text{noinstr}_{Th}(Th) &= \text{noinstr}_{Th}(\langle env, v, stack, store \rangle) \\ &\approx_{v, Th} \text{noinstr}_{Th}(Th'), \end{aligned}$$

so by expanding this property to the complete configuration, and noting that the reduction removes one return, Property P1 holds.

– If $C = \text{event } ev(e_1, \dots, e_{i-1}, [\cdot], v_{i+1}, \dots, v_n)$, we have by Definition 4.4,

$$\begin{aligned} \text{noinstr}_{Th}(Th) &= \text{noinstr}_{Th}(\langle env, v, (env', (e_1, \dots, e_{i-1}, [\cdot], v_{i+1}, \dots, v_n)) :: stack, store \rangle), \\ \text{noinstr}_{Th}(Th') &= \text{noinstr}_{Th}(\langle env', (e_1, \dots, e_{i-1}, v, v_{i+1}, \dots, v_n), stack, store \rangle), \end{aligned}$$

so Property P2 holds.

– If C is neither a return nor an event context, then Property P2 holds.

- The cases of (Context raise2) and (Let ctx raise) are similar to the previous case: Property P1 holds when the context is $\text{return}(\mathcal{M}\mathcal{I}, [\cdot])$; Property P2 holds otherwise.
- Property P2 holds in the other cases.

If $\mathcal{C}\mathcal{I} \rightarrow \mathcal{C}\mathcal{I}_1$ by (Toplevel return), then the program of the current thread is $\text{return}(\mathcal{M}\mathcal{I}, v)$ in $\mathcal{C}\mathcal{I}$ and the only differences between $\mathcal{C}\mathcal{I}$ and $\mathcal{C}\mathcal{I}_1$ are that the program of the current thread is v and the set of callable modules is changed in $\mathcal{C}\mathcal{I}_1$. Therefore, $\text{noinstr}_{\mathcal{C}\mathcal{I}}(\mathcal{C}\mathcal{I}) = \text{noinstr}_{\mathcal{C}\mathcal{I}}(\mathcal{C}\mathcal{I}_1)$, and the reduction removes one return, so Property P1 holds.

If $\mathcal{C}\mathcal{I} \rightarrow \mathcal{C}\mathcal{I}_1$ by (Toplevel event), then the program of the current thread is $\text{event } ev(v_1, \dots, v_n)$ in $\mathcal{C}\mathcal{I}$ and the only differences between $\mathcal{C}\mathcal{I}$ and $\mathcal{C}\mathcal{I}_1$ are that the program of the current thread is (v_1, \dots, v_n) and the list of executed events is updated in $\mathcal{C}\mathcal{I}_1$. Therefore, $\text{noinstr}_{\mathcal{C}\mathcal{I}}(\mathcal{C}\mathcal{I}) = \text{noinstr}_{\mathcal{C}\mathcal{I}}(\mathcal{C}\mathcal{I}_1)$, and the reduction removes one event, so Property P1 holds.

Property P2 holds for `addthread`, `schedule`, and global store related reductions. In the case of `addthread`, we use Assumption 4.1.

We have proved that Property 2' holds. Property 2 also holds, since it is a special case of Property 2'.

Let us now prove:

3. If $\mathcal{C} \approx_v \text{noinstr}_{\mathcal{C}\mathcal{I}}(\mathcal{C}\mathcal{I})$ and $\mathcal{C}\mathcal{I}$ cannot reduce, then \mathcal{C} cannot reduce.

We need to prove that $\text{noinstr}_{\mathcal{C}\mathcal{I}}(\mathcal{C}\mathcal{I})$ cannot reduce. We can then use Lemma 4.3 to conclude. We distinguish cases depending on the program or expression in the current thread of $\mathcal{C}\mathcal{I}$. If this program or expression was of the form $C[e]$ for some program or expression minimal evaluation context C , or $\text{return}(\mathcal{M}\mathcal{I}, v)$, or $\text{event } ev(v_1, \dots, v_n)$, the configuration $\mathcal{C}\mathcal{I}$ could reduce. In all other cases, $\text{noinstr}_{\mathcal{C}\mathcal{I}}$ does not change the form of the possible reductions (since it transforms

tagged functions into functions that behave exactly in the same way). Property 3 is true.

Let us now prove Property 1 by induction on $n_{\text{ev,ret}}(\mathcal{CI})$. Let us suppose that $\mathcal{C} \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI})$ and $\mathcal{C}_1, \dots, \mathcal{C}_n$ are pairwise distinct configurations such that for all $i \leq n$, we have $\mathcal{C} \rightarrow_{p_i} \mathcal{C}_i$ with $\sum_{i \leq n} p_i = 1$.

The configuration \mathcal{CI} must reduce, otherwise, by Property 3, the configuration \mathcal{C} would also not reduce. Let $\mathcal{CI} \rightarrow_{p'_i} \mathcal{CI}_i$ for $i \leq n'$ be all the reductions possible from \mathcal{CI} . By Property 2', we are either in case P1 or in case P2.

In case P1, \mathcal{CI} reduces into only one configuration \mathcal{CI}_1 such that $\mathcal{C} \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI}_1)$, and $n_{\text{ev,ret}}(\mathcal{CI}_1) < n_{\text{ev,ret}}(\mathcal{CI})$. By induction hypothesis, there exist pairwise distinct instrumented configurations $\mathcal{CI}'_1, \dots, \mathcal{CI}'_n$ such that for all $i \leq n$, we have $\mathcal{CI}_1 \xrightarrow{p_i} \mathcal{CI}'_i$ and $\mathcal{C}_i \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI}'_i)$. As there is only one reduction from \mathcal{CI} to \mathcal{CI}_1 with probability 1, we can conclude that Property 1 holds in this case.

In case P2, there exist pairwise distinct configurations $\mathcal{C}'_1, \dots, \mathcal{C}'_{n'}$ such that for all $i \leq n'$, we have $\mathcal{C} \rightarrow_{p'_i} \mathcal{C}'_i$ and $\mathcal{C}'_i \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI}_i)$. Since $\sum_{i \leq n} p_i = 1$ and $\sum_{i \leq n'} p'_i = 1$, the reductions $\mathcal{C} \rightarrow_{p_i} \mathcal{C}_i$ ($i \leq n$) are all possible reductions of \mathcal{C} and the reductions $\mathcal{C} \rightarrow_{p'_i} \mathcal{C}'_i$ ($i \leq n'$) are also all possible reductions of \mathcal{C} , so they are the same reductions. Therefore, $n = n'$ and there exists a bijection α from $\{1, \dots, n\}$ to $\{1, \dots, n\}$ such that $p_i = p'_{\alpha(i)}$, $\mathcal{C}_i = \mathcal{C}'_{\alpha(i)} \approx_v \text{noinstr}_{\mathcal{CI}}(\mathcal{CI}_{\alpha(i)})$, and $\mathcal{CI} \rightarrow_{p_i} \mathcal{CI}_{\alpha(i)}$. By renumbering the configurations \mathcal{CI}_i ($i \leq n$), we can conclude that Property 1 holds in this case. \square

C Proof of Proposition 6.5

Definition C.1 Let $Th \stackrel{\text{def}}{=} \langle env, pe, stack, store \rangle$ and $Th' \stackrel{\text{def}}{=} \langle env', pe', stack', store' \rangle$ be two threads, such that the domains of $store$ and $store'$ are disjoint.

We define $\text{plug}(Th, Th') \stackrel{\text{def}}{=} \langle env, pe, stack @ stack', store \cup store' \rangle$.

Definition C.2 A well-formed thread $Th = \langle env, pe, stack, store \rangle$ is a thread such that:

1. all store locations $l \in S_l$ that occur in the thread Th are bound in the store: $l \in \text{Dom}(store)$,
2. pe and $stack$ do not contain global store locations, nor return, event, schedule, or addthread operations.

Lemma C.3 If Th is a well-formed thread and $Th \rightarrow_p Th'$, then for all Th'' such that the domains of the stores of Th'' and of Th are disjoint, after renaming the fresh locations introduced in $Th \rightarrow_p Th'$ so that they do not occur in Th'' , we have $\text{plug}(Th, Th'') \rightarrow_p \text{plug}(Th', Th'')$ and the domains of the stores of Th'' and of Th' are disjoint.

Proof By reviewing the reduction rules, we have Property (P1): if $env, pe, stack \xrightarrow{L} env', pe', stack'$, then for every stack $stack''$, we have $env, pe, stack @ stack'' \xrightarrow{L} env', pe', stack' @ stack''$.

Let $Th \stackrel{\text{def}}{=} \langle env, pe, stack, store \rangle$ and $Th' \stackrel{\text{def}}{=} \langle env', pe', stack', store' \rangle$. Let us prove that, if $Th \rightarrow_p Th'$, then for every $Th'' \stackrel{\text{def}}{=} \langle env'', pe'', stack'', store'' \rangle$ such that $\text{Dom}(store) \cap \text{Dom}(store'') = \emptyset$, we have the reduction $plug(Th, Th'') = \langle env, pe, stack @ stack'', store \cup store'' \rangle \rightarrow_p plug(Th', Th'') = \langle env', pe', stack' @ stack'', store' \cup store'' \rangle$ with $\text{Dom}(store') \cap \text{Dom}(store'') = \emptyset$. We distinguish cases on the label L present in rule (Thread).

- if L is empty, then by (Store empty), $store = store'$. We conclude by Property (P1) and rule (Thread).
- if L is $!l = v$, by (Store lookup), the location l is in the domain of the store $store$, and $store(l) = v$, and $store = store'$. We also have $(store \cup store'')(l) = v$, so $store \cup store'' \xrightarrow{!l=v} store' \cup store''$. We conclude by Property (P1) and rule (Thread).
- if L is $l := v$, then by (Store assign), the location l is in the domain of the store $store$, and $store' = store[l \mapsto v]$. The domain of the store $store \cup store''$ also contains l , so $store \cup store'' \xrightarrow{l:=v} store' \cup store''$. We conclude by Property (P1) and rule (Thread).
- if L is $\text{ref } v = l$, then by (Store alloc), $l \notin \text{Dom}(store)$. By reviewing the uses of L of the form $\text{ref } v = l$ in the reduction rules, we can deduce that $pe = \text{ref } v$ and $pe' = l$. By Property 1 of Definition C.2, the location l does not occur in Th . Let us take a location $l' \in S_l$ that is not in $\text{Dom}(store) \cup \text{Dom}(store'')$; we rename l into l' . As $l' \notin \text{Dom}(store)$, the thread Th' becomes $\langle env', l', stack', store[l' \mapsto v] \rangle$, which is in the same equivalence class as Th' , so we still designate this thread by Th' . Let $L' \stackrel{\text{def}}{=} (\text{ref } v = l')$. By Property (P1), $env, pe, stack @ stack'' \xrightarrow{L'} env', l', stack' @ stack''$ and, since $l' \notin \text{Dom}(store) \cup \text{Dom}(store'')$, we have $store \cup store'' \xrightarrow{L'} store[l' \mapsto v] \cup store''$. We conclude that $plug(Th, Th'') = \langle env, pe, stack @ stack'', store \cup store'' \rangle \rightarrow plug(Th', Th'') = \langle env', l', stack' @ stack'', store[l' \mapsto v] \cup store'' \rangle$ by rule (Thread). \square

Lemma C.4 *Let Th be a well-formed thread. If $Th \rightarrow_p Th'$, then Th' is also well-formed.*

Proof Let us prove that both properties of Definition C.2 are preserved.

- The only rule that may add new locations in the thread is (Store alloc), which creates a new location l and also adds it in the domain of the store. So Property 1 is preserved for Th' .
- By looking at the reduction rules, we can see that no rule can create global store locations or `return`, `event`, `schedule`, or `addthread` operations. So Property 2 is preserved for Th' .

Therefore, the thread Th' is also well-formed. \square

Lemma C.5 *If v is a value of the type of the argument of the primitive s , then the thread $\langle \emptyset, env_{\text{prim}}(s) v, [], \emptyset \rangle$ is well-formed.*

Proof The thread $Th_0^s = \langle \emptyset, program_{\text{prim}};;, [], \emptyset \rangle$ is well-formed, since it contains no locations in S_l and by Assumption 6.1, it contains no schedule, `addthread`, `return`, or event operations and no global store locations.

By Assumption 6.2, $Th_0^s \rightarrow^* Th = \langle env_{\text{prim}}, \varepsilon, [], \emptyset \rangle$, so by Lemma C.4, Th is also well-formed. Therefore, the thread $\langle \emptyset, env_{\text{prim}}(s) v, [], \emptyset \rangle$ is well-formed, since by Assumption 6.3, v contains no locations and no `return`, `event`, `schedule`, or `addthread` operations since it contains no closure, and env_{prim} contains no locations and no `return`, `event`, `schedule`, or `addthread` operations since Th is well-formed. \square

Proof (of Proposition 6.5) By Assumption 6.4, we have reductions of the form

$$Th_1 \stackrel{\text{def}}{=} \langle \emptyset, env_{\text{prim}}(s) v, [], \emptyset \rangle \rightarrow_p^* Th_1' \stackrel{\text{def}}{=} \langle env', v', [], store'_1 \rangle.$$

Let $Th_2 \stackrel{\text{def}}{=} \langle env, (), stack, store \rangle$. By Lemma C.5, Th_1 is well-formed, so by Lemmas C.3 and C.4,

$$Th = \text{plug}(Th_1, Th_2) \rightarrow_p^* \text{plug}(Th_1', Th_2) = \langle env', v', stack, store'_1 \cup store \rangle.$$

Letting $store' \stackrel{\text{def}}{=} store'_1 \cup store$, we obtain exactly the desired reductions $Th \rightarrow_p^* \langle env', v', stack, store' \rangle$, and $store' \supseteq store$. \square

D Proof of Lemma 6.10

Let us first prove lemmas useful to prove Lemma 6.10.

Lemma D.1 (Write file) *Let \mathcal{C} be an OCaml configuration. If $\mathcal{C}_{Th}(\mathcal{C}) = \langle env, \mathbb{G}_{\text{file}}(x[\tilde{a}]), stack, store \rangle$, $env(\mathbb{G}_{\text{var}}(x)) = \mathbb{G}_{\text{val}T_x}(a)$, $env \supseteq env_{\text{prim}}$, and $\mathcal{C}_{\text{globalstore}}(\mathcal{C}) \supseteq \text{globalstore}(E, \mathcal{T})$, then we have $\mathcal{C} \rightarrow^* \mathcal{C}'$ where*

- $\mathcal{C}' = \mathcal{C}[\text{Th} \mapsto \langle env, (), stack, store' \rangle, \text{globalstore} \mapsto \text{globalstore}']$,
- $store' \supseteq store$,
- $\text{globalstore}' \supseteq \text{globalstore}(E[x[\tilde{a}] \mapsto a], \mathcal{T})$,
- $\text{globalstore}'(l) = \mathcal{C}_{\text{globalstore}}(\mathcal{C})(l)$ for all $l \notin S_{\text{priv}}$.

Proof If $(x[\tilde{a}], f) \in \text{files}$ for some f , then we have $x[\tilde{a}] = x[]$ and $\mathbb{G}_{\text{file}}(x[\tilde{a}]) = (f := \mathbb{G}_{\text{ser}}(T_x) \mathbb{G}_{\text{var}}(x))$, so

$$\mathcal{C} \rightarrow \mathcal{C}[\text{Th} \mapsto \langle env, \mathbb{G}_{\text{ser}}(T_x) \mathbb{G}_{\text{var}}(x), stack', store \rangle]$$

$$\text{where } stack' \stackrel{\text{def}}{=} (env, f := [\cdot]) :: stack$$

$$\rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env, env_{\text{prim}}(\mathbb{G}_{\text{ser}}(T_x)) \mathbb{G}_{\text{val}T_x}(a), stack', store \rangle]$$

$$\rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env', \text{ser}(T_x, a), stack', store' \rangle] \quad \text{by Proposition 6.5}$$

$$\rightarrow \mathcal{C}[\text{Th} \mapsto \langle env, f := \text{ser}(T_x, a), stack, store' \rangle]$$

$$\rightarrow \mathcal{C}' = \mathcal{C}[\text{Th} \mapsto \langle env, (), stack, store' \rangle, \text{globalstore} \mapsto \text{globalstore}']$$

where

$$\begin{aligned} \text{globalstore}' &\stackrel{\text{def}}{=} \mathcal{C}_{\text{globalstore}}(\mathcal{C})[f \mapsto \text{ser}(T_x, a)] \\ &\supseteq \text{globalstore}(E, \mathcal{T})[f \mapsto \text{ser}(T_x, a)] \\ &\supseteq \text{globalstore}(E[x[] \mapsto a], \mathcal{T}). \end{aligned}$$

The modified location f is in S_{priv} , so for all $l \notin S_{\text{priv}}$, we have $\text{globalstore}'(l) = \mathcal{C}_{\text{globalstore}}(\mathcal{C})(l)$. We have $\text{store}' \supseteq \text{store}$ by Proposition 6.5.

Otherwise, we have $\mathbb{G}_{\text{file}}(x[\tilde{a}]) = ()$ and $\mathcal{C}' = \mathcal{C}$, so

$$\begin{aligned} \text{globalstore}' &= \mathcal{C}_{\text{globalstore}}(\mathcal{C}) \\ &\supseteq \text{globalstore}(E, \mathcal{T}) = \text{globalstore}(E[x[\tilde{a}] \mapsto a], \mathcal{T}), \end{aligned}$$

and for all $l \notin S_{\text{priv}}$, we have $\text{globalstore}'(l) = \mathcal{C}_{\text{globalstore}}(\mathcal{C})(l)$. \square

Definition D.2 (Deserialized OCaml values for tables) *Let Tbl be a table of type $T_1 \times \dots \times T_l$. The OCaml value that corresponds to an element (b_1, \dots, b_l) of this table is*

$$\mathbb{G}_{\text{val}T_1, \dots, T_l}(b_1, \dots, b_l) \stackrel{\text{def}}{=} (\mathbb{G}_{\text{val}T_1}(a_1), \dots, \mathbb{G}_{\text{val}T_l}(a_l)).$$

Let $t = [a_1; \dots; a_k]$ be a list of elements of table Tbl . The corresponding OCaml list is

$$\mathbb{G}_{\text{tbl}deser}(Tbl, t) \stackrel{\text{def}}{=} [\mathbb{G}_{\text{val}T_1, \dots, T_l}(a_1); \dots; \mathbb{G}_{\text{val}T_1, \dots, T_l}(a_k)].$$

Definition D.3 (Function filter) *Let E be a CryptoVerif environment, M a CryptoVerif boolean term, (x_1, \dots, x_k) a tuple of variables, and t a list of tuples of CryptoVerif values of type $T_{x_1} \times \dots \times T_{x_k}$. We let $\text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), t)$ be the list of tuples (a_1, \dots, a_k) in t such that the term M is true when the variables $x_1[\tilde{a}], \dots, x_k[\tilde{a}]$ are bound to a_1, \dots, a_k in the environment E , respectively:*

$$\begin{aligned} \text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), t) &\stackrel{\text{def}}{=} \\ &[(a_1, \dots, a_k) \in t \mid E[x_1[\tilde{a}] \mapsto a_1, \dots, x_k[\tilde{a}] \mapsto a_k], M \Downarrow \text{true}] \end{aligned}$$

Let us recall that our fold function on lists \mathbb{G}_{fold} is defined in Figure 16 as follows:

$$\mathbb{G}_{\text{fold}} \stackrel{\text{def}}{=} f \rightarrow \text{function } a \rightarrow \text{function } [] \rightarrow a \mid x :: l \rightarrow f (\text{fold } f a l) x$$

Lemma D.4 *Suppose that*

$$\begin{aligned} Th &= \langle \text{env}, \text{fold } c [] (\mathbb{G}_{\text{tbl}}(Tbl, t)), \text{stack}, \text{store} \rangle \\ c &= \text{function}[\text{env}', \\ &\quad a \rightarrow \text{function } x \rightarrow \text{try } (c' x) :: a \text{ with Match_failure} \rightarrow a] \\ c' &= \mathbb{G}_{\text{test}}((x_1, \dots, x_k), M) \end{aligned}$$

$$\begin{aligned} env' \supseteq env_{\text{prim}} \cup \{ \mathbb{G}_{\text{var}}(x) \mapsto \mathbb{G}_{\text{val}T_x}(b) \mid x[\tilde{a}] \in \text{fv}(M) \setminus \{x_1[\tilde{a}], \dots, x_k[\tilde{a}]\}, \\ E(x[\tilde{a}']) = b \} \end{aligned}$$

$$\begin{aligned} env(\text{fold}) = env'(\text{fold}) = \text{letrec}[env_0, \{\text{fold} \mapsto \mathbb{G}_{\text{fold}}\} \text{ in } \text{fold}] \\ \text{where } t \text{ is a list of CryptoVerif values of type } T_{x_1} \times \dots \times T_{x_k} \\ \text{and all occurrences of } x_1, \dots, x_k \text{ in } M \text{ have indices } \tilde{a}. \end{aligned}$$

Then $Th \rightarrow^* Th'$ such that

$$Th' = \langle env'', \mathbb{G}_{\text{tbldeser}}(\text{Tbl}, \text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), t)), \text{stack}, \text{store}' \rangle$$

for some env'' and store' such that $\text{store}' \supseteq \text{store}$.

Proof We prove this lemma by induction on the length of t .

In the base case, $t = []$, so $\mathbb{G}_{\text{tbl}}(\text{Tbl}, t) = []$, and

$$\begin{aligned} Th &= \langle env, \text{fold } c \ [] \ [], \text{stack}, \text{store} \rangle \\ &\rightarrow^* \langle env_1, (\text{match } c \text{ with } \mathbb{G}_{\text{fold}}) \ [] \ [], \text{stack}, \text{store} \rangle \\ &\quad \text{where } env_1 \stackrel{\text{def}}{=} env_0[\text{fold} \mapsto env'(\text{fold})] \\ &\rightarrow^* \langle env'', \text{match } [] \ \text{with } [] \mapsto a \mid x :: l \rightarrow f \ (\text{fold } f \ a \ l) \ x, \text{stack}, \text{store} \rangle \\ &\quad \text{where } env'' \stackrel{\text{def}}{=} env_1[f \mapsto c, a \mapsto []] \\ &\rightarrow^* Th' \stackrel{\text{def}}{=} \langle env'', [], \text{stack}, \text{store} \rangle \end{aligned}$$

For any $E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}])$, we have $\text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), []) = []$. So the lemma is correct for the base case.

In the inductive case, let $t = b :: t'$. Let $y = \mathbb{G}_{\text{tbl}}(\text{Tbl}, b)$ and $l' = \mathbb{G}_{\text{tbl}}(\text{Tbl}, t')$, so $\mathbb{G}_{\text{tbl}}(\text{Tbl}, t) = y :: l'$. Let $b = (a_1, \dots, a_k)$ and $y = (d_1, \dots, d_k)$, where each d_i corresponds to a_i . Let $(d'_1, \dots, d'_k) = \mathbb{G}_{\text{val}T_1, \dots, T_k}(b)$, where Tbl is a table of type $T_1 \times \dots \times T_k$.

$$\begin{aligned} Th &= \langle env, \text{fold } c \ [] \ (y :: l'), \text{stack}, \text{store} \rangle \\ &\rightarrow^* \langle env_1, (\text{match } c \ \text{with } \mathbb{G}_{\text{fold}}) \ [] \ (y :: l'), \text{stack}, \text{store} \rangle \\ &\quad \text{where } env_1 \stackrel{\text{def}}{=} env_0[\text{fold} \mapsto env'(\text{fold})] \\ &\rightarrow^* \langle env_2, \text{match } y :: l' \ \text{with } [] \mapsto a \mid x :: l \rightarrow f \ (\text{fold } f \ a \ l) \ x, \text{stack}, \text{store} \rangle \\ &\quad \text{where } env_2 \stackrel{\text{def}}{=} env_1[f \mapsto c, a \mapsto []] \\ &\rightarrow^* \langle env_3, f \ (\text{fold } f \ a \ l) \ x, \text{stack}, \text{store} \rangle \\ &\quad \text{where } env_3 \stackrel{\text{def}}{=} env_2[x \mapsto y, l \mapsto l'] \\ &\rightarrow^* \langle env_3, \text{fold } f \ a \ l, \text{stack}_1, \text{store} \rangle \text{ where } \text{stack}_1 \stackrel{\text{def}}{=} (env_3, f \ [\cdot] \ y) :: \text{stack} \\ &\quad \text{(arguments are evaluated from right to left)} \\ &\rightarrow^* \langle env_3, \text{fold } c \ [] \ l', \text{stack}_1, \text{store} \rangle \\ &\rightarrow^* \langle env_4, \mathbb{G}_{\text{tbldeser}}(\text{Tbl}, t'), \text{stack}_1, \text{store}_1 \rangle \quad \text{by induction hypothesis} \\ &\quad \text{where } t'' \stackrel{\text{def}}{=} \text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), t') \end{aligned}$$

$$\begin{aligned}
& \rightarrow \langle env_3, f(\mathbb{G}_{tbldeser}(Tbl, t'')) y, stack, store_1 \rangle \\
& \rightarrow^* \langle env_3, c(\mathbb{G}_{tbldeser}(Tbl, t'')) y, stack, store_1 \rangle \\
& \rightarrow^* \langle env_5, \text{try}(c' x) :: a \text{ with Match_failure} \rightarrow a, stack, store_1 \rangle \\
& \quad \text{where } env_5 \stackrel{\text{def}}{=} env'[a \mapsto \mathbb{G}_{tbldeser}(Tbl, t''), x \mapsto y] \\
& \rightarrow \langle env_5, (c' x) :: a, stack_2, store_1 \rangle \\
& \quad \text{where } stack_2 \stackrel{\text{def}}{=} (env_5, \text{try}[\cdot] \text{ with Match_failure} \rightarrow a) :: stack \\
& \rightarrow \langle env_5, c' x, stack_3, store_1 \rangle \text{ where } stack_3 \stackrel{\text{def}}{=} (env_5, [\cdot] :: a) :: stack_2 \\
& \rightarrow^* \langle env_5[\mathbb{G}_{var}(x_1) \mapsto d_1, \dots, \mathbb{G}_{var}(x_k) \mapsto d_k], \\
& \quad \text{let } \mathbb{G}_{var}(x_1) = \mathbb{G}_{deser}(T_{x_1}) \mathbb{G}_{var}(x_1) \text{ in } \dots \\
& \quad \text{let } \mathbb{G}_{var}(x_k) = \mathbb{G}_{deser}(T_{x_k}) \mathbb{G}_{var}(x_k) \text{ in} \\
& \quad \text{if } (\mathbb{G}_M(M)) \text{ then } (\mathbb{G}_{var}(x_1), \dots, \mathbb{G}_{var}(x_k)) \text{ else raise Match_failure,} \\
& \quad stack_3, store_1 \rangle \\
& \rightarrow^* \langle env_6, \\
& \quad \text{if } (\mathbb{G}_M(M)) \text{ then } (\mathbb{G}_{var}(x_1), \dots, \mathbb{G}_{var}(x_k)) \text{ else raise Match_failure,} \\
& \quad stack_3, store_2 \rangle \text{ where } env_6 \stackrel{\text{def}}{=} env_5[\mathbb{G}_{var}(x_1) \mapsto d'_1, \dots, \mathbb{G}_{var}(x_k) \mapsto d'_k] \\
& \quad \text{by Proposition 6.5 applied } k \text{ times} \\
& \rightarrow Th_1 \stackrel{\text{def}}{=} \langle env_6, \mathbb{G}_M(M), stack_4, store_2 \rangle \\
& \quad \text{where } stack_4 \stackrel{\text{def}}{=} (env_6, \text{if}[\cdot] \text{ then } (\mathbb{G}_{var}(x_1), \dots, \mathbb{G}_{var}(x_k)) \\
& \quad \quad \text{else raise Match_failure}) :: stack_3
\end{aligned}$$

The environment env_6 contains env_{prim} and $env(E[x_1[\tilde{a}] \mapsto a_1, \dots, x_k[\tilde{a}] \mapsto a_k], M)$. Let r be the CryptoVerif value such that $E[x_1[\tilde{a}] \mapsto a_1, \dots, x_k[\tilde{a}] \mapsto a_k], M \Downarrow r$. So by Lemma 6.8,

$$Th_1 \rightarrow^* Th_2 \stackrel{\text{def}}{=} \langle env_7, \mathbb{G}_{valbool}(r), stack_4, store_3 \rangle$$

and by Lemma 6.8, Proposition 6.5, and the induction hypothesis, we have $store_3 \supseteq store_2 \supseteq store_1 \supseteq store$.

Let us suppose that $r = \text{true}$. In this case, $\text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), t) = b :: \text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), t')$. Let us denote $t''' \stackrel{\text{def}}{=} \text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), t)$.

$$\begin{aligned}
Th_2 &= \langle env_7, \text{true}, stack_4, store_3 \rangle \\
& \rightarrow^* \langle env_6, \text{if true then } (\mathbb{G}_{var}(x_1), \dots, \mathbb{G}_{var}(x_k)) \text{ else raise Match_failure,} \\
& \quad stack_3, store_3 \rangle \\
& \rightarrow \langle env_6, (\mathbb{G}_{var}(x_1), \dots, \mathbb{G}_{var}(x_k)), stack_3, store_3 \rangle \\
& \rightarrow^* \langle env_6, (d'_1, \dots, d'_k), stack_3, store_3 \rangle \\
& \rightarrow \langle env_5, (d'_1, \dots, d'_k) :: a, stack_2, store_3 \rangle \\
& \rightarrow^* \langle env_5, \mathbb{G}_{tbldeser}(Tbl, t'''), stack_2, store_3 \rangle \\
& \quad \text{since } \mathbb{G}_{tbldeser}(Tbl, t''') = (d'_1, \dots, d'_k) :: (\mathbb{G}_{tbldeser}(Tbl, t''))
\end{aligned}$$

following reduction:

$$E, \text{if } M \text{ then } P_1 \text{ else } P_2, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow E, P_1, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}.$$

On the OCaml side, we implement the CryptoVerif *bool* type with booleans in OCaml, and we have $\mathbb{G}_{\text{valbool}}(\text{true}) = \text{true}$ and $\mathbb{G}_{\text{valbool}}(\text{false}) = \text{false}$. Let $Th = \mathcal{C}_{Th}(\mathcal{C})$.

$$\begin{aligned} Th &= \langle env, \text{if } \mathbb{G}_M(M) \text{ then } \mathbb{G}(P_1) \text{ else } \mathbb{G}(P_2), stack, store \rangle \\ &\rightarrow \langle env, \mathbb{G}_M(M), stack', store \rangle \\ &\quad \text{where } stack' \stackrel{\text{def}}{=} (env, \text{if } [\cdot] \text{ then } \mathbb{G}(P_1) \text{ else } \mathbb{G}(P_2)) :: stack \\ &\rightarrow^* \langle env, \text{true}, stack', store' \rangle \quad \text{by Lemma 6.8} \\ &\rightarrow \langle env, \text{if true then } \mathbb{G}(P_1) \text{ else } \mathbb{G}(P_2), stack, store' \rangle \\ &\rightarrow Th' \stackrel{\text{def}}{=} \langle env, \mathbb{G}(P_1), stack, store' \rangle \end{aligned}$$

By (Globalstore1) and (Toplevel), we obtain $\mathcal{C} \rightarrow^* \mathcal{C}' \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto Th']$. This sequence of reductions describes the set of traces \mathcal{CTS}_1 that corresponds to the CryptoVerif reduction. We have $\Pr[\mathcal{CTS}_1] = 1$.

The CryptoVerif environment E and tables \mathcal{T} , the OCaml environment env and global store $globalstore$, and the events on both sides are unchanged. By Lemma 6.8, we have $store' \supseteq store$, so this construct satisfies the lemma.

- The insert construct:

On the CryptoVerif side, let us suppose that $E, M_i \Downarrow a_i$. We have the following reduction:

$$E, \text{insert } Tbl(M_1, \dots, M_k); P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow E, P', \mathcal{T}', \mathcal{Q}, \mathcal{R}, \mathcal{E},$$

where $\mathcal{T}' \stackrel{\text{def}}{=} \mathcal{T}[Tbl \mapsto (a_1, \dots, a_k) :: \mathcal{T}(Tbl)]$. Let the type of the table Tbl be $T_1 \times \dots \times T_k$.

On the OCaml side, there exists a unique f such that $(Tbl, f) \in \text{tables}$. Let $globalstore = \mathcal{C}_{globalstore}(\mathcal{C})$. Since $globalstore \supseteq globalstore(E, \mathcal{T})$, we have

$$globalstore(f) = t \text{ where } t \stackrel{\text{def}}{=} \mathbb{G}_{tbl}(Tbl, \mathcal{T}(Tbl)).$$

Let $t' \stackrel{\text{def}}{=} \mathbb{G}_{tbl}(Tbl, (a_1, \dots, a_k)) :: t$. By definition of \mathbb{G}_{tbl} , we have $t' = \mathbb{G}_{tbl}(Tbl, \mathcal{T}'(Tbl))$. Let $globalstore' = globalstore[f \mapsto t']$.

$$\begin{aligned} \mathcal{C} &= \mathcal{C}[\text{Th} \mapsto \langle env, f := e :: (!f); \mathbb{G}(P'), stack, store \rangle] \\ &\quad \text{where } e \stackrel{\text{def}}{=} (\mathbb{G}_{\text{ser}}(T_1) \mathbb{G}_M(M_1), \dots, \mathbb{G}_{\text{ser}}(T_k) \mathbb{G}_M(M_k)) \\ &\rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env, e :: t, stack', store \rangle] \\ &\quad \text{where } stack' \stackrel{\text{def}}{=} (env, f := [\cdot]) :: (env, [\cdot]; \mathbb{G}(P')) :: stack \end{aligned}$$

$$\begin{aligned}
& \rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env, \mathbb{G}_{\text{tbl}}(Tbl, (a_1, \dots, a_k)) :: t, stack', store' \rangle] \\
& \quad \text{by Lemma 6.8 and Proposition 6.5} \\
& \rightarrow \mathcal{C}[\text{Th} \mapsto \langle env, f := t', (env, []; \mathbb{G}(P')) :: stack, store' \rangle] \\
& \rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env, () ; \mathbb{G}(P'), stack, store' \rangle, \text{globalstore} \mapsto \text{globalstore}'] \\
& \rightarrow \mathcal{C}' \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto \langle env, \mathbb{G}(P'), stack, store' \rangle, \text{globalstore} \mapsto \text{globalstore}']
\end{aligned}$$

This sequence of reductions describes the set of traces \mathcal{CTS}_1 that corresponds to the CryptoVerif reduction. We have $\Pr[\mathcal{CTS}_1] = 1$.

The global store is modified so that $\text{globalstore}' \supseteq \text{globalstore}(E, \mathcal{T}')$ and $\text{globalstore}'(l) = \mathcal{C}_{\text{globalstore}}(\mathcal{C})(l)$ for all $l \notin S_{\text{priv}}$, and the environments and events are unchanged on both sides. Moreover, by Proposition 6.5 and Lemma 6.8, we have $\text{store}' \supseteq \text{store}$, so this construct satisfies the lemma.

- The `get` construct:

On the CryptoVerif side, let us consider a CryptoVerif configuration such that its program is

$$P = \text{get } Tbl(x_1[\tilde{a}], \dots, x_k[\tilde{a}]) \text{ suchthat } M \text{ in } P' \text{ else } P''.$$

Let the type of the table Tbl be $T_1 \times \dots \times T_k$.

We have two cases depending on whether there is a value in the table Tbl that satisfies M or not. Let $l' \stackrel{\text{def}}{=} \text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), \mathcal{T}(Tbl)) = [b_1, \dots, b_m]$. This list contains every element of $\mathcal{T}(Tbl)$ such that M is true.

If l' is empty, then:

$$E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow E, P'', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}.$$

If l' is not empty, then there is a reduction for each element $b = (a_1, \dots, a_k)$ in l' ,

$$E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_{p_b} E_b, P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E},$$

where $p_b \stackrel{\text{def}}{=} \sum_{\{j \in \{1, \dots, m\} \mid b_j = b\}} \text{among}(\{1, \dots, m\}, j)$ and $E_b \stackrel{\text{def}}{=} E[x_1[\tilde{a}] \mapsto a_1, \dots, x_k[\tilde{a}] \mapsto a_k]$.

On the OCaml side, let us denote

$$\begin{aligned}
e & \stackrel{\text{def}}{=} \text{if } l = [] \text{ then } \mathbb{G}(P') \text{ else} \\
& \quad \text{let } (\mathbb{G}_{\text{var}}(x_1), \dots, \mathbb{G}_{\text{var}}(x_k)) = \text{random}_l \text{ in} \\
& \quad (\mathbb{G}_{\text{file}}(x_1[\tilde{a}]); \dots; \mathbb{G}_{\text{file}}(x_k[\tilde{a}]); \mathbb{G}(P)) \\
e_1 & \stackrel{\text{def}}{=} \text{try } (\mathbb{G}_{\text{test}}((x_1, \dots, x_k), M) \ x) :: a \text{ with Match_failure} \rightarrow a
\end{aligned}$$

There exists a unique f such that $(Tbl, f) \in \text{tables}$, and we have

$$\begin{aligned}
\mathcal{C} & = \mathcal{C}[\text{Th} \mapsto \langle env, \text{let } l = e_2 \text{ in } e, stack, store \rangle] \\
& \quad \text{where } e_2 \stackrel{\text{def}}{=} \text{read_table}(f, \mathbb{G}_{\text{test}}((x_1, \dots, x_k), M)) \\
& \quad = \text{let rec } fold = \text{function } \mathbb{G}_{\text{fold}} \text{ in} \\
& \quad \quad fold \ (\text{function } a \rightarrow x \rightarrow e_1) \ [] \ !f
\end{aligned}$$

$$\begin{aligned}
&\rightarrow \mathcal{C}[\text{Th} \mapsto \langle env, e_2, stack', store \rangle] \\
&\quad \text{where } stack'' \stackrel{\text{def}}{=} (env, \text{let } l = [\cdot] \text{ in } e) :: stack \\
&\rightarrow \mathcal{C}[\text{Th} \mapsto \langle env', fold \text{ (function } a \rightarrow x \rightarrow e_1) [] !f, stack', store \rangle] \\
&\quad \text{where } env' \stackrel{\text{def}}{=} env[fold \mapsto \text{letrec}[env, \{fold \mapsto \mathbb{G}_{\text{fold}}\} \text{ in } fold]] \\
&\rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env', e_3, stack', store \rangle] \\
&\quad \text{where } e_3 \stackrel{\text{def}}{=} \\
&\quad \quad fold \text{ function}[env', a \rightarrow \text{function } x \rightarrow e_1] [] \mathbb{G}_{\text{tbl}}(Tbl, \mathcal{T}(Tbl)) \\
&\rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env'', \mathbb{G}_{\text{tbldeser}}(Tbl, l'), stack', store' \rangle] \quad \text{by Lemma D.4} \\
&\quad \text{since } l' = \text{filter}(E, M, (x_1[\tilde{a}], \dots, x_k[\tilde{a}]), \mathcal{T}(Tbl)) \\
&\quad \rightarrow \mathcal{C}[\text{Th} \mapsto \langle env, \text{let } l = \mathbb{G}_{\text{tbldeser}}(Tbl, l') \text{ in } e, stack, store' \rangle] \\
&\rightarrow \mathcal{C}_1 \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto \langle env'', e, stack, store' \rangle] \\
&\quad \text{where } env'' \stackrel{\text{def}}{=} env[l \mapsto \mathbb{G}_{\text{tbldeser}}(Tbl, l')]
\end{aligned}$$

Now, if l' is empty, then $env''(l) = []$, so

$$\mathcal{C}_1 \rightarrow^* \mathcal{C}' \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto \langle env'', \mathbb{G}(P''), stack, store' \rangle]$$

The sequence of reductions $\mathcal{C} \rightarrow^* \mathcal{C}_1 \rightarrow^* \mathcal{C}'$ describes the set of traces \mathcal{CTS}_1 that corresponds to the unique CryptoVerif reduction that can happen when l' is empty. We have $\Pr[\mathcal{CTS}_1] = 1$.

The CryptoVerif environment E is unchanged and the OCaml environment env'' is an extension of env , so we have $env'' \supseteq env_{\text{prim}} \cup env(E, P')$. The CryptoVerif tables, the global store, and the events on both sides are unchanged. By Lemma D.4, we have $store' \supseteq store$. So, in this case, the get construct satisfies the lemma.

If l' is not empty, then let $b = (a_1, \dots, a_k)$ be any element of l' , and let $v = \mathbb{G}_{\text{val}_{T_1, \dots, T_k}}(Tbl, b) = (\mathbb{G}_{\text{val}_{T_1}}(a_1), \dots, \mathbb{G}_{\text{val}_{T_k}}(a_k))$. We have $env''(l) = \mathbb{G}_{\text{tbldeser}}(Tbl, l')$. Let $env''(l) = [v_1; \dots; v_m]$. The set $S \stackrel{\text{def}}{=} \{j \in \{1, \dots, m\} \mid v = v_j\}$ is equal to the set $\{j \in \{1, \dots, m\} \mid b = b_j\}$, because the function $b \mapsto \mathbb{G}_{\text{val}_{T_1, \dots, T_k}}(Tbl, b)$ is injective. Hence, we have $p_b = \sum_{j \in S} \text{among}(\{1, \dots, m\}, j)$.

$$\begin{aligned}
&\mathcal{C}_1 \rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env'', e_4, stack, store' \rangle] \\
&\quad \text{where } e_4 \stackrel{\text{def}}{=} \text{let } (\mathbb{G}_{\text{var}}(x_1), \dots, \mathbb{G}_{\text{var}}(x_k)) = \text{random}_l \text{ in} \\
&\quad \quad (\mathbb{G}_{\text{file}}(x_1[\tilde{a}]); \dots; \mathbb{G}_{\text{file}}(x_k[\tilde{a}]); \mathbb{G}(P')) \\
&\rightarrow_{p_b}^* \mathcal{C}[\text{Th} \mapsto \langle env'', e_5, stack, store'' \rangle] \quad \text{by Proposition 6.5} \\
&\quad \text{where } e_5 \stackrel{\text{def}}{=} \text{let } (\mathbb{G}_{\text{var}}(x_1), \dots, \mathbb{G}_{\text{var}}(x_k)) = v \text{ in} \\
&\quad \quad (\mathbb{G}_{\text{file}}(x_1[\tilde{a}]); \dots; \mathbb{G}_{\text{file}}(x_k[\tilde{a}]); \mathbb{G}(P'))
\end{aligned}$$

$$\begin{aligned}
& \rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env''', \mathbb{G}_{\text{file}}(x_1[\tilde{a}]); \dots; \mathbb{G}_{\text{file}}(x_k[\tilde{a}]); \mathbb{G}(P'), stack, store'' \rangle] \\
& \quad \text{where } env''' \stackrel{\text{def}}{=} \\
& \quad \quad env''[\mathbb{G}_{\text{var}}(x_1) \mapsto \mathbb{G}_{\text{val}T_1}(a_1), \dots, \mathbb{G}_{\text{var}}(x_k) \mapsto \mathbb{G}_{\text{val}T_k}(a_k)] \\
& \rightarrow^* \mathcal{C}'_b \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto \langle env''', \mathbb{G}(P'), stack, store''' \rangle, \text{globalstore} \mapsto \text{globalstore}'] \\
& \quad \quad \quad \text{by Lemma D.1}
\end{aligned}$$

The sequence of reductions $\mathcal{C} \rightarrow^* \mathcal{C}_1 \rightarrow^* \mathcal{C}'_b$ describes the set of traces \mathcal{CTS}_b that corresponds to the CryptoVerif reduction in which the element $b = (a_1, \dots, a_k)$ of l' is chosen. We have $\Pr[\mathcal{CTS}_b] = p_b$.

By Lemma D.1, $\text{globalstore}' \supseteq \text{globalstore}(E_b, \mathcal{T})$, and $\text{globalstore}'$ and $\mathcal{C}_{\text{globalstore}}(\mathcal{C})$ are equal on all locations not in S_{priv} , since $E_b = E[x_1[\tilde{a}] \mapsto a_1, \dots, x_k[\tilde{a}] \mapsto a_k]$. Since the OCaml environment is $env''' = env[l \mapsto \dots, \mathbb{G}_{\text{var}}(x_1) \mapsto \mathbb{G}_{\text{val}T_1}(a_1), \dots, \mathbb{G}_{\text{var}}(x_k) \mapsto \mathbb{G}_{\text{val}T_k}(a_k)]$, we have $env''' \supseteq env_{\text{prim}} \cup env(E_b, P')$. The events are unchanged on both sides. By Lemma D.1, Proposition 6.5, and Lemma D.4, we have $store''' \supseteq store'' \supseteq store' \supseteq store$. So, in this case, the get construct also satisfies the lemma.

- The event construct:

On the CryptoVerif side, let us suppose that $E, M_j \Downarrow a_j$ for all $j \leq l$. We have the following reduction:

$$E, \text{event } ev(M_1, \dots, M_l); P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow E, P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}',$$

where $\mathcal{E}' \stackrel{\text{def}}{=} ev(a_1, \dots, a_l) :: \mathcal{E}$. Let us denote $T_1 \times \dots \times T_l$ the type of the event ev .

On the OCaml side, let us denote

$$events' \stackrel{\text{def}}{=} \mathbb{G}_{\text{ev}}(\mathcal{E}') = ev(\mathbb{G}_{\text{val}T_1}(a_1), \dots, \mathbb{G}_{\text{val}T_l}(a_l)) :: events.$$

We have

$$\begin{aligned}
\mathcal{C} &= \mathcal{C}[\text{Th} \mapsto \langle env, e; \mathbb{G}(P'), stack, store \rangle] \\
& \quad \text{where } e \stackrel{\text{def}}{=} \text{event } ev(\mathbb{G}_{\text{M}}(M_1), \dots, \mathbb{G}_{\text{M}}(M_l)) \\
& \rightarrow \mathcal{C}[\text{Th} \mapsto \langle env, e, stack', store, \rangle] \\
& \quad \text{where } stack' \stackrel{\text{def}}{=} (env, [\cdot]; \mathbb{G}(P')) :: stack \\
& \rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env, e', stack', store' \rangle] \quad \text{by Lemma 6.8} \\
& \quad \text{where } e' \stackrel{\text{def}}{=} \text{event } ev(\mathbb{G}_{\text{val}T_1}(a_1), \dots, \mathbb{G}_{\text{val}T_l}(a_l)) \\
& \rightarrow \mathcal{C}[\text{Th} \mapsto \langle env, (), stack', store' \rangle, \text{events} \mapsto events'] \\
& \rightarrow^* \mathcal{C}' \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto \langle env, \mathbb{G}(P'), stack, store' \rangle, \text{events} \mapsto events']
\end{aligned}$$

This sequence of reductions describes the set of traces \mathcal{CTS}_1 that corresponds to the CryptoVerif reduction. We have $\Pr[\mathcal{CTS}_1] = 1$.

The CryptoVerif environment E and tables \mathcal{T} and the OCaml environment env and global store $globalstore$ are unchanged. We have $events' = \mathbb{G}_{ev}(\mathcal{E}')$. By Lemma 6.8, we have $store' \supseteq store$, so this construct satisfies the lemma. \square

E Proof of Lemma 6.16 and Proposition 6.17

Proof (of Lemma 6.16) Let us first show by induction on $steps$ that, if $\mathcal{CS} \rightarrow^* \mathcal{CS}'$ in at most $steps$ steps and \mathcal{CS}' does not reduce, or $\mathcal{CS} \rightarrow^* \mathcal{CS}'$ in exactly $steps$ steps, $\mathfrak{C} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$, and $P = x[a] \leftarrow simulate(s[a]); P'$, then $\mathfrak{C}, steps, \mathcal{CS} \rightsquigarrow^* E[x[a] \mapsto simreturn(\mathcal{CS}')] , P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$.

- If $steps = 0$ or \mathcal{CS} does not reduce, then $\mathcal{CS}' = \mathcal{CS}$ and $\mathfrak{C}, steps, \mathcal{CS} \rightsquigarrow E[x[a] \mapsto simreturn(\mathcal{CS}')] , P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$ by (Leave Simulator).
- If $steps > 0$ and $\mathcal{CS} \rightarrow \mathcal{CS}_1$, then $\mathcal{CS}_1 \rightarrow^* \mathcal{CS}'$ in at most $steps - 1$ steps and \mathcal{CS}' does not reduce, or $\mathcal{CS}_1 \rightarrow^* \mathcal{CS}'$ in exactly $steps - 1$ steps, so by induction hypothesis, $\mathfrak{C}, steps - 1, \mathcal{CS}_1 \rightsquigarrow^* E[x[a] \mapsto simreturn(\mathcal{CS})] , P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$. By (Simulator), $\mathfrak{C}, steps, \mathcal{CS} \rightsquigarrow \mathfrak{C}, steps - 1, \mathcal{CS}_1$, so $\mathfrak{C}, steps, \mathcal{CS} \rightsquigarrow^* E[x[a] \mapsto simreturn(\mathcal{CS}')] , P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$.

Let us now prove that, if $\mathfrak{C} \rightarrow_p \mathfrak{C}'$, then there is a trace $\mathfrak{C} \rightsquigarrow_p^* \mathfrak{C}'$ and all intermediate configurations in this trace (if any) are of the form $\mathfrak{C}, steps, \mathcal{CS}$. Let $\mathfrak{C} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$.

- If P is not of the form $x[a] \leftarrow simulate(s[a]); P'$ for any x, a, P' , then $\mathfrak{C} \rightsquigarrow_p \mathfrak{C}'$ by (CryptoVerif).
- If $P = x[a] \leftarrow simulate(s[a]); P'$ for some x, a, P' , then by the semantics of CryptoVerif, $s[a] \in \text{Dom}(E)$, $E(s[a])$ is of type $T_{\mathcal{CS}}$, and $\mathfrak{C}' = E[x[a] \mapsto simulate(E(s[a]))] , P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$. Since $E(s[a])$ is of type $T_{\mathcal{CS}}$, there exists a configuration \mathcal{CS} such that $E(s[a]) = \text{repr}(\mathcal{CS})$. By reduction rule (Enter Simulator), $\mathfrak{C} \rightsquigarrow \mathfrak{C}_1^{cs} = \mathfrak{C}, N_{steps}, \mathcal{CS}$. Moreover, by definition of $simulate$, $\mathcal{CS} \rightarrow^* \mathcal{CS}'$ in at most N_{steps} steps and \mathcal{CS}' does not reduce, or $\mathcal{CS} \rightarrow^* \mathcal{CS}'$ in exactly N_{steps} steps, and $simulate(\text{repr}(\mathcal{CS})) = simreturn(\mathcal{CS}')$. By the result shown above, $\mathfrak{C} \rightsquigarrow \mathfrak{C}_1^{cs} \rightsquigarrow^* E[x[a] \mapsto simreturn(\mathcal{CS}')] , P', \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} = \mathfrak{C}'$.

Finally, let us show that, if \mathfrak{C} does not reduce by \rightarrow , then it does not reduce by \rightsquigarrow either. Let $\mathfrak{C} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$.

- If P is not of the form $x[a] \leftarrow simulate(s[a]); P'$ for any x, a, P' , then the only rule applicable to reduce \mathfrak{C} by \rightsquigarrow is (CryptoVerif), and it cannot be applied because \mathfrak{C} does not reduce by \rightarrow . Hence \mathfrak{C} does not reduce by \rightsquigarrow .
- If $P = x[a] \leftarrow simulate(s[a]); P'$ for some x, a, P' , then either $s[a] \notin \text{Dom}(E)$ or $E(s[a]) \notin T_{\mathcal{CS}}$. The only rule applicable to reduce \mathfrak{C} by \rightsquigarrow is (Enter Simulator), and it does not apply when $s[a] \notin \text{Dom}(E)$ or $E(s[a]) \notin T_{\mathcal{CS}}$.

T_{CS} . Hence \mathfrak{C} does not reduce by \rightsquigarrow . (We could also show that, because the CryptoVerif configurations are well-typed, \mathfrak{C} always reduces when $P = x[a] \leftarrow \text{simulate}(s[a]); P'$.) \square

Proof (of Proposition 6.17) For $b \in \{\text{true}, \text{false}\}$, let $\mathfrak{CT}\mathfrak{S}_b$ be the set of complete CryptoVerif traces using \rightarrow starting at \mathfrak{C} and such that the list of events \mathcal{E} in their last configuration satisfies $D(\mathcal{E}) = b$. By the first property of Lemma 6.16, we can map each trace $\mathfrak{CT} \in \mathfrak{CT}\mathfrak{S}_b$ into a trace $\mathfrak{CT}^{\text{cs}}$ using \rightsquigarrow and starting at \mathfrak{C} , such that the configurations of the form \mathfrak{C} of $\mathfrak{CT}^{\text{cs}}$ are exactly the same as in \mathfrak{CT} and $\Pr[\mathfrak{CT}^{\text{cs}}] = \Pr[\mathfrak{CT}]$.

Let $\mathfrak{CT}\mathfrak{S}_b^{\text{cs}}$ be the set of these traces $\mathfrak{CT}^{\text{cs}}$. Let us show that $\mathfrak{CT}\mathfrak{S}_b^{\text{cs}}$ is the set of complete CryptoVerif traces using \rightsquigarrow starting at \mathfrak{C} and such that the list of events \mathcal{E} in their last configuration satisfies $D(\mathcal{E}) = b$.

The list of events \mathcal{E} in the last configuration of $\mathfrak{CT}^{\text{cs}}$ is the same as in \mathfrak{CT} , so it satisfies $D(\mathcal{E}) = b$. By the second property of Lemma 6.16, since \mathfrak{CT} is complete, $\mathfrak{CT}^{\text{cs}}$ is also complete. Since the mapping from \mathfrak{CT} to $\mathfrak{CT}^{\text{cs}}$ is injective, we have $\Pr[\mathfrak{CT}\mathfrak{S}_b^{\text{cs}}] = \Pr[\mathfrak{CT}\mathfrak{S}_b]$.

Moreover, if a configuration \mathfrak{C}^{cs} reduces by \rightsquigarrow into another configuration, then the sum of the probabilities of all the possible reductions from \mathfrak{C}^{cs} is 1:

$$\sum_{\{\mathfrak{C}^{\text{cs}'}/\mathfrak{C}^{\text{cs}} \rightsquigarrow_{p(\mathfrak{C}^{\text{cs}'})} \mathfrak{C}^{\text{cs}'}\}} p(\mathfrak{C}^{\text{cs}'}) = 1.$$

Indeed, the rules that define \rightsquigarrow are mutually exclusive. If \mathfrak{C}^{cs} reduces by rule (CryptoVerif), then the property holds because it holds for the semantics of CryptoVerif. Otherwise, a single reduction is possible, and it has probability 1.

Using the same property for \rightarrow , the probability of all complete traces using \rightarrow starting from \mathfrak{C} is 1, so $\Pr[\mathfrak{CT}\mathfrak{S}_{\text{true}}] + \Pr[\mathfrak{CT}\mathfrak{S}_{\text{false}}] = 1$. So $\Pr[\mathfrak{CT}\mathfrak{S}_{\text{true}}^{\text{cs}}] + \Pr[\mathfrak{CT}\mathfrak{S}_{\text{false}}^{\text{cs}}] = 1$. Since the sum of the probabilities of all the possible reductions from each configuration by \rightsquigarrow is 1, the probability of all complete traces using \rightsquigarrow starting from \mathfrak{C} is 1, so all these traces are in $\mathfrak{CT}\mathfrak{S}_{\text{true}}^{\text{cs}}$ or $\mathfrak{CT}\mathfrak{S}_{\text{false}}^{\text{cs}}$. Hence all complete CryptoVerif traces using \rightsquigarrow starting at \mathfrak{C} and such that the list of events \mathcal{E} in their last configuration satisfies $D(\mathcal{E}) = b$ are in $\mathfrak{CT}\mathfrak{S}_b^{\text{cs}}$.

So $\Pr[\mathfrak{C}(\rightsquigarrow) : D] = \Pr[\mathfrak{CT}\mathfrak{S}_{\text{true}}^{\text{cs}}] = \Pr[\mathfrak{CT}\mathfrak{S}_{\text{true}}] = \Pr[\mathfrak{C} : D]$. \square

F Proof of Lemmas 6.28 and 6.29

Proof (of Lemma 6.28) Let $\mathfrak{C}_0^{\text{cs}} \stackrel{\text{def}}{=} \mathfrak{C}^0(Q_0, \text{program}_0)$.

$\mathfrak{C}_0^{\text{cs}} = \emptyset$, let $x[] : \text{bitstring} = O_{\text{start}}()$ in $\text{return}(x[])$ else end, $\mathcal{T}_0, \mathcal{Q}_0, \emptyset, []$

where $\mathcal{Q}_0 \stackrel{\text{def}}{=} \{Q_{\text{start}}(Q_0, \text{program}_0)\} \cup \bigcup_{a \leq N_{\text{rand}} + \text{calls}} \{Q_{\text{loop}}\{a/i'\}\}$
 $\cup \text{reduce}(Q_0)$

$$\begin{aligned}
& \rightsquigarrow \emptyset, P_1, \mathcal{T}_0, \mathcal{Q}_1, \mathcal{R}_1, [] \\
& \quad \text{where } \mathcal{Q}_1 \stackrel{\text{def}}{=} \mathcal{Q}_0 \setminus \{Q_{\text{start}}(Q_0, \text{program}_0)\}, \\
& \quad \quad P_1 \stackrel{\text{def}}{=} s_0[] \leftarrow s_0(Q_0, \text{program}_0); P_2, \\
& \quad \quad P_2 \stackrel{\text{def}}{=} \text{let } r[] : T_{CS} = \text{loop } O_{\text{loop}}[1](s_0) \text{ in end else end,} \\
& \quad \quad \mathcal{R}_1 \stackrel{\text{def}}{=} [x[], \text{return}(x[]), \text{end}] \\
& \rightsquigarrow E_1, P_2, \mathcal{T}_0, \mathcal{Q}_1, \mathcal{R}_1, [] \\
& \quad \text{where } E_1 \stackrel{\text{def}}{=} \{s_0[] \mapsto s_0(Q_0, \text{program}_0)\} \\
& \rightsquigarrow E_1, P_3, \mathcal{T}_0, \mathcal{Q}_1, \mathcal{R}_1, [] \\
& \quad \text{where } P_3 \stackrel{\text{def}}{=} \text{let } (r'_{1,r}[] : T_{CS}, b_{1,r}[] : \text{bool}) = O_{\text{loop}}[1](s_0) \\
& \quad \quad \quad \text{in } P_{\text{return-loop}}(1) \text{ else end} \\
& \rightsquigarrow E_2, P_{\text{loop}}\{1/i'\}, \mathcal{T}_0, \mathcal{Q}_2, \mathcal{R}_{\text{loop}}(1), [] \\
& \quad \text{where } E_2 \stackrel{\text{def}}{=} E_1[s[1] \mapsto s_0(Q_0, \text{program}_0)], \\
& \quad \quad \mathcal{Q}_2 \stackrel{\text{def}}{=} \mathcal{Q}_1 \setminus \{Q_{\text{loop}}\{1/i'\}\}, \\
& \rightsquigarrow \mathfrak{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E_2, P_{\text{loop}}\{1/i'\}, \mathcal{T}_0, \mathcal{Q}_2, \mathcal{R}_{\text{loop}}(1), [], N_{\text{steps}}, CS_0 \\
& \quad \text{where } CS_0 \stackrel{\text{def}}{=} ([\langle \emptyset, \text{program}_0, [], \emptyset \rangle], \text{globalstore}_0, 1), \mathcal{RI}_0, \emptyset
\end{aligned}$$

We have

$$C_0(Q_0, \text{program}_0) = [\langle \emptyset, \text{program}_0, [], \emptyset \rangle, \text{globalstore}_0, 1, \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q_0), []].$$

Let \mathcal{CT} be the trace that consists only of the configuration $C_0(Q_0, \text{program}_0)$. Let us prove that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}$. Properties 1, 2, and 4 hold. The set $\mathcal{O}(\mathcal{RI}_0)$ contains all the oracles that can be called at the beginning, and

$$\mathcal{Q}_2 = \bigcup_{2 \leq a \leq N_{\text{rand}} + \text{calls}} \{Q_{\text{loop}}\{a/i'\}\} \cup \text{reduce}(Q_0),$$

so Property 3 holds. As mentioned in Section 3.2.6, the initial program program_0 does not contain locations in S_l , so Property 5 holds. As also mentioned in Section 3.2.6, program_0 contains no closure, and as mentioned in Section 5, program_0 contains no tagged function, no return, and no event except in parts $\text{program}(\mu_{\text{role}})$ inside `addthread`. So Property 6(b) holds. By Assumption 5.1, Property 7 holds. The global store globalstore_0 maps each $l \in S_g$ to its initial value initval_l and $\text{globalstore}(E_2, \mathcal{T}_0)$ maps each $f \in S_{\text{priv}}$ to its initial value initval_f (the empty string "" when $(x[], f) \in \text{files}$ and the empty list [] when $(\text{Tbl}, f) \in \text{tables}$), so Properties 8, 9, and 10 hold. The module set $\mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q_0)$ and the role set \mathcal{RI}_0 correspond by definition of \mathcal{RI}_0 , so Property 11 holds. The event lists are empty on both sides, so Property 12 holds. The sets $\mathcal{O}^\infty(\mathcal{I})$ with $\mathcal{I} \stackrel{\text{def}}{=} \emptyset$ and $\mathcal{O}_{\text{call}}(CS_0)$ are both empty, so Property 13 holds. The sets $\mathcal{O}_{\text{call-repl}}(Th_1^s)$, $\mathcal{O}^\infty(\mathcal{R}_{\text{init-function}}(Th_1^s))$, and $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(Th_1^s))$ where Th_1^s is the current thread of CS_0 are also empty, so Property 14 holds. We have $0 +$

$N_{\text{steps}} \geq N_{\text{steps}}$, so Property 15 holds. We have $\alpha = 1$, so Property 16 holds. The set \mathcal{I} is empty, so Property 17 holds. For all $\text{role}[[a', +\infty[, \tilde{a}]] \in \mathcal{R}\mathcal{I}_0$, we have $a' = 1$, so Property 18 holds. Therefore, $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}$. \square

The following two lemmas serve to prove Property 4 of the invariant.

Lemma F.1 *If $E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E} \rightarrow_p E', P', \mathcal{Q}', \mathcal{T}', \mathcal{R}', \mathcal{E}'$, and $\text{fv}(P) \cup \text{fv}(\mathcal{Q}) \cup \text{fv}(\mathcal{R}) \subseteq \text{Dom}(E)$, then $\text{fv}(P') \cup \text{fv}(\mathcal{Q}') \cup \text{fv}(\mathcal{R}') \subseteq \text{Dom}(E')$.*

Proof This result is easily proved by cases on the applied reduction rule. \square

We denote by $\mathfrak{C}^{\text{cs}} = E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}$, *rest* an extended CryptoVerif configuration in which *rest* is either nothing or *steps*, \mathcal{CS} .

Lemma F.2 *If $E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, \text{rest} \rightsquigarrow_p E', P', \mathcal{Q}', \mathcal{T}', \mathcal{R}', \mathcal{E}', \text{rest}'$ and $\text{fv}(P) \cup \text{fv}(\mathcal{Q}) \cup \text{fv}(\mathcal{R}) \subseteq \text{Dom}(E)$, then $\text{fv}(P') \cup \text{fv}(\mathcal{Q}') \cup \text{fv}(\mathcal{R}') \subseteq \text{Dom}(E')$.*

Proof This result is easily proved by cases on the applied reduction rule. By Lemma F.1, the rule (CryptoVerif) preserves the invariant. The rules (Enter Simulator) and (Simulator) leave the environment and the set of free variables unchanged. The rule (Leave Simulator) introduces a new free variable and adds it to the environment. \square

The following lemma shows that a correct closure always remains correct during execution.

Lemma F.3 *Suppose that $\text{fv}(\mathcal{Q}_0) \subseteq \text{Dom}(E)$, $\mathcal{Q}_0 \leftrightarrow \mathcal{R}\mathcal{I}, \mathcal{I}$, $\mathcal{Q}'_0 \leftrightarrow \mathcal{R}\mathcal{I}', \mathcal{I}'$, R is an oracle reference of the form $O'[a']$ when oracle O' is not under replication and $O'[_-, \tilde{a}'']$ when O' is under replication, and one of the following two situations occurs:*

1. $E' \supseteq E$, $\mathcal{I}' = \mathcal{I} - \{O[\tilde{a}]\}$, $\mathcal{Q}'_0 \supseteq \mathcal{Q}_0 \setminus \{\mathcal{Q}_0(O[\tilde{a}])\}$, $\tau'_0 = \tau_0$, and l'_{tok} is a restriction of l_{tok} such that, if $R = O'[a']$, then $R \in \text{Dom}(l'_{\text{tok}})$.
2. $E' \supseteq E$, $\mathcal{I}' \supseteq \mathcal{I}$, $\mathcal{Q}'_0 \supseteq \mathcal{Q}_0$, $l'_{\text{tok}} \supseteq l_{\text{tok}}$, $\tau'_0 \supseteq \tau_0$, if $R = O'[\tilde{a}']$, then $O'[\tilde{a}'] \notin \mathcal{I}' \setminus \mathcal{I}$ and $O'[\tilde{a}'] \in \text{Dom}(l_{\text{tok}})$, and if $R = O'[_-, \tilde{a}'']$, then for all a , $O'[[a, +\infty[, \tilde{a}'']] \notin \mathcal{I}' \setminus \mathcal{I}$ and $O'[_-, \tilde{a}''] \in \text{Dom}(\tau_0)$.

Then $\text{correctclosure}(R, \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_0) \supseteq \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_0)$.

Proof Let us consider the first situation.

- Case $R = O[\tilde{a}]$. Oracle O is not under replication. Since $\mathcal{I} - \{O[\tilde{a}]\}$ is defined, we have $O[\tilde{a}] \in \mathcal{I}$, and since $\mathcal{I}' = \mathcal{I} - \{O[\tilde{a}]\}$, we have $O[\tilde{a}] \notin \mathcal{I}'$. We also have $l'_{\text{tok}}(O[\tilde{a}]) = l_{\text{tok}}(O[\tilde{a}])$. So, by definition of correctclosure ,

$$\begin{aligned}
& \text{correctclosure}(O[\tilde{a}], \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_0) \\
&= \{\text{tagfunction}^{O, \tau}[\text{env}, \text{pm}_{\text{false}}(Q)] \mid \\
&\quad \text{for any } Q, \text{env}(\text{token}) = l'_{\text{tok}}(O[\tilde{a}])\} \\
&\supseteq \{\text{tagfunction}^{O, \tau}[\text{env}, \text{pm}_{\text{false}}(\mathcal{Q}(O[\tilde{a}]))] \mid \\
&\quad \text{env} \supseteq \text{env}_{\text{prim}} \cup \text{env}(E, \mathcal{Q}(O[\tilde{a}]), \text{env}(\text{token}) = l_{\text{tok}}(O[\tilde{a}])\} \\
&\supseteq \text{correctclosure}(O[\tilde{a}], \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_0).
\end{aligned}$$

- Case $R = O[_-, \widetilde{a}'']$ where $\widetilde{a} = a', \widetilde{a}''$ for some a' . Oracle O is under replication. Since $\mathcal{I} - \{O[\widetilde{a}]\}$ is defined, $O[[a', +\infty[, \widetilde{a}'']] \in \mathcal{I}$, and since $\mathcal{I}' = \mathcal{I} - \{O[\widetilde{a}]\}$, we have $O[[a' + 1, +\infty[, \widetilde{a}'']] \in \mathcal{I}'$.

Suppose that $a' < N_O$. Since $\mathcal{Q}_0 \leftrightarrow \mathcal{R}\mathcal{I}, \mathcal{I}$, there exist Q and i such that i does not occur in $\text{fv}(Q)$, $\mathcal{Q}_0(O[a', \widetilde{a}'']) = Q\{a'/i\}$, and $\mathcal{Q}_0(O[a' + 1, \widetilde{a}'']) = Q\{a' + 1/i\}$. It is easy to see that $pm_{\text{true}}(Q\{a' + 1/i\}) = pm_{\text{true}}(Q\{a'/i\})$, since the translation into OCaml does not depend on the indices. Moreover, $\text{fv}(Q\{a' + 1/i\}) = \text{fv}(Q\{a'/i\})$ since i does not occur in $\text{fv}(Q)$, so $\text{env}(E, Q\{a' + 1/i\}) = \text{env}(E, Q\{a'/i\})$. Since $\text{fv}(\mathcal{Q}_0) \subseteq \text{Dom}(E)$ and E' is an extension of E , we have $\text{env}(E', Q\{a' + 1/i\}) = \text{env}(E, Q\{a' + 1/i\})$. So, by definition of `correctclosure`,

$$\begin{aligned}
& \text{correctclosure}(O[_-, \widetilde{a}''], \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_O) \\
&= \{ \text{tagfunction}^{O, \tau}[\text{env}, pm_{\text{true}}(\mathcal{Q}'_0(O[a' + 1, \widetilde{a}'']))] \mid \\
&\quad \tau = \tau'_O(O[_-, \widetilde{a}'']), \text{env} \supseteq \text{env}_{\text{prim}} \cup \text{env}(E', \mathcal{Q}'_0(O[a' + 1, \widetilde{a}''])) \} \\
&= \{ \text{tagfunction}^{O, \tau}[\text{env}, pm_{\text{true}}(\mathcal{Q}_0(O[a', \widetilde{a}'']))] \mid \\
&\quad \tau = \tau_O(O[_-, \widetilde{a}'']), \text{env} \supseteq \text{env}_{\text{prim}} \cup \text{env}(E, \mathcal{Q}_0(O[a', \widetilde{a}''])) \} \\
&= \text{correctclosure}(O[_-, \widetilde{a}''], \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_O).
\end{aligned}$$

Suppose that $a' = N_O$. By definition of `correctclosure`,

$$\begin{aligned}
& \text{correctclosure}(O[_-, \widetilde{a}''], \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_O) \\
&= \{ \text{tagfunction}^{O, \tau}[\text{env}, pm_{\text{true}}(Q)] \mid \tau = \tau'_O(O[_-, \widetilde{a}'']), \text{for any } Q, \text{env} \} \\
&\supseteq \{ \text{tagfunction}^{O, \tau}[\text{env}, pm_{\text{true}}(\mathcal{Q}_0(O[a', \widetilde{a}'']))] \mid \\
&\quad \tau = \tau_O(O[_-, \widetilde{a}'']), \text{env} \supseteq \text{env}_{\text{prim}} \cup \text{env}(E, \mathcal{Q}_0(O[a', \widetilde{a}''])) \} \\
&\supseteq \text{correctclosure}(O[_-, \widetilde{a}''], \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_O).
\end{aligned}$$

Suppose that $a' > N_O$. We have $\text{correctclosure}(O[_-, \widetilde{a}''], \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_O) = \text{correctclosure}(O[_-, \widetilde{a}''], \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_O)$ since $E, \mathcal{Q}_0, l_{\text{tok}}$ are not used and $\tau'_O = \tau_O$.

- Other cases. All references to $\mathcal{Q}'_0(O'[\widetilde{a}'])$ in the definition of `correctclosure` satisfy $O'[\widetilde{a}'] \neq O[\widetilde{a}]$. In this case, we have $\mathcal{Q}'_0(O'[\widetilde{a}']) = \mathcal{Q}_0(O'[\widetilde{a}'])$. Since $\text{fv}(\mathcal{Q}_0) \subseteq \text{Dom}(E)$ and E' is an extension of E , we have $\text{env}(E', \mathcal{Q}'_0(O'[\widetilde{a}'])) = \text{env}(E', \mathcal{Q}_0(O'[\widetilde{a}'])) = \text{env}(E, \mathcal{Q}_0(O'[\widetilde{a}']))$. Moreover, when $R = O'[\widetilde{a}']$, we have $l'_{\text{tok}}(O'[\widetilde{a}']) = l_{\text{tok}}(O'[\widetilde{a}'])$. Hence, by definition of `correctclosure`, $\text{correctclosure}(R, \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_O) = \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_O)$.

Let us now consider the second situation.

- Case $R = O'[\tilde{a}']$. Since $O'[\tilde{a}'] \notin \mathcal{I}' \setminus \mathcal{I}$, we have $O'[\tilde{a}'] \in \mathcal{I}'$ if and only if $O'[\tilde{a}'] \in \mathcal{I}$. We have $l'_{\text{tok}}(O'[\tilde{a}']) = l_{\text{tok}}(O'[\tilde{a}'])$.

If $O'[\tilde{a}'] \notin \mathcal{I}$, these points are sufficient to conclude that $\text{correctclosure}(R, \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_0) = \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_0)$.

If $O'[\tilde{a}'] \in \mathcal{I}$, there is an oracle $O'[\tilde{a}']$ in \mathcal{Q}_0 ; since $\mathcal{Q}'_0 \supseteq \mathcal{Q}_0$, we have $\mathcal{Q}'_0(O'[\tilde{a}']) = \mathcal{Q}_0(O'[\tilde{a}'])$. Since $\text{fv}(\mathcal{Q}_0) \subseteq \text{Dom}(E)$ and E' is an extension of E , we have $\text{env}(E', \mathcal{Q}'_0(O'[\tilde{a}'])) = \text{env}(E', \mathcal{Q}_0(O'[\tilde{a}'])) = \text{env}(E, \mathcal{Q}_0(O'[\tilde{a}']))$. So $\text{correctclosure}(R, \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_0) = \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_0)$.

- Case $R = O'[_\perp, \tilde{a}'']$. Since for all a , $O'[[a, +\infty[, \tilde{a}'']] \notin \mathcal{I}' \setminus \mathcal{I}$, we have $O'[[a', +\infty[, \tilde{a}'']] \in \mathcal{I}'$ if and only if $O'[[a', +\infty[, \tilde{a}'']] \in \mathcal{I}$.

If there is no a' such that $O'[[a', +\infty[, \tilde{a}'']] \in \mathcal{I}$, this point is sufficient to conclude that $\text{correctclosure}(R, \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_0) = \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_0)$.

If $O'[[a', +\infty[, \tilde{a}'']] \in \mathcal{I}$ and $a' \leq N_{O'}$, then there is an oracle $O'[a', \tilde{a}'']$ in \mathcal{Q}_0 ; since $\mathcal{Q}'_0 \supseteq \mathcal{Q}_0$, we have $\mathcal{Q}'_0(O'[a', \tilde{a}'']) = \mathcal{Q}_0(O'[a', \tilde{a}''])$. Since $\text{fv}(\mathcal{Q}_0) \subseteq \text{Dom}(E)$ and E' is an extension of E , we obtain $\text{env}(E', \mathcal{Q}'_0(O'[a', \tilde{a}''])) = \text{env}(E', \mathcal{Q}_0(O'[a', \tilde{a}''])) = \text{env}(E, \mathcal{Q}_0(O'[a', \tilde{a}'']))$. Moreover, we have $\tau'_0(O'[_\perp, \tilde{a}'']) = \tau_0(O'[_\perp, \tilde{a}''])$. So $\text{correctclosure}(R, \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_0) = \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_0)$.

If $O'[[a', +\infty[, \tilde{a}'']] \in \mathcal{I}$ and $a' > N_{O'}$, then we have $\tau'_0(O'[_\perp, \tilde{a}'']) = \tau_0(O'[_\perp, \tilde{a}''])$, so $\text{correctclosure}(R, \mathcal{I}', E', \mathcal{Q}'_0, l'_{\text{tok}}, \tau'_0) = \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}_0, l_{\text{tok}}, \tau_0)$. \square

Let P'_{loop} be the process from line 8 to line 20 of Figure 19. Let P^j_{loop} be the process from line 13 to line 15 for the if $o = o_j$ branch. Let P^R_{loop} be the process from line 19 to line 20. The expansion of the let construct with pattern-matching introduces a fresh variable. Let us denote $xs[i']$ the variable created for the let matching on line 7, $xa_j[i']$ and $xi_j[i']$ the variables created on lines 11 and 12 for oracle number j .

Proof (of Lemma 6.29) Let us consider \mathfrak{C}^{cs} and \mathcal{CT} such that $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$. Let $\mathfrak{C}^{\text{cs}} = E, P_{\text{loop}}\{\alpha/i'\}, \mathcal{T}, \mathcal{Q}, \mathcal{R}_{\text{loop}}(\alpha), \mathcal{E}, \text{steps}, \mathcal{CS}$ and let \mathcal{C} be the last configuration of \mathcal{CT} . Let

$$\begin{aligned} \mathcal{CS} &= ([Th_1, \dots, Th_n], \text{globalstore}^s, tj), \mathcal{RI}, \mathcal{I}, \\ \mathcal{C} &= [Th'_1, \dots, Th'_n], \text{globalstore}^o, tj, \mathcal{MI}, \text{events}, \\ Th_{tj} &= Th^s = \langle \text{env}^s, pe^s, \text{stack}^s, \text{store}^s \rangle, \\ Th'_{tj} &= Th^o = \langle \text{env}^o, pe^o, \text{stack}^o, \text{store}^o \rangle. \end{aligned}$$

We use the exponent s for the elements of the simulator configuration and the exponent o for the elements of the OCaml configuration.

Let us first distinguish cases depending on whether Property 6(a) or Property 6(b) is satisfied for the current thread.

Case 1. Property 6(a) is satisfied for the current thread, that is, we are at the beginning of the initialization of a protocol thread. There exists a program $program'$ such that

$$Th^s = \langle \emptyset, program_{\text{prim}};; program'(\text{role}_1[\widetilde{a}_1]);; \dots;; program'(\text{role}_m[\widetilde{a}_m]);; program', [], \emptyset \rangle.$$

There is no closure, no tagged function $\text{tagfunction}^t pm$, no event, and no return in $program'$, except in $program(\mu_{\text{role}})$ in arguments of `addthread`. The OCaml thread verifies $Th^o = \text{replaceinitpm}(Th^s)$, so

$$Th^o = \langle \emptyset, program_{\text{prim}};; program(\mu_{\text{role}_1});; \dots;; program(\mu_{\text{role}_m});; program', [], \emptyset \rangle.$$

By Assumption 6.2, there is exactly one complete thread trace \mathcal{TT} that begins at $\langle \emptyset, program_{\text{prim}};; [], \emptyset \rangle$, and the last thread of this trace is $\langle env_{\text{prim}}, \varepsilon, [], \emptyset \rangle$. So there is no call to the `random` function inside the initialization of the primitives. Let $\mathcal{TT}(\text{definitions})$ be the trace \mathcal{TT} where, in each thread, we replace the empty definition list ε by definitions . As no OCaml reduction rule depends on the contents of a definition list, the trace $\mathcal{TT}(\text{definitions})$ is a valid trace for any definition list definitions . So, by taking definitions the definitions after $program_{\text{prim}}$ in Th^s and Th^o ,

$$Th^s \rightarrow^* Th_e^s \stackrel{\text{def}}{=} \langle env_{\text{prim}}, program'(\text{role}_1[\widetilde{a}_1]);; \dots;; program'(\text{role}_m[\widetilde{a}_m]);; program', [], \emptyset \rangle$$

$$Th^o \rightarrow^* Th_e^o \stackrel{\text{def}}{=} \langle env_{\text{prim}}, program(\mu_{\text{role}_1});; \dots;; program(\mu_{\text{role}_m});; program', [], \emptyset \rangle$$

in exactly the same number of steps.

Let l_j ($j \leq m$) be m distinct locations. For $j \leq m + 1$, let $Th_j^s \stackrel{\text{def}}{=} \langle env_j^s, program_j, [], store_j^s \rangle$ where $env_j^s \stackrel{\text{def}}{=} env_{\text{prim}} \cup \{\mu_{\text{role}_i}.\text{init} \mapsto c_i \mid i < j\}$ with $c_i \stackrel{\text{def}}{=} \text{tagfunction}^{\text{role}_i, \tau_i}[env_i, pm'_{\text{role}_i}[\widetilde{a}_i]]$ and $env_i \stackrel{\text{def}}{=} env_i^s[\text{token} \mapsto l_i]$, $program_j \stackrel{\text{def}}{=} program'(\text{role}_j[\widetilde{a}_j]);; \dots;; program'(\text{role}_m[\widetilde{a}_m]);; program'$ for $j \leq m$ and $program_j \stackrel{\text{def}}{=} program'$ for $j = m + 1$, and $store_j^s \stackrel{\text{def}}{=} \{l_i \mapsto \text{true} \mid i < j\}$. For $j \leq m$, the thread Th_j^s reduces as follows:

$$\begin{aligned} Th_j^s &= \langle env_j^s, \text{let } \mu_{\text{role}_j}.\text{init} = e_j^s;; program_{j+1}, [], store_j^s \rangle \\ &\quad \text{where } e_j^s \stackrel{\text{def}}{=} \text{let } token = \text{ref true in tagfunction}^{\text{role}_j} pm'_{\text{role}_j}[\widetilde{a}_j] \\ &\rightarrow \langle env_j^s, e_j^s, stack_j^s, store_j^s \rangle \\ &\quad \text{where } stack_j^s \stackrel{\text{def}}{=} [env_j^s, \text{let } \mu_{\text{role}_j}.\text{init} = [];; program_{j+1}] \\ &\rightarrow^* \langle env_j, \text{tagfunction}^{\text{role}_j} pm'_{\text{role}_j}[\widetilde{a}_j], stack_j^s, store_{j+1}^s \rangle \\ &\quad \text{since } env_j = env_j^s[\text{token} \mapsto l_j] \text{ and } store_{j+1}^s = store_j^s[l_j \mapsto \text{true}] \end{aligned}$$

$$\begin{aligned}
& \rightarrow \langle env_j, c_j, stack_j^s, store_{j+1}^s \rangle \\
\rightarrow^* Th_{j+1}^s &= \langle env_{j+1}^s, program_{j+1}, [], store_{j+1}^s \rangle \\
& \text{since } env_{j+1}^s = env_j^s[\mu_{role_j} \cdot \text{init} \mapsto c_j]
\end{aligned}$$

Let $Th_j^o \stackrel{\text{def}}{=} \text{replaceinitpm}(Th_j^s)$. We have $Th_j^o = \langle env_j^o, program'_j, [], store_j^s \rangle$ where env_j^o is the environment env_j^s in which we replace $pm'_{role_i[\tilde{a}_i]}$ with $pm_{\mu_{role_i}}$ for all $i < j$, $program'_j \stackrel{\text{def}}{=} program(\mu_{role_j}); \dots; program(\mu_{role_m}); program'$ for $j \leq m$, and $program'_j \stackrel{\text{def}}{=} program'$ for $j = m + 1$. For $j \leq m$, the thread Th_j^o reduces as follows:

$$\begin{aligned}
Th_j^o &= \langle env_j^o, \text{let } \mu_{role_j} \cdot \text{init} = e_j^o;; program'_{j+1}, [], store_j^s \rangle \\
& \text{where } e_j^o \stackrel{\text{def}}{=} \text{let } token = \text{ref true in tagfunction}^{role_j} pm_{\mu_{role_j}} \\
& \rightarrow \langle env_j^o, e_j^o, stack_j^o, store_j^s \rangle \\
& \text{where } stack_j^o \stackrel{\text{def}}{=} [env_j^o, \text{let } \mu_{role_j} \cdot \text{init} = [];; program'_{j+1}] \\
\rightarrow^* & \langle env'_j, \text{tagfunction}^{role_j} pm_{\mu_{role_j}}, stack_j^o, store_{j+1}^s \rangle \\
& \text{where } env'_j \stackrel{\text{def}}{=} env_j^o[token \mapsto l_j] \\
& \rightarrow \langle env'_j, c'_j, stack_j^o, store_{j+1}^s \rangle \\
& \text{where } c'_j \stackrel{\text{def}}{=} \text{tagfunction}^{role_j, \tau_j} [env'_j, pm_{\mu_{role_j}}] \\
\rightarrow^* & Th_{j+1}^o = \langle env_{j+1}^o, program'_{j+1}, [], store_{j+1}^s \rangle \\
& \text{since } env_{j+1}^o = env_j^o[\mu_{role_j} \cdot \text{init} \mapsto c'_j]
\end{aligned}$$

Moreover, $Th_e^s = Th_1^s$ and $Th_e^o = Th_1^o$, so $Th^s \rightarrow^* Th_{m+1}^s$ and $Th^o \rightarrow^* Th_{m+1}^o$. There are exactly the same number of steps in both traces. Let $steps^s$ be this number of steps.

Let \mathcal{CT}_1 be the extension of the trace \mathcal{CT} until $\mathcal{C}[\text{Th} \mapsto Th_{m+1}^o]$. Since $Th^s \rightarrow^* Th_{m+1}^s$ without using (Random), we have $\mathcal{CS} \rightarrow^* \mathcal{CS}[\text{Th} \mapsto Th_{m+1}^s]$ by (Globalstore1), (Toplevel), and (Simulator toplevel). Furthermore, by definition of N_{steps} , all traces of the OCaml program have at most N_{steps} steps, so in particular $|\mathcal{CT}_1| = |\mathcal{CT}| + steps^s \leq N_{\text{steps}}$. Hence, by Property 15, $steps \geq N_{\text{steps}} - |\mathcal{CT}| \geq steps^s$. So, with $\mathfrak{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, steps - steps^s, \mathcal{CS}[\text{Th} \mapsto Th_{m+1}^s]$, we have $\mathfrak{C}^{\text{cs}} \rightsquigarrow^+ \mathfrak{C}_1^{\text{cs}}$ by (Simulator) since $steps$ remains positive during the reduction. (More generally, the same reasoning shows that, if the simulator trace has at most as many steps as the OCaml trace, then the extended CryptoVerif configuration can reduce by (Simulator) because $steps$ remains positive by Property 15. We shall not detail this point in the other cases.)

Let us prove that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}_1$. Properties 1, 2, 3, 4, 7, 8, 9, 10, 11, 12 are inherited from $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$. As mentioned in Section 3.2.6, there are no local store locations in the initial program, so there are no local store locations in $program'$, so the locations l_1, \dots, l_m are the only local store locations present in Th_{m+1}^s , and they are all in the domain of $store_{m+1}^s$. So Property 5

holds. Let l_{tok} be the empty function. The set $\mathcal{O}_{\text{call}}(Th_{m+1}^s)$ is empty. We have that $\text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th_{m+1}^s), l_{\text{tok}}) = \emptyset$ and $Th_{m+1}^o = \text{replaceinitpm}(Th_{m+1}^s) \in \text{replacecalls}(\text{replaceinitpm}(Th_{m+1}^s), \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_{\mathcal{O}})$, so Property 6(b)i holds for the current thread. Using the function $l_{\text{init-tok}}$ that maps $\text{role}_j[\tilde{a}_j]$ to l_j for $j \leq m$, Property 6(b)ii holds for the current thread. The environment of the tagged closures that we created contains env_{prim} , so Property 6(b)iii holds for the current thread. Since there is no tagged function, no event and no return in $\text{program}'$ except in $\text{program}(\mu_{\text{role}})$ in arguments of addthread , Property 6(b)iv holds for the current thread. Threads that are not the current thread did not change, so Property 6 holds. The only change in the oracle sets is that the roles $\text{role}_j[\tilde{a}_j]$ are transferred from $\mathcal{R}_{\text{init-function}}(Th^s)$ to $\mathcal{R}_{\text{init-closure}}(Th^s)$, so Properties 13 and 14 are preserved. We have

$$|\mathcal{CT}_1| + \text{steps} - \text{steps}^s = |\mathcal{CT}| + \text{steps}^s + \text{steps} - \text{steps}^s \geq N_{\text{steps}}$$

so Property 15 holds. Properties 16, 17, and 18 are preserved, because all components of these inequalities are unchanged. Therefore, we have proved that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}_1$.

Case 2. Property 6(b) is satisfied for the current thread. We now distinguish cases on the form of the simulator configuration \mathcal{CS} .

Case 2.1. The current expression of \mathcal{CS} is $pe^s = \text{call}(O_j[\tilde{a}]) (v_1, \dots, v_{m_j})$ and \mathcal{CS} cannot reduce, that is, the configuration \mathcal{CS} makes a successful call to $O_j[\tilde{a}]$, an oracle not under replication. By definition of simreturn , $\text{simreturn}(\mathcal{CS}) \stackrel{\text{def}}{=} (\text{repr}(\mathcal{CS}), o_j, \tilde{a}, \text{args})$ where $\text{args} \stackrel{\text{def}}{=} (b_1, \dots, b_{m_j})$ and $b_k \stackrel{\text{def}}{=} \mathbb{G}_{\text{val}T_{j,k}}^{-1}(v_k)$ for $k \leq m_j$.

So \mathfrak{C}^{cs} reduces in several steps into the configuration $E_1, P'_{\text{loop}}\{\alpha/i'\}, \mathcal{T}, \mathcal{Q}, \mathcal{R}_{\text{loop}}(\alpha), \mathcal{E}$ that corresponds to line 8 where

$$\begin{aligned} E_1 &\stackrel{\text{def}}{=} E[xs[\alpha] \mapsto (\text{repr}(\mathcal{CS}), o_j, \tilde{a}, \text{args}), \\ &\quad s'[\alpha] \mapsto \text{repr}(\mathcal{CS}), o[\alpha] \mapsto o_j, i[\alpha] \mapsto \tilde{a}, \text{args}[\alpha] \mapsto \text{args}]. \end{aligned}$$

Let $a'_1, \dots, a'_{n_j} \stackrel{\text{def}}{=} \tilde{a}$. As $E_1(o[\alpha]) = o_j$, this configuration reduces in several steps into the configuration $E_2, P'_{\text{loop}}\{\alpha/i'\}, \mathcal{T}, \mathcal{Q}, \mathcal{R}_{\text{loop}}(\alpha), \mathcal{E}$ where

$$\begin{aligned} E_2 &\stackrel{\text{def}}{=} E_1[xa_j[\alpha] \mapsto \text{args}, a_{j,1}[\alpha] \mapsto b_1, \dots, a_{j,m_j}[\alpha] \mapsto b_{m_j}, \\ &\quad xi_j[\alpha] \mapsto \tilde{a}, i_{j,1}[\alpha] \mapsto a'_1, \dots, i_{j,n_j}[\alpha] \mapsto a'_{n_j}]. \end{aligned}$$

The oracle $O_j[\tilde{a}]$ is in \mathcal{I} , otherwise \mathcal{CS} could reduce using (FailedCall1). By Property 3 of the invariant, there exists \mathcal{Q}_0 such that $\mathcal{Q}_0 \leftrightarrow \mathcal{I}, \mathcal{RI}$ and $\mathcal{Q} = \{Q_{\text{loop}}\{\alpha/i'\} \mid \alpha < a \leq N_{\text{rand+calls}}\} \cup \mathcal{Q}_0$. Let $Q \stackrel{\text{def}}{=} \mathcal{Q}_0(O_j[\tilde{a}])$. The oracle definition Q is of the form $O_j[\tilde{a}](x_1[\tilde{a}] : T_{j,1}, \dots, x_{m_j}[\tilde{a}] : T_{j,m_j}) := P_O$. The previous configuration reduces in one step into $\mathfrak{C} \stackrel{\text{def}}{=} E_3, P_O, \mathcal{T}, \mathcal{Q}_1, \mathcal{R}_1, \mathcal{E}$ where

$$E_3 \stackrel{\text{def}}{=} E_2[x_1[\tilde{a}] \mapsto b_1, \dots, x_{m_j}[\tilde{a}] \mapsto b_{m_j}]$$

$$\begin{aligned}
\mathcal{R}_1 &\stackrel{\text{def}}{=} ((r_{j,1}, \dots, r_{j,m'_j}), \text{return}(\text{simulate}'_{O_j}(s', (r_{j,1}, \dots, r_{j,m'_j})), \text{true}), \\
&\quad \text{return}(\text{simulate}''_{O_j}(s'), \text{true})) :: \mathcal{R}_{loop}(\alpha) \\
\mathcal{Q}_1 &\stackrel{\text{def}}{=} \mathcal{Q} \setminus \{Q\}
\end{aligned}$$

Let us now look at \mathcal{C} . By the invariant, there exists an injection l_{tok} that satisfies Property 6(b)i. The current expression pe^o is of the form $c(v_1, \dots, v_{m_j})$, where $c \in \text{correctclosure}(O_j[\tilde{a}], \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_O)$, that is, $c = \text{tagfunction}^{O_j, \tau}[env_1^o, pm_{\text{false}}(Q)]$ where $env_1^o \supseteq env_{\text{prim}} \cup env(E, Q)$ and $env_1^o(\text{token}) = l_{\text{tok}}(O_j[\tilde{a}])$. By the same property, $store^o(l_{\text{tok}}(O_j[\tilde{a}])) = \text{true}$.

$$\begin{aligned}
Th^o &= \langle env^o, c(v_1, \dots, v_{m_j}), stack^o, store^o \rangle \\
&\rightarrow \langle env_1^o, \text{match}(v_1, \dots, v_{m_j}) \text{ with } pm_{\text{false}}(Q), stack^o, store^o \rangle \\
\rightarrow Th_1^o &\stackrel{\text{def}}{=} \langle env_2^o, e, stack^o, store^o \rangle \\
&\quad \text{where } env_2^o \stackrel{\text{def}}{=} env_1^o[\mathbb{G}_{\text{var}}(x_1) \mapsto v_1, \dots, \mathbb{G}_{\text{var}}(x_{m_j}) \mapsto v_{m_j}] \text{ and} \\
&\quad e \stackrel{\text{def}}{=} \text{if } (!\text{token}) \&\& \\
&\quad \quad (\mathbb{G}_{\text{pred}}(T_{j,1}) \mathbb{G}_{\text{var}}(x_1)) \&\& \dots \&\& (\mathbb{G}_{\text{pred}}(T_{j,m_j}) \mathbb{G}_{\text{var}}(x_{m_j})) \\
&\quad \text{then } (\text{token} := \text{false}; e') \text{ else raise Bad_Call} \\
e' &\stackrel{\text{def}}{=} \mathbb{G}_{\text{file}}(x_1[\tilde{a}]); \dots; \mathbb{G}_{\text{file}}(x_{m_j}[\tilde{a}]); \mathbb{G}(P_O)
\end{aligned}$$

For all $k \leq m_j$, there exists b_k such that $\mathbb{G}_{\text{val}T_{j,k}}(b_k) = v_k$, so $\mathbb{G}_{\text{pred}}(T_{j,k}) \mathbb{G}_{\text{var}}(x_k)$ evaluates to true using Proposition 6.5. Moreover, $env_2^o(\text{token}) = l_{\text{tok}}(O_j[\tilde{a}])$. So, the configuration \mathcal{C} reduces as follows

$$\begin{aligned}
\mathcal{C} &\rightarrow^* \mathcal{C}' \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto Th_1^o] \\
&\rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env_2^o, \text{token} := \text{false}; e', stack^o, store_1^o \rangle] \\
&\rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env_2^o, e', stack^o, store_2^o \rangle] \\
&\quad \text{where } store_2^o \stackrel{\text{def}}{=} store_1^o[l_{\text{tok}}(O_j[\tilde{a}]) \mapsto \text{false}] \\
&\rightarrow^* \mathcal{C}_1 \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto \langle env_2^o, \mathbb{G}(P_O), stack^o, store_3^o \rangle, \text{globalstore} \mapsto \text{globalstore}_1^o] \\
&\quad \text{by Lemma D.1 applied } m_j \text{ times}
\end{aligned}$$

where $store_3^o \supseteq store_2^o$, $store_1^o \supseteq store^o$, $\text{globalstore}_1^o \supseteq \text{globalstore}(E_3, \mathcal{T})$ since $\text{globalstore}^o \supseteq \text{globalstore}(E, \mathcal{T})$ by Property 9 of the invariant, and $\text{globalstore}_1^o(l) = \text{globalstore}^o(l)$ for all $l \notin S_{\text{priv}}$.

We prove that for any traces $\mathfrak{C}\mathfrak{I}_1, \dots, \mathfrak{C}\mathfrak{I}_m$ beginning at \mathfrak{C} such that $\sum_{i \leq m} \Pr[\mathfrak{C}\mathfrak{I}_i] = 1$, none of these traces is a prefix of another, and there is no intermediate configuration inside any of these traces with a return, end, call, or loop as current process, there exist m disjoint sets of OCaml traces $\mathcal{C}\mathcal{T}\mathcal{S}_1, \dots, \mathcal{C}\mathcal{T}\mathcal{S}_m$ all starting at \mathcal{C}_1 such that none of these traces is a prefix of another of these traces, $\Pr[\mathcal{C}\mathcal{T}\mathcal{S}_i] = \Pr[\mathfrak{C}\mathfrak{I}_i]$ for all $i \leq m$, and if \mathcal{C}_4 is the last configuration of a trace $\mathcal{C}\mathcal{T}' \in \mathcal{C}\mathcal{T}\mathcal{S}_i$, then $\mathcal{C}_4 = \mathcal{C}[\text{Th} \mapsto Th_4^o, \text{globalstore} \mapsto$

$globalstore_4^o, events \mapsto events_4]$ where

$$\begin{aligned}
Th_4^o &= \langle env_4^o, \mathbb{G}(P_4), stack^o, store_4^o \rangle \text{ with} \\
env_4^o &\supseteq env_{\text{prim}} \cup env(E_4, P_4) \text{ and } store_4^o \supseteq store_3^o, \\
globalstore_4^o &\supseteq globalstore(E_4, \mathcal{T}_4), \\
globalstore_4^o(l) &= globalstore^o(l) \text{ for all } l \notin S_{\text{priv}}, \\
events_4 &= \mathbb{G}_{\text{ev}}(\mathcal{E}_4),
\end{aligned}$$

and the last configuration of $\mathfrak{C}\mathfrak{T}_i$ is $E_4, P_4, \mathcal{T}_4, \mathcal{Q}_1, \mathcal{R}_1, \mathcal{E}_4$.

The proof proceeds by induction on the total length of the traces $\mathfrak{C}\mathfrak{T}_1, \dots, \mathfrak{C}\mathfrak{T}_m$. In the base case, $m = 1$ and $\mathfrak{C}\mathfrak{T}_1$ is the trace that consists only of the configuration \mathfrak{C} . Let $\mathcal{C}\mathcal{T}\mathcal{S}_1$ consist of the single trace that contains just the configuration \mathcal{C}_1 . We have $env_2^o \supseteq env_{\text{prim}} \cup env(E_3, P_O)$ since $env^o \supseteq env_{\text{prim}} \cup env(E, Q)$, the variables $x_1[\tilde{a}], \dots, x_{m_j}[\tilde{a}]$ are added on the CryptoVerif side, and $\mathbb{G}_{\text{var}}(x_1), \dots, \mathbb{G}_{\text{var}}(x_{m_j})$ are added correspondingly on the OCaml side. As shown above, $globalstore_1^o \supseteq globalstore(E_3, \mathcal{T})$ and $globalstore_1^o(l) = globalstore^o(l)$ for all $l \notin S_{\text{priv}}$. By Property 12 of the invariant, $events = \mathbb{G}_{\text{ev}}(\mathcal{E})$. So the property holds for the base case. The inductive case follows from Lemma 6.10.

Let us take the maximal CryptoVerif traces $\mathfrak{C}\mathfrak{T}_1, \dots, \mathfrak{C}\mathfrak{T}_n$ that begin at \mathfrak{C} and that contain no return, end, call, or loop as current process in intermediate configurations. Let $\mathcal{C}\mathcal{T}\mathcal{S}_1, \dots, \mathcal{C}\mathcal{T}\mathcal{S}_n$ the trace sets as defined above. The final configurations of the CryptoVerif traces $\mathfrak{C}\mathfrak{T}_i$ contain either return or end, since the oracle $O_j[\tilde{a}]$ does not contain loop or call constructs. Let us take one such trace $\mathfrak{C}\mathfrak{T}_i$ and a trace $\mathcal{C}\mathcal{T}' \in \mathcal{C}\mathcal{T}\mathcal{S}_i$. Let \mathfrak{C}_4 and \mathcal{C}_4 be the last configurations of $\mathfrak{C}\mathfrak{T}_i$ and $\mathcal{C}\mathcal{T}'$ respectively. Let $\mathfrak{C}_4 = E_4, P_4, \mathcal{T}_4, \mathcal{Q}_1, \mathcal{R}_1, \mathcal{E}_4$. We distinguish cases on the form of P_4 .

- If $P_4 = \text{end}$,

$$\begin{aligned}
\mathfrak{C}_4 &\rightsquigarrow E_4, \text{return}(\text{simulate}_{O_j}''(s'[\alpha]), \text{true}), \mathcal{T}_4, \mathcal{Q}_1, \mathcal{R}_{\text{loop}}(\alpha), \mathcal{E}_4 \\
&\rightsquigarrow E_5, P_{\text{return-loop}}(\alpha), \mathcal{T}_4, \mathcal{Q}_1, [x[]], \text{return}(x[]), \text{end}], \mathcal{E}_4 \\
&\text{where } E_5 \stackrel{\text{def}}{=} E_4[r'_{\alpha,r}[]] \mapsto s, b_{\alpha,r} \mapsto \text{true}], \\
&\quad s \stackrel{\text{def}}{=} \text{repr}(\mathcal{C}\mathcal{S}'), \\
&\quad \mathcal{C}\mathcal{S}' \text{ is } \mathcal{C}\mathcal{S} \text{ in which the current expression is replaced} \\
&\quad \text{with } \text{raise Match_failure} \text{ and the set } \mathcal{I} \text{ is replaced with} \\
&\quad \mathcal{I}' \stackrel{\text{def}}{=} \mathcal{I} - (O_j[\tilde{a}]) \\
&\rightsquigarrow E_5, P_5, \mathcal{T}_4, \mathcal{Q}_1, [x[]], \text{return}(x[]), \text{end}], \mathcal{E}_4 \\
&\quad \text{where } P_5 \stackrel{\text{def}}{=} \text{let } r[] : T_{\mathcal{C}\mathcal{S}} = \text{loop } O_{\text{loop}}[\alpha + 1](r'_{\alpha,r}[]) \\
&\quad \quad \text{in end else end} \\
&\rightsquigarrow E_5, P_6, \mathcal{T}_4, \mathcal{Q}_1, [x[]], \text{return}(x[]), \text{end}], \mathcal{E}_4 \\
&\quad \text{where } P_6 \stackrel{\text{def}}{=} \text{let } (r'_{\alpha+1,r}[] : T_{\mathcal{C}\mathcal{S}}, b_{\alpha+1,r}[] : \text{bool}) = \\
&\quad \quad O_{\text{loop}}[\alpha + 1](r'_{\alpha,r}) \text{ in } P_{\text{return-loop}}(\alpha + 1) \text{ else end}
\end{aligned}$$

$\rightsquigarrow E_6, P_{loop}\{\alpha + 1/i'\}, \mathcal{T}_4, \mathcal{Q}_2, \mathcal{R}_{loop}(\alpha + 1), \mathcal{E}_4$
 where $E_6 \stackrel{\text{def}}{=} E_5[s[\alpha + 1] \mapsto s]$,
 $\mathcal{Q}_2 \stackrel{\text{def}}{=} \mathcal{Q}_1 \setminus \{Q_{loop}\{\alpha + 1/i'\}\}$
 (Let \mathcal{CT}'' be the trace \mathcal{CT} followed by \mathcal{CT}' . By definition of $N_{\text{rand+calls}}$, $N_{\text{rand+calls}} \geq \left(N_{\text{rand}}(\mathcal{CT}'') + \sum_{O,\tau} N_{\text{calls}}(O, \tau, \mathcal{CT}'')\right) + 1 = \left(N_{\text{rand}}(\mathcal{CT}) + \sum_{O,\tau} N_{\text{calls}}(O, \tau, \mathcal{CT})\right) + 2$ since \mathcal{CT}'' makes one more call to O_j than \mathcal{CT} . So, by Property 16, $N_{\text{rand+calls}} \geq \alpha + 1$. So, by Property 3, $Q_{loop}\{\alpha + 1/i'\} \in \mathcal{Q}$, so $Q_{loop}\{\alpha + 1/i'\} \in \mathcal{Q}_1$.)
 $\rightsquigarrow \mathfrak{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E_6, P_{loop}\{\alpha + 1/i'\}, \mathcal{T}_4, \mathcal{Q}_2, \mathcal{R}_{loop}(\alpha + 1), \mathcal{E}_4, N_{\text{steps}}, \mathcal{CS}'$

By definition of the translation of end, the current expression of \mathcal{C}_4 is raise Match_failure. Let \mathcal{CT}'' be the trace \mathcal{CT} followed by \mathcal{CT}' . The last configuration of \mathcal{CT}'' is \mathcal{C}_4 .

Let us prove that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}''$. By the form of $\mathfrak{C}_1^{\text{cs}}$ and \mathcal{C}_4 , Properties 1 and 2 hold. The set \mathcal{Q}_2 is the set \mathcal{Q} where we removed the oracles $O_j[\tilde{a}]$ and $O_{loop}[\alpha + 1]$. We have $\mathcal{I}' = \mathcal{I} - (O_j[\tilde{a}])$, so Property 3 is preserved. Property 4 is an immediate consequence of Lemma F.2. No new locations were created in the simulator, and the domains of stores can only grow, so Property 5 is preserved.

For all threads $tj' \neq tj$, the thread tj' does not change so, to prove Property 6, we just have to show that Property 6(b)i is preserved; the other elements of Property 6 are obviously preserved. Suppose that thread tj' satisfies Property 6(b)i initially, with a function l_{tok} . By Lemma F.3, Item 1, for all call(R) that occur in $Th''_{tj'} \stackrel{\text{def}}{=} \text{replaceinitpm}(Th_{tj'})$, we have $\text{correctclosure}(R, \mathcal{I}', E_6, \mathcal{Q}_2, l_{\text{tok}}, \tau_0) \supseteq \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_0)$, so $\text{replacecalls}(Th''_{tj'}, \mathcal{I}', E_6, \mathcal{Q}_2, l_{\text{tok}}, \tau_0) \supseteq \text{replacecalls}(Th''_{tj'}, \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_0)$. Furthermore, the oracle $O_j[\tilde{a}]$ is in $\mathcal{O}_{\text{call}}(Th_{tj})$, so by Property 14 of the invariant, it is not in $\mathcal{O}_{\text{call}}(Th_{tj'})$. Hence, $\mathcal{O}_{\text{call}}(Th_{tj'}) \cap \mathcal{I}' = \mathcal{O}_{\text{call}}(Th_{tj'}) \cap \mathcal{I}$ and $\mathcal{O}_{\text{call}}(Th_{tj'}) \setminus \mathcal{I}' = \mathcal{O}_{\text{call}}(Th_{tj'}) \setminus \mathcal{I}$, so $\text{gettokens}(\mathcal{I}', \mathcal{O}_{\text{call}}(Th_{tj'}), l_{\text{tok}}) = \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th_{tj'}), l_{\text{tok}})$. So thread tj' continues to verify Property 6(b)i with the same function l_{tok} .

Let us now consider the current thread. The current thread of the simulator is $Th_1^s = \langle env^s, \text{raise Match_failure}, stack^s, store^s \rangle$ and the current thread on the OCaml side is $Th_4^o = \langle env_4^o, \text{raise Match_failure}, stack^o, store_4^o \rangle$ where $store_4^o \supseteq store_3^o$. By Property 6(b)i, there exist $store_5^o$ and l_{tok} such that

$$\begin{aligned}
 & \langle env^o, pe^o, stack^o, store_5^o \rangle \\
 & \in \text{replacecalls}(\text{replaceinitpm}(Th^s), \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_0) \\
 & store_5^o \cup \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th^s), l_{\text{tok}}) \subseteq store^o.
 \end{aligned}$$

Let us denote $Th'' \stackrel{\text{def}}{=} \text{replaceinitpm}(Th^s)$ and $Th_1'' \stackrel{\text{def}}{=} \text{replaceinitpm}(Th_1^s)$. The thread Th_1'' is the thread Th'' in which the current expression is

replaced with `raise Match_failure`. This is an exceptional value, so the definition of `replacecalls` allows any environment in the threads it returns, hence

$$\begin{aligned} Th_5^o &\stackrel{\text{def}}{=} \langle env_4^o, \text{raise Match_failure}, stack^o, store_5^o \rangle \\ &\in \text{replacecalls}(Th_1'', \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_0). \end{aligned}$$

Let $l'_{\text{tok}} \stackrel{\text{def}}{=} l_{\text{tok}|_{\mathcal{O}_{\text{call}}(Th_1^s)}}$. By Lemma F.3, Item 1, for all `call(R)` that occur in Th_1'' , we have $\text{correctclosure}(R, \mathcal{I}', E_6, \mathcal{Q}_2, l'_{\text{tok}}, \tau_0) \supseteq \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_0)$, so $Th_5^o \in \text{replacecalls}(Th_1'', \mathcal{I}', E_6, \mathcal{Q}_2, l'_{\text{tok}}, \tau_0)$. We have

$$\begin{aligned} &store_5^o \cup \text{gettokens}(\mathcal{I}', \mathcal{O}_{\text{call}}(Th_1^s), l'_{\text{tok}}) \\ &\subseteq store_5^o \cup \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th^s), l_{\text{tok}})[l_{\text{tok}}(O_j[\tilde{a}]) \mapsto \text{false}] \\ &\subseteq store^o[l_{\text{tok}}(O_j[\tilde{a}]) \mapsto \text{false}] \\ &\subseteq store_1^o[l_{\text{tok}}(O_j[\tilde{a}]) \mapsto \text{false}] = store_2^o \subseteq store_3^o \subseteq store_4^o, \end{aligned}$$

so Property 6(b)i holds with the function l'_{tok} . Properties 6(b)ii, 6(b)iii, 6(b)iv are clearly preserved, so Property 6 holds.

Properties 7, 8, and 11 are also preserved. Properties 9, 10, and 12 hold because they are kept inside Lemma 6.10. We have $\mathcal{O}^\infty(\mathcal{I}') = \mathcal{O}^\infty(\mathcal{I}) \setminus \{O_j[\tilde{a}]\}$. If there remains no occurrence of `call(Oj[\tilde{a}])` in the thread Th_1^s , then $\mathcal{O}_{\text{call}}(Th_1^s) = \mathcal{O}_{\text{call}}(Th^s) \setminus \{O_j[\tilde{a}]\}$ and $\mathcal{O}_{\text{call}}(\mathcal{CS}') = \mathcal{O}_{\text{call}}(\mathcal{CS}) \setminus \{O_j[\tilde{a}]\}$, so $\mathcal{O}^\infty(\mathcal{I}') \cup \mathcal{O}_{\text{call}}(\mathcal{CS}') = \mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{CS}) \setminus \{O_j[\tilde{a}]\}$. Otherwise, $\mathcal{O}_{\text{call}}(Th_1^s) = \mathcal{O}_{\text{call}}(Th^s)$ and $\mathcal{O}_{\text{call}}(\mathcal{CS}') = \mathcal{O}_{\text{call}}(\mathcal{CS})$, so $\mathcal{O}^\infty(\mathcal{I}') \cup \mathcal{O}_{\text{call}}(\mathcal{CS}') = \mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{CS})$. We also have $\mathcal{O}_{\text{call-repl}}(Th_1^s) = \mathcal{O}_{\text{call-repl}}(Th^s)$, $\mathcal{O}^\infty(\mathcal{R}_{\text{init-function}}(Th_1^s)) = \mathcal{O}^\infty(\mathcal{R}_{\text{init-function}}(Th^s))$, and $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(Th_1^s)) = \mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(Th^s))$, so Property 14 is preserved. The set $\mathcal{O}^\infty(\mathcal{I}') \cup \mathcal{O}_{\text{call}}(\mathcal{CS}')$ is included in $\mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{CS})$ and the set $\text{willbeavailable}(\mathcal{CS}')$ is included in $\text{willbeavailable}(\mathcal{CS})$, so Property 13 is preserved. We have $|\mathcal{CT}''| + N_{\text{steps}} \geq N_{\text{steps}}$, so Property 15 holds. The number of calls to O_j increases by 1 and α increases by 1, so Property 16 is preserved. Properties 17 and 18 are preserved, because all components of these inequalities are unchanged. So $\mathcal{C}_1^{\text{cs}} \equiv \mathcal{CT}''$ in this case.

- If $P_4 = \text{return}(M_1, \dots, M_{m'_j}); Q'$, let c_i ($i \leq m'_j$) be the `CryptoVerif` values such that $E_4, M_i \Downarrow c_i$.

$$\begin{aligned} \mathcal{C}_4 &\rightsquigarrow E_5, \text{return}(\text{simulate}'_{O_j}(s'[\alpha], (r_{j,1}[\alpha], \dots, r_{j,m'_j}[\alpha])), \text{true}), \mathcal{T}_4, \mathcal{Q}_1, \\ &\mathcal{R}, \mathcal{E}_4 \end{aligned}$$

$$\text{where } E_5 \stackrel{\text{def}}{=} E_4[r_{j,1} \mapsto c_1, \dots, r_{j,m'_j} \mapsto c_{m'_j}]$$

$$\rightsquigarrow^* \mathcal{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E_6, P_{\text{loop}}\{\alpha + 1/i'\}, \mathcal{T}_4, \mathcal{Q}_2, \mathcal{R}_{\text{loop}}(\alpha + 1), \mathcal{E}_4, N_{\text{steps}}, \mathcal{CS}'$$

$$\text{where } E_6 \supseteq E_5[s[\alpha + 1] \mapsto \text{repr}(\mathcal{CS}')],$$

$$\mathcal{Q}_2 \stackrel{\text{def}}{=} \mathcal{Q}_1 \cup \text{reduce}(Q') \setminus \{Q_{\text{loop}}\{\alpha + 1/i'\}\}$$

$$\text{repr}(\mathcal{CS}') = \text{simulate}'_{O_j}(\text{repr}(\mathcal{CS}), (c_1, \dots, c_{j,m'_j}))$$

where we show that $Q_{loop}\{\alpha + 1/i'\} \in \mathcal{Q}_1$ using Property 16 as in the case $P_4 = \text{end}$.

Suppose that $\text{reduce}'(Q') = [(Q_1, b_1), \dots, (Q_l, b_l)]$ and $\text{oracles}'(Q') = [O'_1[\tilde{a}_1], \dots, O'_l[\tilde{a}_l]]$. A thread $\langle \text{env}, \mathbb{G}_O(Q_i, b_i), \text{stack}, \text{store} \rangle$ where $\text{env} \supseteq \text{env}_{\text{prim}} \cup \text{env}(E_4, P_4)$ reduces into $\langle \text{env}', c(Q_i, b_i), \text{stack}, \text{store}' \rangle$ where $c(Q_i, b_i) \stackrel{\text{def}}{=} \text{tagfunction}^{O'_i, \tau_i}[\text{env}', pm_{b_i}(Q_i)]$ and

- if $b_i = \text{false}$, then $\text{env}' = \text{env}[\text{token} \mapsto l_i]$ and $\text{store}' = \text{store}[l_i \mapsto \text{true}]$ where l_i is a fresh location: $l_i \notin \text{Dom}(\text{store})$;
- if $b_i = \text{true}$, then $\text{env}' = \text{env}$ and $\text{store}' = \text{store}$.

So in both cases, $\text{env}' \supseteq \text{env}_{\text{prim}} \cup \text{env}(E_4, P_4)$.

Let $Th_4^o = \langle \text{env}_4^o, \mathbb{G}(P_4), \text{stack}^o, \text{store}_4^o \rangle$ be the current thread of \mathcal{C}_4 . We have $\text{env}_4^o \supseteq \text{env}_{\text{prim}} \cup \text{env}(E_4, P_4)$ and $\text{store}_4^o \supseteq \text{store}_3^o$.

$$\begin{aligned} Th_4^o &= \langle \text{env}_4^o, (\mathbb{G}_O(Q_1, b_1), \dots, \mathbb{G}_O(Q_l, b_l), \mathbb{G}_M(M_1), \dots, \mathbb{G}_M(M_{m'_j})), \\ &\quad \text{stack}^o, \text{store}_4^o \rangle \\ &\rightarrow^* \langle \text{env}_4^o, (\mathbb{G}_O(Q_1, b_1), \dots, \mathbb{G}_O(Q_l, b_l), \mathbb{G}_{\text{val}T'_{j,1}}(c_1), \dots, \\ &\quad \mathbb{G}_{\text{val}T'_{j,m'_j}}(c_{m'_j})), \text{stack}^o, \text{store}_5^o \rangle \end{aligned}$$

by Lemma 6.8 applied m'_j times

$$\begin{aligned} \rightarrow^* Th_5^o &\stackrel{\text{def}}{=} \langle \text{env}_4^o, (c(Q_1, b_1), \dots, c(Q_l, b_l), \mathbb{G}_{\text{val}T'_{j,1}}(c_1), \dots, \mathbb{G}_{\text{val}T'_{j,m'_j}}(c_{m'_j})), \\ &\quad \text{stack}^o, \text{store}_6^o \rangle \end{aligned}$$

where $\text{store}_6^o \stackrel{\text{def}}{=} \text{store}_5^o \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} \supseteq \text{store}_5^o \supseteq \text{store}_4^o$.

Let \mathcal{CT}'' be the trace \mathcal{CT} followed by \mathcal{CT}' and extended until $\mathcal{C}_5 \stackrel{\text{def}}{=} \mathcal{C}_4[\text{Th} \mapsto Th_5^o]$. Let \mathcal{I}' be the set \mathcal{I} of \mathcal{CS}' .

Let us now prove that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}''$. We define $\tau'_O \stackrel{\text{def}}{=} \tau_O \cup \{O'_i[_, \tilde{a}] \mapsto \tau_i \mid b_i = \text{true}\}$. (This function is well defined, because $O'_i[_, \tilde{a}] \notin \text{Dom}(\tau_O)$. Indeed, for any a' , $O'_i[a', \tilde{a}] \in \text{willbeavailable}(\mathcal{CS})$, so by Property 13, $O'_i[a', \tilde{a}] \notin \mathcal{O}^\infty(\mathcal{I})$, hence for any a' , $O'_i[[a', +\infty], \tilde{a}] \notin \mathcal{I}$. The main reason why we introduced the set $\mathcal{O}^\infty(\mathcal{I})$ is that at this point, we are able to distinguish between an oracle under replication that has not been called yet and an oracle whose calls have been exhausted. If we used the set $\mathcal{O}(\mathcal{I})$ instead here, we would not be able to conclude that there is no oracle $O'_i[[a', +\infty], \tilde{a}]$ in \mathcal{I} : if $a' > N_{O'_i}$, then $\mathcal{O}(\{O'_i[[a', +\infty], \tilde{a}]\}) = \emptyset$.) By the form of $\mathfrak{C}_1^{\text{cs}}$ and \mathcal{C}_5 , Properties 1 and 2 hold. The set \mathcal{Q}_2 is the set \mathcal{Q} where we removed the oracles $O_j[\tilde{a}]$ and $O_{loop}[\alpha + 1]$ and where we added the new oracles $\text{reduce}'(Q')$. By definition of $\text{simulate}'_{O_j}$, the set \mathcal{I}' is the set \mathcal{I} where we removed $O_j[\tilde{a}]$ and added the elements of $\text{oracles}'(Q')$. So Property 3 is preserved. We also have $\mathcal{Q}_2(O'_i[\tilde{a}_i]) = Q_i$ for $i \leq l$. Property 4 is an immediate consequence of Lemma F.2. No new locations were created

in the simulator, and the domains of stores can only grow, so Property 5 is preserved.

For all threads $tj' \neq tj$, the thread tj' does not change so, to prove Property 6, we just have to show that Property 6(b)i is preserved; the other elements of Property 6 are obviously preserved. Suppose that thread tj' satisfies Property 6(b)i initially, with a function l_{tok} . By Lemma F.3, Items 1 and 2, for all $\text{call}(R)$ that occur in $Th''_{tj'} \stackrel{\text{def}}{=} \text{replaceinitpm}(Th_{tj'})$, we have $\text{correctclosure}(R, \mathcal{I}', E_6, \mathcal{Q}_2, l_{\text{tok}}, \tau'_0) \supseteq \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_0)$, so $\text{replacecalls}(Th''_{tj'}, \mathcal{I}', E_6, \mathcal{Q}_2, l_{\text{tok}}, \tau'_0) \supseteq \text{replacecalls}(Th''_{tj'}, \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_0)$. Furthermore, the oracle $O_j[\tilde{a}]$ is in $\mathcal{O}_{\text{call}}(Th_{tj})$ and the oracles $O'_i[\tilde{a}_i]$ that are not under replication are in $\text{willbeavailable}(\mathcal{CS})$, so by Properties 13 and 14 of the invariant, they are not in $\mathcal{O}_{\text{call}}(Th_{tj'})$. Hence, $\mathcal{O}_{\text{call}}(Th_{tj'}) \cap \mathcal{I}' = \mathcal{O}_{\text{call}}(Th_{tj'}) \cap \mathcal{I}$ and $\mathcal{O}_{\text{call}}(Th_{tj'}) \setminus \mathcal{I}' = \mathcal{O}_{\text{call}}(Th_{tj'}) \setminus \mathcal{I}$, so $\text{gettokens}(\mathcal{I}', \mathcal{O}_{\text{call}}(Th_{tj'}), l_{\text{tok}}) = \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th_{tj'}), l_{\text{tok}})$. So thread tj' continues to verify Property 6(b)i with the same function l_{tok} .

Let us now consider the current thread. The current thread of the simulator is $Th_1^s = \langle env^s, pe_1^s, stack^s, store^s \rangle$ where $pe_1^s \stackrel{\text{def}}{=} (\text{call}(O'_1[\tilde{a}_1]), \dots, \text{call}(O'_l[\tilde{a}_l]), \mathbb{G}_{\text{val}T'_{j,1}}(c_1), \dots, \mathbb{G}_{\text{val}T'_{j,m'_j}}(c_{m'_j}))$ and the current thread on the OCaml side is $Th_5^o = \langle env_4^o, pe_6^o, stack^o, store_6^o \rangle$ where $pe_6^o \stackrel{\text{def}}{=} (c(Q_1, b_1), \dots, c(Q_l, b_l), \mathbb{G}_{\text{val}T'_{j,1}}(c_1), \dots, \mathbb{G}_{\text{val}T'_{j,m'_j}}(c_{m'_j}))$. By Property 6(b)i, there exist $store_7^o$ and l_{tok} such that

$$\begin{aligned} & \langle env^o, pe^o, stack^o, store_7^o \rangle \\ & \in \text{replacecalls}(\text{replaceinitpm}(Th^s), \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_0) \\ & store_7^o \cup \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th^s), l_{\text{tok}}) \subseteq store^o. \end{aligned}$$

Let us denote $Th'' \stackrel{\text{def}}{=} \text{replaceinitpm}(Th^s)$ and $Th''_1 \stackrel{\text{def}}{=} \text{replaceinitpm}(Th_1^s)$. The thread Th''_1 is the thread Th'' where the current expression is replaced with pe_1^s . Let $Th_2^s \stackrel{\text{def}}{=} \langle env^s, (), stack^s, store^s \rangle$ be a thread intermediate between Th^s and Th_1^s , in which the result of the call has not been inserted yet in the thread. When $b_i = \text{false}$, $O'_i[\tilde{a}_i]$ is in $\text{willbeavailable}(\mathcal{CS})$, so by Property 13, $O'_i[\tilde{a}_i]$ is not in $\mathcal{O}_{\text{call}}(Th^s)$, so we can define $l'_{\text{tok}} \stackrel{\text{def}}{=} l_{\text{tok}}|_{\mathcal{O}_{\text{call}}(Th_2^s)} \cup \{O'_i[\tilde{a}_i] \mapsto l_i \mid b_i = \text{false}\}$ and l'_{tok} is an extension of $l_{\text{tok}}|_{\mathcal{O}_{\text{call}}(Th_2^s)}$. By Lemma F.3, Items 1 and 2, for all $\text{call}(R)$ that occur in $Th''_2 \stackrel{\text{def}}{=} \text{replaceinitpm}(Th_2^s)$, we have $\text{correctclosure}(R, \mathcal{I}', E_6, \mathcal{Q}_2, l'_{\text{tok}}, \tau'_0) \supseteq \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_0)$. Moreover, $c(Q_i, b_i) \in \text{correctclosure}(O'_i[\tilde{a}_i], \mathcal{I}', E_6, \mathcal{Q}_2, l'_{\text{tok}}, \tau'_0)$ for $i \leq l$, and pe_1^s is a value so replacecalls allows any environment in the threads it returns, so $\langle env_4^o,$

$pe_6^o, stack^o, store_7^o \in \text{replacecalls}(Th_1'', \mathcal{I}', E_6, \mathcal{Q}_2, l'_{\text{tok}}, \tau'_0)$. We have

$$\begin{aligned}
& store_7^o \cup \text{gettokens}(\mathcal{I}', \mathcal{O}_{\text{call}}(Th_1^s), l'_{\text{tok}}) \\
& \subseteq store_7^o \cup \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th^s), l_{\text{tok}})[l_{\text{tok}}(O_j[\tilde{a}]) \mapsto \text{false}] \\
& \quad \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} \\
& \subseteq store^o[l_{\text{tok}}(O_j[\tilde{a}]) \mapsto \text{false}] \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} \\
& \subseteq store_1^o[l_{\text{tok}}(O_j[\tilde{a}]) \mapsto \text{false}] \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} \\
& \subseteq store_2^o \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} \\
& \subseteq store_5^o \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} = store_6^o,
\end{aligned}$$

so Property 6(b)i holds with the function l'_{tok} . Properties 6(b)ii, 6(b)iii, 6(b)iv are preserved, so Property 6 holds.

Properties 7, 8, and 11 are also preserved. Properties 9, 10, 12 hold because they are kept inside Lemma 6.10. The oracles coming from $\text{oracles}'(Q')$ are removed from $\text{willbeavailable}(\mathcal{CS})$ and added to $\mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{CS})$. The oracle $O_j[\tilde{a}]$ is removed from $\mathcal{O}^\infty(\mathcal{I})$; it is also removed from $\mathcal{O}_{\text{call}}(\mathcal{CS})$ if there remains no occurrence of $\text{call}(O_j[\tilde{a}])$ in the thread Th_1^s . So Property 13 is preserved. The oracles coming from $\text{oracles}'(Q')$ are added to $\mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{CS})$ and to $\mathcal{O}_{\text{call-repl}}(Th_1^s)$ or $\mathcal{O}_{\text{call}}(Th_1^s)$ depending on whether they are under replication or not. The oracle $O_j[\tilde{a}]$ is removed from $\mathcal{O}_{\text{call}}(Th_1^s)$ if and only if it is removed from $\mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{CS})$. So Property 14 is preserved. We have $|\mathcal{CT}''| + N_{\text{steps}} \geq N_{\text{steps}}$, so Property 15 holds. The number of calls to O_j increases by 1 and α increases by 1, so Property 16 is preserved. For the oracles $O'_i[\tilde{a}_i]$ ($i \leq l$), when O'_i is under replication, $O'_i[[1, +\infty[\tilde{a}]]$ is added to \mathcal{I} ; Property 17 is obviously satisfied for these oracles because the number of calls to these oracles is not negative. Property 17 for the previously present oracles and Property 18 are preserved, because all components of these inequalities are unchanged. So $\mathcal{E}_1^{\text{CS}} \equiv \mathcal{CT}''$ in this case.

- If $P_4 = \text{return}(M_1, \dots, M_{m'_j})$; Q' , the CryptoVerif process reduces in exactly the same manner as above. The configuration \mathcal{CS}' is the configuration \mathcal{CS} in which we replace the current expression pe^s with $(\mathbb{G}_{\text{val}T'_{j,1}}(c_1), \dots, \mathbb{G}_{\text{val}T'_{j,m'_j}}(c_{m'_j}))$, the set \mathcal{I} with $\mathcal{I}' \stackrel{\text{def}}{=} \mathcal{I} \setminus \{O_j[\tilde{a}]\}$, and the set \mathcal{RI} with $\mathcal{RI}' \stackrel{\text{def}}{=} \mathcal{RI} \cup \{\text{role}[\tilde{a}] \mid (\mu_{\text{role}}, \text{false}) \in \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q')\} \cup \{\text{role}[[1, +\infty[\tilde{a}]] \mid (\mu_{\text{role}}, \text{true}) \in \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q')\}$.

Let Th_4^o be the current thread of \mathcal{C}_4 . We have

$$\begin{aligned}
Th_4^o &= \langle env_4^o, \text{return}(\mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q'), (\mathbb{G}_{\mathcal{M}}(M_1), \dots, \mathbb{G}_{\mathcal{M}}(M_{m'_j}))), \\
& \quad stack^o, store_4^o \rangle \\
&\rightarrow \langle env_4^o, (\mathbb{G}_{\mathcal{M}}(M_1), \dots, \mathbb{G}_{\mathcal{M}}(M_{m'_j})), stack_1^o, store_4^o \rangle \\
& \quad \text{where } stack_1^o \stackrel{\text{def}}{=} (env_4^o, \text{return}(\mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q'), [\cdot])) :: stack^o
\end{aligned}$$

$$\rightarrow^* Th_5^o \stackrel{\text{def}}{=} \langle env_4^o, (\mathbb{G}_{\text{val}T'_{j,1}}(c_1), \dots, \mathbb{G}_{\text{val}T'_{j,m'_j}}(c_{m'_j})), stack_1^o, store_5^o \rangle$$

by Lemma 6.8 applied m'_j times

where $store_5^o \supseteq store_4^o$ and $\mathcal{C}_4 \rightarrow^* \mathcal{C}_4[\text{Th} \mapsto Th_5^o] \rightarrow^* \mathcal{C}_5 \stackrel{\text{def}}{=} \mathcal{C}_4[\text{Th} \mapsto Th_6^o, \text{MI} \mapsto \mathcal{M}\mathcal{I}']$ where $Th_6^o \stackrel{\text{def}}{=} \langle env_4^o, (\mathbb{G}_{\text{val}T'_{j,1}}(c_1), \dots, \mathbb{G}_{\text{val}T'_{j,m'_j}}(c_{m'_j})), stack^o, store_5^o \rangle$ and $\mathcal{M}\mathcal{I}' \stackrel{\text{def}}{=} \mathcal{M}\mathcal{I} \cup \mathbb{G}_{\text{get}_{\mathcal{M}\mathcal{I}}}(Q')$. Let $\mathcal{C}\mathcal{T}''$ be the trace $\mathcal{C}\mathcal{T}$ followed by $\mathcal{C}\mathcal{T}'$ and extended until \mathcal{C}_5 .

Let us now prove that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{C}\mathcal{T}''$. By the form of $\mathfrak{C}_1^{\text{cs}}$ and \mathcal{C}'' , Properties 1 and 2 hold. The set \mathcal{Q}_2 is the set \mathcal{Q} from which we removed the oracles $O_j[\tilde{a}]$ and $O_{\text{loop}}[\alpha + 1]$ and to which we added the new oracles of $\text{reduce}'(Q')$. The set \mathcal{I}' is the set \mathcal{I} from which we removed $O_j[\tilde{a}]$. We added the elements of $\text{oracles}'(Q')$ to $\mathcal{O}(\mathcal{R}\mathcal{I}')$. So Property 3 is preserved. Properties 4 to 10, 12, and 15 to 18 are proved as in the case $P_4 = \text{end}$. We added matching elements in $\mathcal{M}\mathcal{I}'$ and in $\mathcal{R}\mathcal{I}'$, so Property 11 is preserved. The oracles coming from $\text{oracles}'(Q')$ are removed from $\text{willbeavailable}(\mathcal{C}\mathcal{S})$ and added to $\mathcal{O}^\infty(\mathcal{R}\mathcal{I})$. The oracle $O_j[\tilde{a}]$ is removed from $\mathcal{O}^\infty(\mathcal{I})$; it is also removed from $\mathcal{O}_{\text{call}}(\mathcal{C}\mathcal{S})$ if there remains no occurrence of $\text{call}(O_j[\tilde{a}])$ in the thread Th_1^s . So Property 13 is preserved. The oracle $O_j[\tilde{a}]$ is removed from $\mathcal{O}_{\text{call}}(Th_1^s)$ if and only if it is removed from $\mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{C}\mathcal{S})$, and the other oracle sets of Property 14 are unchanged, so Property 14 is preserved. So $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{C}\mathcal{T}''$ in this case.

Case 2.2. The current expression of $\mathcal{C}\mathcal{S}$ is $pe^s = \text{call}(O_j[_, \tilde{a}]) (v_1, \dots, v_{m_j})$ and $\mathcal{C}\mathcal{S}$ cannot reduce, that is, the configuration $\mathcal{C}\mathcal{S}$ makes a successful call to $O_j[_, \tilde{a}]$, an oracle under replication. We prove this case by a reasoning similar to the previous case.

We show that a copy of the oracle $O_j[_, \tilde{a}]$ is available in \mathcal{Q} using Property 17, as follows. By Property 6(b)i, $pe^o = \text{tagfunction}^{O_j, \tau} [env_1^o, pm_{\text{true}}(Q)] (v_1, \dots, v_{m_j})$, with $\tau = \tau_{\mathcal{O}}(O_j[_, \tilde{a}])$ and $O[[a', +\infty[_, \tilde{a}]] \in \mathcal{I}$ for some a' . Let $\mathcal{C}\mathcal{T}'$ be the extension of $\mathcal{C}\mathcal{T}$ with one step. By definition of N_{O_j} , we have $N_{O_j} \geq N_{\text{calls}}(O_j, \tau, \mathcal{C}\mathcal{T}') = N_{\text{calls}}(O_j, \tau, \mathcal{C}\mathcal{T}) + 1$, so by Property 17, $a' \leq N_{O_j}$, so $O_j[a', \tilde{a}] \in \mathcal{O}(\mathcal{I})$, so by Property 3, \mathcal{Q}_0 contains a process of the form $O_j[a', \tilde{a}](x_1[a', \tilde{a}] : T_{j,1}, \dots, x_{m_j}[a', \tilde{a}] : T_{j,m_j}) := P_{\mathcal{O}}$.

Due to the call, the index a' such that $O[[a', +\infty[_, \tilde{a}]] \in \mathcal{I}$ increases by 1 and the number of calls to the closure with tag O_j, τ increases by 1, so Property 17 is preserved.

Case 2.3. The current expression of $\mathcal{C}\mathcal{S}$ is $pe^s = \text{random}()$, that is, the configuration $\mathcal{C}\mathcal{S}$ samples a random boolean. By Property 6, the current expression of \mathcal{C} is $pe^o = \text{random}()$. For $b \in \{\text{true}, \text{false}\}$, $\mathcal{C} \rightarrow_{1/2} \mathcal{C}_b$ where $\mathcal{C}_b \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto \langle env^o, b, stack^o, store^o \rangle]$. Let $\mathcal{C}\mathcal{T}_b$ be the extension of the trace $\mathcal{C}\mathcal{T}$ until \mathcal{C}_b .

The configuration $\mathcal{C}\mathcal{S}$ cannot reduce, and $\text{simreturn}(\mathcal{C}\mathcal{S}) = (\text{repr}(\mathcal{C}\mathcal{S}), o_R, (), ())$. Let us denote $s \stackrel{\text{def}}{=} \text{repr}(\mathcal{C}\mathcal{S})$. The simulator configuration reduces in the

following way for a CryptoVerif value $b \in \{\text{true}, \text{false}\}$.

$$\begin{aligned}
\mathfrak{C}^{\text{cs}} &\rightsquigarrow^* E_1, P'_{loop}\{\alpha/i'\}, \mathcal{T}, \mathcal{Q}, \mathcal{R}_{loop}(\alpha), \mathcal{E} \\
&\quad \text{where } E_1 \stackrel{\text{def}}{=} E[xs[\alpha] \mapsto (s, o_R, (), ()), s'[\alpha] \mapsto s, o[\alpha] \mapsto o_R, \\
&\quad \quad i[\alpha] \mapsto (), args[\alpha] \mapsto ()] \\
&\rightsquigarrow^* E_1, P_{loop}^R\{\alpha/i'\}, \mathcal{T}, \mathcal{Q}, \mathcal{R}_{loop}(\alpha), \mathcal{E} \\
&\rightsquigarrow_{1/2} E_2, \text{return}(\text{simulate}_R(s'[\alpha], b_R[\alpha]), \text{true}), \mathcal{T}, \mathcal{Q}, \mathcal{R}_{loop}(\alpha), \mathcal{E} \\
&\quad \text{where } E_2 \stackrel{\text{def}}{=} E_1[b_R[\alpha] \mapsto b] \\
&\rightsquigarrow^* \mathfrak{C}_b^{\text{cs}} \stackrel{\text{def}}{=} E_3, P_{loop}\{\alpha + 1/i'\}, \mathcal{T}, \mathcal{Q}_1, \mathcal{R}_{loop}(\alpha + 1), \mathcal{E}, N_{\text{steps}}, \mathcal{CS}_b \\
&\quad \text{where } E_3 \supseteq E_2[s[\alpha + 1] \mapsto \text{repr}(\mathcal{CS}_b)], \\
&\quad \quad \mathcal{Q}_1 \stackrel{\text{def}}{=} \mathcal{Q} \setminus \{Q_{loop}\{\alpha + 1/i'\}\}, \\
&\quad \quad \mathcal{CS}_b \stackrel{\text{def}}{=} \mathcal{CS}[\text{Th} \mapsto \langle env^s, \mathbb{G}_{\text{valbool}}(b), stack^s, store^s \rangle]
\end{aligned}$$

We verify that $Q_{loop}\{\alpha + 1/i'\} \in \mathcal{Q}$ using Property 16, as follows. By definition of $N_{\text{rand+calls}}$, $N_{\text{rand+calls}} \geq \left(N_{\text{rand}}(\mathcal{CT}_b) + \sum_{O, \tau} N_{\text{calls}}(O, \tau, \mathcal{CT}_b)\right) + 1 = \left(N_{\text{rand}}(\mathcal{CT}) + \sum_{O, \tau} N_{\text{calls}}(O, \tau, \mathcal{CT})\right) + 2$ since \mathcal{CT}_b makes one more random number generation than \mathcal{CT} . So by Property 16, $N_{\text{rand+calls}} \geq \alpha + 1$. So by Property 3, $Q_{loop}\{\alpha + 1/i'\} \in \mathcal{Q}$. Moreover, we have $\mathbb{G}_{\text{valbool}}(b) = b$, so $\mathcal{CS}_b = \mathcal{CS}[\text{Th} \mapsto \langle env^s, b, stack^s, store^s \rangle]$.

In this step, α becomes $\alpha + 1$, the number of random number generations in the trace increases by 1, the current thread is modified exactly in the same manner on both sides, and the other threads, the oracle sets, the global store, and the events are left unchanged, so it is easy to see that $\mathfrak{C}_{\text{true}}^{\text{cs}} \equiv \mathcal{CT}_{\text{true}}$ and $\mathfrak{C}_{\text{false}}^{\text{cs}} \equiv \mathcal{CT}_{\text{false}}$.

Case 2.4. The configuration \mathcal{CS} does not reduce, and does not make a call to an oracle nor sample a random boolean. In this case, $\text{simreturn}(\mathcal{CS}) = (\text{repr}(\mathcal{CS}), o_S, (), ())$. Let us denote $s \stackrel{\text{def}}{=} \text{repr}(\mathcal{CS})$. The simulator configuration reduces in the following way.

$$\begin{aligned}
\mathfrak{C}^{\text{cs}} &\rightsquigarrow^* E_1, P'_{loop}\{\alpha/i'\}, \mathcal{T}, \mathcal{Q}, \mathcal{R}_{loop}(\alpha), \mathcal{E} \\
&\quad \text{where } E_1 \stackrel{\text{def}}{=} E[xs[\alpha] \mapsto (s, o_S, (), ()), s'[\alpha] \mapsto s, o[\alpha] \mapsto o_S, i[\alpha] \mapsto (), \\
&\quad \quad args[\alpha] \mapsto ()] \\
&\rightsquigarrow E_1, \text{return}(s'[\alpha], \text{false}), \mathcal{T}, \mathcal{Q}, \mathcal{R}_{loop}(\alpha), \mathcal{E} \\
&\rightsquigarrow E_2, P_{\text{return-loop}}(\alpha), \mathcal{T}, \mathcal{Q}, \mathcal{R}_1, \mathcal{E} \\
&\quad \text{where } E_2 \stackrel{\text{def}}{=} E_1[r'_{\alpha, r}[] \mapsto s, b_{\alpha, r}[] \mapsto \text{false}], \\
&\quad \quad \mathcal{R}_1 \stackrel{\text{def}}{=} [x[], \text{return}(x[]), \text{end}] \\
&\rightsquigarrow E_2, r[] \leftarrow r'_{\alpha, r}[]; \text{end}, \mathcal{T}, \mathcal{Q}, \mathcal{R}_1, \mathcal{E} \\
&\rightsquigarrow E_3, \text{end}, \mathcal{T}, \mathcal{Q}, \mathcal{R}_1, \mathcal{E} \\
&\quad \text{where } E_3 \stackrel{\text{def}}{=} E_2[r[] \mapsto s]
\end{aligned}$$

$$\rightsquigarrow \mathfrak{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E_3, \text{end}, \mathcal{T}, \mathcal{Q}, [], \mathcal{E}$$

This configuration cannot reduce. By Property 6, the OCaml configuration \mathcal{C} also cannot reduce. (If it could reduce, then the simulator configuration \mathcal{CS} would reduce by the same rule as the OCaml configuration.) Moreover, by Property 12 of the invariant, $\text{events} = \mathbb{G}_{\text{ev}}(\mathcal{E})$, so this case satisfies the second point of Lemma 6.29.

Case 2.5. The current expression of \mathcal{CS} is of the form $pe^s = \text{tagfunction}^{t,\tau}[\text{env}, pm] v$, that is, \mathcal{CS} calls a tagged closure. By Property 6(b)iii, the only tagged closures present in the current thread are of the form $\text{tagfunction}^{\text{role},\tau}[\text{env}_1^s, pm'_{\text{role}[\tilde{a}]}]$ for a given role $\text{role}[\tilde{a}]$, with $\text{env}_{\text{prim}} \subseteq \text{env}_1^s$. Such a closure corresponds to the initialization of the role $\text{role}[\tilde{a}]$. Since our programs are well-typed, and these closures expect an argument of type unit, the current expression of \mathcal{CS} is $pe^s = \text{tagfunction}^{\text{role},\tau}[\text{env}_1^s, pm'_{\text{role}[\tilde{a}]}] ()$.

Let us denote by Q_i, b_i for $i \leq m$ the oracles present in $\text{reduce}'(Q(\text{role})[\tilde{a}])$, and let $\tilde{a}_i = \tilde{a}$ or $_$, \tilde{a} such that $O'_i[\tilde{a}_i]$ is the oracle associated to Q_i, b_i in $\text{oracles}'(Q(\text{role})[\tilde{a}])$.

By Property 6(b)i,

$$pe^o = \text{tagfunction}^{\text{role},\tau}[\text{env}_1^o, pm_{\mu_{\text{role}}}] () .$$

By Property 6(b)ii, $\text{env}_1^s(\text{token}) = l_{\text{init-tok}}(\text{role}[\tilde{a}])$ is a location. Let us denote l this location. By Property 6(b)i, we have $\text{env}_1^o(\text{token}) = \text{env}_1^s(\text{token}) = l$. By Property 5, l is in the domain of store^s . By Property 6(b)i, l is also in the domain of store^o .

Let $x_1[], \dots, x_k[]$ be the free variables of the role role .

Let us denote

$$\begin{aligned} pe_e^s &\stackrel{\text{def}}{=} (\text{call}(O'_1[\tilde{a}_1]), \dots, \text{call}(O'_m[\tilde{a}_m])) \\ pe_e^o &\stackrel{\text{def}}{=} \mathbb{G}_{\text{read}}(x_1[]) \text{ in } \dots \text{ in } \mathbb{G}_{\text{read}}(x_k[]) \text{ in} \\ &\quad (\mathbb{G}_{\text{O}}(Q_1, b_1), \dots, \mathbb{G}_{\text{O}}(Q_m, b_m)) \end{aligned}$$

The simulator reduces as follows:

$$\begin{aligned} Th^s &= \langle \text{env}^s, \text{tagfunction}^{\text{role},\tau}[\text{env}_1^s, pm'_{\text{role}}] (), \text{stack}^s, \text{store}^s \rangle \\ &\rightarrow \langle \text{env}_1^s, \text{match} () \text{ with } pm'_{\text{role}}, \text{stack}^s, \text{store}^s \rangle \\ \rightarrow Th_1^s &\stackrel{\text{def}}{=} \langle \text{env}_1^s, pe_1^s, \text{stack}^s, \text{store}^s \rangle \\ &\quad \text{where } pe_1^s \stackrel{\text{def}}{=} \text{if } !\text{token} \text{ then } (\text{token} := \text{false}; pe_e^s) \text{ else raise Bad_Call} \end{aligned}$$

and the OCaml side reduces as follows:

$$\begin{aligned} Th^o &= \langle \text{env}^o, \text{tagfunction}^{\text{role},\tau}[\text{env}_1^o, pm_{\mu_{\text{role}}}] (), \text{stack}^o, \text{store}^o \rangle \\ &\rightarrow \langle \text{env}_1^o, \text{match} () \text{ with } pm_{\mu_{\text{role}}}, \text{stack}^o, \text{store}^o \rangle \\ \rightarrow Th_1^o &\stackrel{\text{def}}{=} \langle \text{env}_1^o, pe_1^o, \text{stack}^o, \text{store}^o \rangle \\ &\quad \text{where } pe_1^o \stackrel{\text{def}}{=} \text{if } !\text{token} \text{ then } (\text{token} := \text{false}; pe_e^o) \text{ else raise Bad_Call} \end{aligned}$$

- If $store^s(l) = \text{false}$, then by Property 6(b)i, $store^o(l) = \text{false}$, so

$$\begin{aligned} Th_1^s \rightarrow^* Th_2^s &\stackrel{\text{def}}{=} \langle env_1^s, \text{raise Bad_Call}, stack^s, store^s \rangle \\ Th_1^o \rightarrow^* Th_2^o &\stackrel{\text{def}}{=} \langle env_1^o, \text{raise Bad_Call}, stack^s, store^s \rangle \end{aligned}$$

Let \mathcal{CT}_1 be the extension of the trace \mathcal{CT} until $\mathcal{C}[\text{Th} \mapsto Th_2^o]$, $\mathcal{CS}_1 \stackrel{\text{def}}{=} \mathcal{CS}[\text{Th} \mapsto Th_2^s]$, $steps^s$ the number of steps of the trace $\mathcal{CS} \rightarrow^* \mathcal{CS}_1$, and $\mathfrak{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, steps - steps^s, \mathcal{CS}_1$. We have $\mathfrak{C}_1^{\text{cs}} \rightsquigarrow^+ \mathfrak{C}_1^{\text{cs}}$.

Let us prove that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}_1$. As the current expression is an exceptional value, `replacecalls` allows any environment in its image. Moreover, the other elements of the configuration are the same and \mathcal{I} did not change, so Property 6 is preserved. The number of steps in the reduction is the same on both sides, so Property 15 is preserved. All other properties of Definition 6.27 are trivially inherited from $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}$.

- Otherwise, $store^s(l) = \text{true}$. By Property 6(b)i, $store^o(l) = \text{true}$.

On the simulator side:

$$\begin{aligned} Th_1^s \rightarrow^* Th_3^s &\stackrel{\text{def}}{=} \langle env_1^s, pe_e^s, stack^s, store_1^s \rangle \\ &\text{where } store_1^s \stackrel{\text{def}}{=} store^s[l \mapsto \text{false}] \end{aligned}$$

By Property 4, the variables $x_1[], \dots, x_k[]$ are present in the environment E . Let a'_1, \dots, a'_k be the values of these variables in the environment E . By Property 9, $globalstore(E, \mathcal{T}) \subseteq globalstore^o$, so $globalstore^o(f_i) = \text{ser}(T_{x_i}, a'_i)$ where $(x_i[], f_i) \in \text{files}$ for all $i \leq k$. Let $\mathcal{C}_1 \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto Th_1^o]$. We have

$$\begin{aligned} \mathcal{C}_1 &\rightarrow^* \mathcal{C}[\text{Th} \mapsto \langle env_1^o, pe_e^o, stack^o, store_1^o \rangle] \\ &\text{where } store_1^o \stackrel{\text{def}}{=} store^o[l \mapsto \text{false}] \\ &\rightarrow^* \mathcal{C}_2 \stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto \langle env_2^o, pe_2^o, stack^o, store_2^o \rangle] \\ &\text{where } pe_2^o \stackrel{\text{def}}{=} (\mathbb{G}_O(Q_1, b_1), \dots, \mathbb{G}_O(Q_m, b_m)), \\ &\quad env_2^o \stackrel{\text{def}}{=} env_1^o[\mathbb{G}_{\text{var}}(x_1) \mapsto \mathbb{G}_{\text{val}T_{x_1}}(a'_1), \dots, \\ &\quad \quad \quad \mathbb{G}_{\text{var}}(x_k) \mapsto \mathbb{G}_{\text{val}T_{x_k}}(a'_k)], \\ &\quad store_2^o \supseteq store_1^o \end{aligned}$$

by Proposition 6.5 applied k times to show the correctness of the deserialization primitives.

Let l_1, \dots, l_m be pairwise distinct locations that are not in $\text{Dom}(store_2^o)$ and τ_1, \dots, τ_m be pairwise distinct fresh tags. By the same reasoning as

in Case 2.1, sub-case $P_4 = \text{return}(M_1, \dots, M_{m'_j}); Q'$, we have

$$\begin{aligned} \mathcal{C}_2 \rightarrow^* \mathcal{C}_3 &\stackrel{\text{def}}{=} \mathcal{C}[\text{Th} \mapsto \langle \text{env}_2^o, pe_3^o, \text{stack}^o, \text{store}_3^o \rangle] \\ &\text{where } pe_3^o \stackrel{\text{def}}{=} (\text{tagfunction}^{O'_1, \tau_1}[\text{env}_{c,1}^o, pm_{b_1}(Q_1)], \dots, \\ &\quad \text{tagfunction}^{O'_m, \tau_m}[\text{env}_{c,m}^o, pm_{b_m}(Q_m)]), \\ &\quad \text{store}_3^o \stackrel{\text{def}}{=} \text{store}_2^o \cup \{l_i \mapsto \text{true} \mid i \leq m, b_i = \text{false}\} \end{aligned}$$

where, for all $i \leq m$, $\text{env}_{c,i}^o$ is env_2^o when b_i is true and $\text{env}_2^o[\text{token} \mapsto l_i]$ otherwise.

Let \mathcal{CT}_2 be an extension of the trace \mathcal{CT} until \mathcal{C}_3 . Let $\mathcal{CS}_3 \stackrel{\text{def}}{=} \mathcal{CS}[\text{Th} \mapsto \text{Th}_3^s]$. Let steps^s be the number of steps of $\mathcal{CS} \rightarrow^* \mathcal{CS}_3$. Let $\mathfrak{C}_2^{\text{cs}} \stackrel{\text{def}}{=} E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, \text{steps} - \text{steps}^s, \mathcal{CS}_3$. We have $\mathfrak{C}^{\text{cs}} \rightsquigarrow^+ \mathfrak{C}_2^{\text{cs}}$.

Let us prove that $\mathfrak{C}_2^{\text{cs}} \equiv \mathcal{CT}_2$. We define τ'_O as τ_O except that for all $i \leq m$, if $b_i = \text{true}$, then $\tau'_O(O'_i[_, \tilde{a}]) = \tau_i$. Properties 1, 2, 3, 4, 5, 7, 8, 9, 10, 11, 12, 16, 18 hold because they hold for $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$.

Let us prove Property 6. First, we prove Property 6(b) for the current thread. For all $i \leq m$, the free variables of Q_i are contained in $\{x_1[], \dots, x_k[]\}$, so $\text{env}_{c,i}^o \supseteq \text{env}(E, Q_i)$. Moreover, by Properties 6(b)iii and 6(b)i, $\text{env}_{\text{prim}} \subseteq \text{env}_1^o$, so $\text{env}_{c,i}^o \supseteq \text{env}_2^o \supseteq \text{env}_1^o \supseteq \text{env}_{\text{prim}}$. We have $\text{role}[\tilde{a}] \in \mathcal{R}_{\text{init-closure}}(\text{Th}^s)$. By Property 14, $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(\text{Th}^s))$ is included in $\mathcal{O}^\infty(\mathcal{I}) \cup \mathcal{O}_{\text{call}}(\mathcal{CS})$, and furthermore $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(\text{Th}^s))$ is disjoint from $\mathcal{O}_{\text{call}}(\text{Th}_i)$ for all $i \leq n$, so from $\mathcal{O}_{\text{call}}(\mathcal{CS})$, so $\mathcal{O}^\infty(\{\text{role}[\tilde{a}]\})$ is included in $\mathcal{O}^\infty(\mathcal{I})$. Hence, when O'_i is not under replication (that is, $b_i = \text{false}$), $O'_i[\tilde{a}_i] \in \mathcal{I}$, and when O'_i is under replication, $\tilde{a}_i = _, \tilde{a}$ and $O'_i[[1, +\infty[\tilde{a}]] \in \mathcal{I}$. By Property 3, when O'_i is not under replication, $Q_i = \mathcal{Q}(O'_i[\tilde{a}_i])$, and when O'_i is under replication, $Q_i = \mathcal{Q}(O'_i[1, \tilde{a}])$.

By Property 6(b)i, there exist store_4^o and l_{tok} such that

$$\begin{aligned} &\langle \text{env}^o, pe^o, \text{stack}^o, \text{store}_4^o \rangle \\ &\in \text{replacecalls}(\text{replaceinitpm}(\text{Th}^s), \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_O) \\ &\text{store}_4^o \cup \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(\text{Th}^s), l_{\text{tok}}) \subseteq \text{store}^o. \end{aligned}$$

Since $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(\text{Th}^s))$ is disjoint from $\mathcal{O}_{\text{call}}(\mathcal{CS})$ as noticed above, the oracles $O'_i[\tilde{a}_i]$ are not present in $\mathcal{O}_{\text{call}}(\mathcal{CS})$. So we can define the injective function $l'_{\text{tok}} \stackrel{\text{def}}{=} l_{\text{tok}} \cup \{O'_i[\tilde{a}_i] \mapsto l_i \mid i \leq m, b_i = \text{false}\}$. By Lemma F.3, Item 2, for all $\text{call}(R)$ that occur in $\text{replaceinitpm}(\text{Th}^s)$, $\text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}, l'_{\text{tok}}, \tau'_O) \supseteq \text{correctclosure}(R, \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_O)$, noticing that, when $i \leq m$ and $b_i = \text{true}$, $O'_i[N_{O'_i} + 1, \tilde{a}] \in \mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(\text{Th}^s))$, so by Property 14, $O'_i[N_{O'_i} + 1, \tilde{a}] \notin \mathcal{O}_{\text{call-repl}}(\text{Th}^s)$, so $\text{call}(O'_i[_, \tilde{a}])$ does not occur in $\text{replaceinitpm}(\text{Th}^s)$, so the transformation of τ_O into τ'_O does not affect the computation of these correct closures. Moreover, $\text{tagfunction}^{O'_i, \tau_i}[\text{env}_{c,i}^o, pm_{b_i}(Q_i)] \in \text{correctclosure}(O'_i[\tilde{a}_i], \mathcal{I}, E, \mathcal{Q}, l'_{\text{tok}}, \tau'_O)$ for $i \leq m$ and pe_e^s is a value so replacecalls allows any environment in the threads it returns,

so $\langle env_2^o, pe_3^o, stack^o, store_4^o[l \mapsto \text{false}] \rangle \in \text{replacecalls}(\text{replaceinitpm}(Th_3^s), \mathcal{I}, E, \mathcal{Q}, l'_{\text{tok}}, \tau'_O)$. We have

$$\begin{aligned}
& store_4^o[l \mapsto \text{false}] \cup \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th_3^s), l'_{\text{tok}}) \\
& \subseteq store_4^o[l \mapsto \text{false}] \cup \text{gettokens}(\mathcal{I}, \mathcal{O}_{\text{call}}(Th^s), l_{\text{tok}}) \\
& \quad \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} \\
& \subseteq store^o[l \mapsto \text{false}] \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} \\
& \subseteq store_1^o \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} \\
& \subseteq store_2^o \cup \{l_i \mapsto \text{true} \mid b_i = \text{false}\} = store_3^o,
\end{aligned}$$

so Property 6(b)i holds with the function l'_{tok} . Properties 6(b)ii, 6(b)iii, 6(b)iv are preserved, so Property 6 holds for the current thread. The other threads and $\mathcal{I}, E, \mathcal{Q}$ are unchanged, and as above, the transformation of τ_O into τ'_O does not affect the computation of correct closures in these threads, so Property 6 holds for all threads.

The role $\text{role}[\tilde{a}]$ is removed from $\mathcal{R}_{\text{init-closure}}(Th^s)$, so the elements added to $\mathcal{O}_{\text{call}}(Th^s)$ and $\mathcal{O}_{\text{call-repl}}(Th^s)$ are removed from $\mathcal{O}^\infty(\mathcal{R}_{\text{init-closure}}(Th^s))$, hence Properties 13 and 14 are preserved. There are more steps on the OCaml side than on the CryptoVerif side, so Property 15 is preserved. For the oracles $O'_i[\tilde{a}_i]$ ($i \leq l$), when O'_i is under replication, we have already shown that $O'_i[[1, +\infty[\tilde{a}]] \in \mathcal{I}$; Property 17 is obviously satisfied for these oracles because the number of calls to these oracles is not negative. Property 17 is preserved for the other oracles, because all components of these inequalities are unchanged.

Case 2.6. The current expression of \mathcal{CS} is of the form $pe^s = \text{addthread}(\text{program})$, that is, we add a new thread to the current configuration. By Property 6(b)i, the expression pe^o is $\text{addthread}(\text{program})$, and by Property 6(b)iv, program contains no closure, no tagged function, no event, no return except in parts $\text{program}(\mu_{\text{role}})$, and in $\text{program}(\mu_{\text{role}})$ in arguments of addthread .

Suppose first that program is an attacker program: it does not contain $\text{program}(\mu_{\text{role}})$ except in arguments of addthread . In this case,

$$\begin{aligned}
\mathcal{CS} & \rightarrow \mathcal{CS}_1 \stackrel{\text{def}}{=} ([Th_1, \dots, Th_{t_j-1}, \langle env^s, (), stack^s, store^s \rangle, Th_{t_j+1}, \dots, Th_n, \\
& \quad \langle \emptyset, \text{program}, [], \emptyset \rangle], globalstore^s, t_j), \mathcal{RI}, \mathcal{I} \\
\mathcal{C} & \rightarrow \mathcal{C}_1 \stackrel{\text{def}}{=} [Th'_1, \dots, Th'_{t_j-1}, \langle env^o, (), stack^o, store^o \rangle, Th_{t_j+1}, \dots, Th_n, \\
& \quad \langle \emptyset, \text{program}, [], \emptyset \rangle], globalstore^o, t_j, \mathcal{MI}, \text{events}
\end{aligned}$$

Let \mathcal{CT}_1 be the extension of the trace \mathcal{CT} until \mathcal{C}_1 and $\mathfrak{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, \text{steps} - 1, \mathcal{CS}_1$. We have $\mathfrak{C}^{\text{cs}} \rightsquigarrow \mathfrak{C}_1^{\text{cs}}$. Let us prove that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}_1$. The new thread contains no closures and no tagged functions. It contains no call since program is an OCaml program (not a simulator program), so it satisfies Property 6(b). The other properties are inherited from $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$.

Otherwise, the program $program$ is of the form

$$program_{\text{prim}};; program(\mu_{\text{role}_1});; \dots;; program(\mu_{\text{role}_m});; program',$$

where $program'$ does not contain $program(\mu)$ for any $\mu \in \mathcal{M}_g$. By Assumption 4.1, for $\mathcal{M} \stackrel{\text{def}}{=} \{\mu_{\text{role}_1}, \dots, \mu_{\text{role}_m}\}$, we have $\mathcal{M} \subseteq \mathcal{M}_g$ and $\forall \mu \in \mathcal{M}, \exists b, (\mu, b) \in \mathcal{MI}$. By Property 11, for each $i \leq m$, if role_i is not under replication, then the set \mathcal{RI} contains $\text{role}_i[\tilde{a}]$ for some \tilde{a} , and if role_i is under replication, then the set \mathcal{RI} contains $\text{role}_i[[a', +\infty[\tilde{a}]]$ for some a', \tilde{a} . By Property 13, the oracles present in \mathcal{RI} are not in \mathcal{I} . We can then define

$$\begin{aligned} \tilde{a}_1 &\stackrel{\text{def}}{=} \text{smallest}(\mathcal{RI}, \text{role}_1), \dots, \tilde{a}_m \stackrel{\text{def}}{=} \text{smallest}(\mathcal{RI}, \text{role}_m) \\ \mathcal{RI}'' &\stackrel{\text{def}}{=} \{\text{role}_1[\tilde{a}_1], \dots, \text{role}_m[\tilde{a}_m]\} \\ \mathcal{RI}' &\stackrel{\text{def}}{=} \mathcal{RI} - \mathcal{RI}'' \quad \mathcal{I}' \stackrel{\text{def}}{=} \text{add}(\mathcal{I}, \mathcal{RI}'') \\ program^b &\stackrel{\text{def}}{=} program_{\text{prim}};; program'(\text{role}_1[\tilde{a}_1]);; \dots;; program'(\text{role}_m[\tilde{a}_m]);; \\ &\quad program' \\ \mathcal{MI}' &\stackrel{\text{def}}{=} \{(\mu, \text{false}) \mid \mu \in \mathcal{M} \wedge (\mu, \text{false}) \in \mathcal{MI}\} \end{aligned}$$

We have

$$\begin{aligned} \mathcal{CS} &\rightarrow \mathcal{CS}_2 \stackrel{\text{def}}{=} ([Th_1, \dots, Th_{t_j-1}, \langle env^s, (), stack^s, store^s \rangle, Th_{t_j+1}, \dots, Th_n, \\ &\quad \langle \emptyset, program^b, [], \emptyset \rangle, globalstore^s, t_j), \mathcal{RI}', \mathcal{I}' \\ \mathcal{C} &\rightarrow \mathcal{C}_2 \stackrel{\text{def}}{=} [Th'_1, \dots, Th'_{t_j-1}, \langle env^o, (), stack^o, store^o \rangle, Th_{t_j+1}, \dots, Th_n, \\ &\quad \langle \emptyset, program, [], \emptyset \rangle, globalstore^o, t_j, \mathcal{MI} \setminus \mathcal{MI}', events \end{aligned}$$

Let \mathcal{CT}_2 be the extension of the trace \mathcal{CT} until \mathcal{C}_2 and $\mathfrak{C}_2^{\text{cs}} \stackrel{\text{def}}{=} E, P, T, Q, \mathcal{R}, \mathcal{E}, \text{steps} - 1, \mathcal{CS}_2$. We have $\mathfrak{C}^{\text{cs}} \rightsquigarrow \mathfrak{C}_2^{\text{cs}}$. Let us prove that $\mathfrak{C}_2^{\text{cs}} \equiv \mathcal{CT}_2$. The oracles under replication added to \mathcal{I} are the oracles $O[[1, +\infty[\tilde{a}_i]]$ such that $O[_, \tilde{a}_i] \in \text{oracles}'(Q(\text{role}_i)[\tilde{a}_i])$ for any $i \leq m$. We define τ'_0 as the extension of τ_0 that maps all the oracles $O[_, \tilde{a}_i]$ to fresh distinct tags τ . Properties 1, 2, 4, 8, 9, 10, 12, 15, and 16 are inherited from $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$. By Property 3, $Q = \{Q_{\text{loop}}\{a/i'\} \mid \alpha < a \leq N_{\text{rand+calls}}\} \cup Q_0$ and $Q_0 \leftrightarrow \mathcal{RI}, \mathcal{I}$. If role_i is under replication, then by definition of smallest , $\text{role}_i[[a'_i, +\infty[\tilde{a}''_i]] \in \mathcal{RI}$ with $\tilde{a}_i = a'_i, \tilde{a}''_i$. By definition of N_{role_i} , $N_{\text{role}_i} \geq N_{\text{exec}}(\text{role}_i, \mathcal{CT}_2) = N_{\text{exec}}(\text{role}_i, \mathcal{CT}) + 1$. By Property 18, $a'_i \leq N_{\text{role}_i}$. Therefore, the set $\mathcal{O}(\mathcal{RI})$ contains the first oracles of $\text{role}_i[\tilde{a}_i]$ for $i \leq m$. The set $\mathcal{O}(\mathcal{RI}')$ is the set $\mathcal{O}(\mathcal{RI})$ from which we remove the first oracles of $\text{role}_i[\tilde{a}_i]$ for $i \leq m$ and $\mathcal{O}(\mathcal{I}')$ is the set $\mathcal{O}(\mathcal{I})$ to which we add these oracles. So $Q_0 \leftrightarrow \mathcal{RI}', \mathcal{I}'$ and Property 3 holds. There are no local store locations in $program$, so Property 5 holds. For each thread Th_i of the simulator except the new thread, let us show that Property 6 is preserved. The only changes are that the current expression is replaced with $()$ and that $\mathcal{I}' = \text{add}(\mathcal{I}, \mathcal{RI}'')$, so we just have to show that Property 6(b)i is preserved; the other elements of Property 6 are obviously preserved. By Lemma F.3, Item 2, the correct closures

are preserved. By Property 14, the set $\mathcal{O}_{\text{call}}(Th_i)$ does not contain any of the oracles added to \mathcal{I} , so the tokens are preserved. Hence, Property 6(b)i is preserved. Since $program'$ already occurs in the initial program, it does not contain closures. By Property 6(b)iv, it does not contain tagged functions, events, or returns, except in $program(\mu_{\text{role}})$ in arguments of `addthread`, so Property 6(a) holds for the new thread. By Property 7, $program'$ does not contain any location $l \in S_{\text{priv}}$ except in $program(\mu_{\text{role}})$ in arguments of `addthread`, so Property 7 holds. When role_i is not under replication, we remove one copy of the module μ_{role_i} from the multiset \mathcal{MI} , and correspondingly, we remove $\text{role}_i[\widetilde{a}_i]$ from \mathcal{RI} . When role_i is under replication, we add 1 to the first index of roles $\text{role}_i[\widetilde{a}_i]$ in \mathcal{RI} , and \mathcal{MI} is not affected by this change. (The role role_i can still be called.) So Property 11 is preserved. The first oracles of $\text{role}_1[\widetilde{a}_1], \dots, \text{role}_m[\widetilde{a}_m]$ are transferred from $\mathcal{O}^\infty(\mathcal{RI})$ to $\mathcal{O}^\infty(\mathcal{I})$, so Property 13 is preserved. More precisely, these oracles are added to $\mathcal{O}^\infty(\mathcal{R}_{\text{init-function}}(Th_{n+1}))$, where $Th_{n+1} \stackrel{\text{def}}{=} \langle \emptyset, program^b, [], \emptyset \rangle$ is the new thread, so Property 14 is preserved. For the oracles $O[[1, +\infty[\widetilde{a}_i]$ added to \mathcal{I} , Property 17 is obviously satisfied because the number of calls to an oracle is not negative. Property 17 is preserved for the previously present oracles, because all components of these inequalities are unchanged. For the roles $\text{role}_i[[a'_i, +\infty[\widetilde{a}''_i] \in \mathcal{RI}$, with $\widetilde{a}_i = a'_i, \widetilde{a}''_i$, we have $\text{role}_i[[a'_i + 1, +\infty[\widetilde{a}''_i] \in \mathcal{RI}$; the elements $\text{role}_i[\dots]$ in \mathcal{RI} with indices that do not end with \widetilde{a}''_i are unchanged; and $N_{\text{exec}}(\text{role}_i, \mathcal{CT})$ increases by 1, so Property 18 is preserved for the roles $\text{role}_1, \dots, \text{role}_m$. Property 18 is preserved for the other roles, because all components of the inequalities are unchanged. Therefore, $\mathcal{CS}^s \equiv \mathcal{CT}_2$.

Case 2.7. The current expression of \mathcal{CS} is of the form $pe^s = \text{call}(O[\widetilde{a}]) v$ and \mathcal{CS} reduces, that is, we call an oracle but the call fails. By reduction rule (FailedCall1) or (FailedCall2),

$$Th^s \rightarrow Th_1^s \stackrel{\text{def}}{=} \langle env^s, \text{raise Bad_Call}, stack^s, store^s \rangle,$$

and $\mathcal{CS} \rightarrow \mathcal{CS}_1 \stackrel{\text{def}}{=} \mathcal{CS}[\text{Th} \mapsto Th_1^s]$.

By Property 6(b)i, $pe^o = c v'$, where $c \in \text{correctclosure}(O[\widetilde{a}], \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_O)$.

We suppose that the program is well typed, so the value v is a k -tuple (v_1, \dots, v_k) , where k is the number of arguments of oracle O . Let T_1, \dots, T_k be their CryptoVerif types. Let x_1, \dots, x_k be the CryptoVerif variables that are the arguments of O . By Assumption 6.3, the value v' does not contain closures nor locations, so $v' = v$.

Let us first suppose that the oracle O is under replication. In this case, $\widetilde{a} = _ , \widetilde{a}'$. There exists a'' such that $O[[a'', +\infty[\widetilde{a}'] \in \mathcal{I}$, because otherwise we would have $\text{correctclosure}(O[\widetilde{a}], \mathcal{I}, E, \mathcal{Q}, l_{\text{tok}}, \tau_O) = \emptyset$. The closure c is of the form $\text{tagfunction}^{O, \tau}[env_1^o, pm_{\text{true}}(Q)]$. Let \mathcal{CT}' be the extension of \mathcal{CT} by one step. By definition of N_O , $N_O \geq N_{\text{calls}}(O, \tau, \mathcal{CT}') = N_{\text{calls}}(O, \tau, \mathcal{CT}) + 1$. Hence, by Property (17), $a'' \leq N_O$. Therefore, by definition of correctclosure , $Q = \mathcal{Q}(O[a'', \widetilde{a}'])$. Since (FailedCall2) applies, there exists i such that $\forall a \in T_i$, $v_i \neq \mathbb{G}_{\text{val}T_i}(a)$. By Proposition 6.5, for any environment env , stack $stack$ and

store $store$,

$$\begin{aligned} & \langle env, env_{\text{prim}}(\mathbb{G}_{\text{pred}}(T_i)) v_i, stack, store \rangle \\ \rightarrow^* & \langle env', \text{false}, stack, store' \rangle \text{ where } store' \supseteq store \end{aligned}$$

So,

$$\begin{aligned} Th^o &= \langle env^o, \text{tagfunction}^{O,\tau}[env_1^o, pm_{\text{true}}(Q)] v, stack^o, store^o \rangle \\ \rightarrow^* & \langle env_2^o, pe_2^o, stack^o, store^o \rangle \\ & \text{where } env_2^o \stackrel{\text{def}}{=} env_1^o[\mathbb{G}_{\text{var}}(x_1) \mapsto v_1, \dots, \mathbb{G}_{\text{var}}(x_k) \mapsto v_k], \\ & pe_2^o \stackrel{\text{def}}{=} \text{if } (\mathbb{G}_{\text{pred}}(T_1) \mathbb{G}_{\text{var}}(x_1)) \&\& \dots \&\& (\mathbb{G}_{\text{pred}}(T_k) \mathbb{G}_{\text{var}}(x_k)) \\ & \quad \text{then } (\mathbb{G}_{\text{file}}(x_1[\tilde{i}]); \dots; \mathbb{G}_{\text{file}}(x_k[\tilde{i}]); \mathbb{G}(P)) \\ & \quad \text{else raise Bad_Call} \\ \rightarrow^* & Th_1^o \stackrel{\text{def}}{=} \langle env_2^o, \text{raise Bad_Call}, stack^o, store_1^o \rangle \end{aligned}$$

where $store_1^o \supseteq store^o$ by Proposition 6.5 applied k times.

If the oracle O is not under replication, then (FailedCall1) applies, so either $O[\tilde{a}] \notin \mathcal{I}$ and in this case by Property 6(b)i, $store^o[l_{\text{tok}}(O[\tilde{a}])] = \text{false}$, or there exists i such that $\forall a \in T_i, v_i \neq \mathbb{G}_{\text{val}T_i}(a)$, so we have a reduction similar to the case in which O is under replication.

Let \mathcal{CT}_1 be an extension of the trace \mathcal{CT} until $\mathcal{C}[\text{Th} \mapsto Th_1^o]$ and $\mathfrak{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, \text{steps} - 1, \mathcal{CS}_1$. We have $\mathfrak{C}^{\text{cs}} \rightsquigarrow \mathfrak{C}_1^{\text{cs}}$. Let us prove that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}_1$. The current expression is an exceptional value, so `replacecalls` allows any environment in the current thread, and $store_1^o \supseteq store^o$, so Property 6(b)i is preserved for the current thread. The OCaml side uses more reductions than the simulator side, so Property 15 is preserved. There is one more oracle call, and α and \mathcal{I} are unchanged, so Properties 16 and 17 are preserved. The other properties are inherited from $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$.

Case 2.8. Let us finally deal with the remaining cases. Cases 2.1, 2.2, 2.3, 2.4 present all cases in which \mathcal{CS} does not reduce. Case 2.6 covers the reduction rule (Simulator add thread). So \mathcal{CS} reduces using rule (Simulator toplevel). So let us denote \mathcal{CS}_1 the configuration such that $\mathcal{CS} = \mathcal{C}^s, \mathcal{RI}, \mathcal{I} \rightarrow \mathcal{CS}_1 = \mathcal{C}_1^s, \mathcal{RI}, \mathcal{I}$. Since the case of failed oracle calls is already handled in Case 2.7, $\mathcal{C}^s \rightarrow \mathcal{C}_1^s$ is obtained by rules of the OCaml semantics, not by (FailedCall1) or (FailedCall2).

If $pe^s = \text{schedule}(tj')$, then by Property 6(b)i, $pe^o = \text{schedule}(tj')$, so \mathcal{C}^s and \mathcal{C} reduce in the same way using rules (Toplevel schedule1) or (Toplevel schedule2) for \mathcal{C}^s and (New toplevel schedule1) or (New toplevel schedule2) for \mathcal{C} . Let \mathcal{C}_1 be the configuration such that $\mathcal{C} \rightarrow \mathcal{C}_1$ and \mathcal{CT}_1 be the extension of the trace \mathcal{CT} until \mathcal{C}_1 . Let $\mathfrak{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, \text{steps} - 1, \mathcal{CS}_1$. We have $\mathfrak{C}^{\text{cs}} \rightsquigarrow \mathfrak{C}_1^{\text{cs}}$ and $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}_1$.

In all other cases, \mathcal{C}^s reduces by (Toplevel). By Property 6(b)i, the current thread of the OCaml configuration has the same form as in the simulator

configuration: the semantic rules are parametric in the elements that are replaced by `replaceinitpm` and `replacecalls`, so the OCaml configuration \mathcal{C} reduces by (New toplevel), using a reduction $Th, globalstore \rightarrow_p Th', globalstore'$ obtained by exactly the same semantic rules as on the simulator side. Let \mathcal{C}_1 be the configuration such that $\mathcal{C} \rightarrow \mathcal{C}_1$ and \mathcal{CT}_1 be the extension of the trace \mathcal{CT} until \mathcal{C}_1 . Let $\mathfrak{C}_1^{\text{cs}} \stackrel{\text{def}}{=} E, P, \mathcal{T}, \mathcal{Q}, \mathcal{R}, \mathcal{E}, \text{steps} - 1, \mathcal{CS}_1$. We have $\mathfrak{C}^{\text{cs}} \rightsquigarrow \mathfrak{C}_1^{\text{cs}}$ and we briefly show that $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}_1$.

If the reduction touches a local store location l , then by Properties 6(b)i and 6(b)ii, l cannot be in the image of l_{tok} or $l_{\text{init-tok}}$. Moreover, in all cases, the reduction commutes with `replaceinitpm` and `replacecalls`, so Property 6 holds for $\mathfrak{C}_1^{\text{cs}}$ and \mathcal{CT}_1 . (Since calls to tagged closures are already handled in Case 2.5, we do not consider this case here. This is important, because the reduction would not commute with `replaceinitpm` in this case: `replaceinitpm` replaces the pattern-matching inside the tagged closure before the call, but would not replace it in the reduced configuration.) If the reduction touches the global store, that is, it uses rule (Globalstore2), let l be the concerned location; by Property 7, the location l is not in S_{priv} , and in OCaml the same operation is carried out on l . So in all cases, Properties 7, 8, 9, and 10 hold for $\mathfrak{C}_1^{\text{cs}}$ and \mathcal{CT}_1 . The oracle sets may only decrease, in case a subexpression is removed by reduction, so Properties 13 and 14 are preserved. The reduction is performed in one step on both sides, so Property 15 is preserved. The other properties are inherited from $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$, so $\mathfrak{C}_1^{\text{cs}} \equiv \mathcal{CT}_1$. \square

G Proof of Proposition 6.30

Definition G.1 *The relation $\mathfrak{CT}_1^{\text{cs}}, \dots, \mathfrak{CT}_n^{\text{cs}} \equiv_{\text{t}} \mathcal{CTS}_1, \dots, \mathcal{CTS}_n$ is verified when the following properties hold:*

1. All traces $\mathfrak{CT}_1^{\text{cs}}, \dots, \mathfrak{CT}_n^{\text{cs}}$ start at $\mathfrak{C}_0(Q_0, \text{program}_0)$, and none of these traces is a prefix of another of these traces.
2. The trace sets $\mathcal{CTS}_1, \dots, \mathcal{CTS}_n$ are pairwise disjoint, all traces in these sets start at $\mathfrak{C}_0(Q_0, \text{program}_0)$, and none of these traces is a prefix of another of these traces.
3. $\forall i \leq n, \Pr[\mathfrak{CT}_i^{\text{cs}}] = \Pr[\mathcal{CTS}_i]$.
4. $\sum_{i \leq n} \Pr[\mathfrak{CT}_i^{\text{cs}}] = 1$.
5. For each trace $\mathfrak{CT}_i^{\text{cs}}$, $i \leq n$,
 - (a) either $\mathfrak{CT}_i^{\text{cs}}$ is complete, every trace $\mathcal{CT} \in \mathcal{CTS}_i$ is complete, and the event list \mathcal{E} of the last configuration of $\mathfrak{CT}_i^{\text{cs}}$ and the event list events of the last configuration of \mathcal{CT} verify $\text{events} = \mathbb{G}_{\text{ev}}(\mathcal{E})$,
 - (b) or for every trace $\mathcal{CT} \in \mathcal{CTS}_i$, the last configuration \mathfrak{C}^{cs} of $\mathfrak{CT}_i^{\text{cs}}$ verifies $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$.

The next lemma applies to any traces, so in particular to OCaml traces and CryptoVerif traces.

Lemma G.2 *Let $\mathcal{CT}_1, \dots, \mathcal{CT}_n$ be traces such that none of these traces is a prefix of another of these traces. If $\mathcal{CT}_1'', \dots, \mathcal{CT}_{n'}''$ are extensions of \mathcal{CT}_n such that none of these traces is a prefix of another, then none of the traces $\mathcal{CT}_1, \dots, \mathcal{CT}_{n-1}, \mathcal{CT}_1'', \dots, \mathcal{CT}_{n'}''$ is a prefix of another of these traces.*

In particular, this is true when, for all $i \leq n'$, \mathcal{CT}_i'' is the concatenation of \mathcal{CT}_n and \mathcal{CT}_i' where $\mathcal{CT}_1', \dots, \mathcal{CT}_{n'}'$ are traces that start at the last configuration of \mathcal{CT}_n such that none of these traces is a prefix of another of these traces.

Proof Let us prove the first point. Consider two traces among $\mathcal{CT}_1, \dots, \mathcal{CT}_{n-1}, \mathcal{CT}_1'', \dots, \mathcal{CT}_{n'}''$. If they are both among $\mathcal{CT}_1, \dots, \mathcal{CT}_{n-1}$, they are not prefix of one another by hypothesis. If they are both among $\mathcal{CT}_1'', \dots, \mathcal{CT}_{n'}''$, they are also not prefix of one another by hypothesis. Now consider \mathcal{CT}_i with $i \leq n-1$ and \mathcal{CT}_j'' with $j \leq n'$. If \mathcal{CT}_i was a prefix of \mathcal{CT}_j'' , then either its length is less or equal to the length of \mathcal{CT}_n , so \mathcal{CT}_i would be a prefix of \mathcal{CT}_n , which is impossible by hypothesis, or its length is greater than the length of \mathcal{CT}_n , so \mathcal{CT}_i would be an extension of \mathcal{CT}_n , that is, \mathcal{CT}_n would be a prefix of \mathcal{CT}_i , which is also impossible by hypothesis. If \mathcal{CT}_j'' was a prefix of \mathcal{CT}_i , then a fortiori \mathcal{CT}_n would be a prefix of \mathcal{CT}_i , which is impossible by hypothesis. Hence, none of the traces $\mathcal{CT}_1, \dots, \mathcal{CT}_{n-1}, \mathcal{CT}_1'', \dots, \mathcal{CT}_{n'}''$ is a prefix of another of these traces.

To show the second point, if \mathcal{CT}_i'' was a prefix of \mathcal{CT}_j'' , then \mathcal{CT}_i' would be a prefix of \mathcal{CT}_j' , which is a contradiction. So we can apply the first point in this case. \square

Lemma G.3 *Suppose that $\mathfrak{CT}_1^{\text{cs}}, \dots, \mathfrak{CT}_n^{\text{cs}} \equiv_{\text{t}} \mathcal{CTS}_1, \dots, \mathcal{CTS}_n$. Either all traces $\mathfrak{CT}_1^{\text{cs}}, \dots, \mathfrak{CT}_n^{\text{cs}}$ are complete, or there exist $\mathfrak{CT}_1^{\text{cs}'}, \dots, \mathfrak{CT}_{n'}^{\text{cs}'}$ and $\mathcal{CTS}'_1, \dots, \mathcal{CTS}'_{n'}$ such that there are strictly more reduction steps in traces $\mathfrak{CT}_1^{\text{cs}'}, \dots, \mathfrak{CT}_{n'}^{\text{cs}'}$ than in traces $\mathfrak{CT}_1^{\text{cs}}, \dots, \mathfrak{CT}_n^{\text{cs}}$ and $\mathfrak{CT}_1^{\text{cs}'}, \dots, \mathfrak{CT}_{n'}^{\text{cs}'} \equiv_{\text{t}} \mathcal{CTS}'_1, \dots, \mathcal{CTS}'_{n'}$.*

Proof Suppose that $\mathfrak{CT}_1^{\text{cs}}, \dots, \mathfrak{CT}_n^{\text{cs}} \equiv_{\text{t}} \mathcal{CTS}_1, \dots, \mathcal{CTS}_n$ and there is a trace $\mathfrak{CT}_i^{\text{cs}}$ that is not complete. We can renumber the traces so that the last trace $\mathfrak{CT}_n^{\text{cs}}$ is not complete.

By Property 5(b), the last configuration \mathfrak{C}^{cs} of the trace $\mathfrak{CT}_n^{\text{cs}}$ and all traces $\mathcal{CT} \in \mathcal{CTS}_n$ verify $\mathfrak{C}^{\text{cs}} \equiv \mathcal{CT}$. By Property 2, \mathcal{CT} is a trace beginning at $\mathcal{C}_0(Q_0, \text{program}_0)$. Let us denote $\mathcal{CTS}_n = \{\mathcal{CT}_1, \dots, \mathcal{CT}_m\}$. We can then apply Lemma 6.29 to \mathfrak{C}^{cs} .

- Either there exist n' configurations $\mathfrak{C}_1^{\text{cs}}, \dots, \mathfrak{C}_{n'}^{\text{cs}}$, n' traces $\mathfrak{C}^{\text{cs}} \rightsquigarrow_{p_1}^+ \mathfrak{C}_1^{\text{cs}}, \dots, \mathfrak{C}^{\text{cs}} \rightsquigarrow_{p_{n'}}^+ \mathfrak{C}_{n'}^{\text{cs}}$ such that none of these traces is a prefix of another, $\sum_{i \leq n} p_i = 1$, and for each trace $\mathcal{CT}_j, j \leq m$, there exist n' pairwise disjoint trace sets $\mathcal{CTS}_{j,1}, \dots, \mathcal{CTS}_{j,n'}$ such that all traces in these sets are extensions of \mathcal{CT}_j , none of these traces is a prefix of another, and for each trace $\mathcal{CT} \in \mathcal{CTS}_{j,i}, \mathfrak{C}_i^{\text{cs}'} \equiv \mathcal{CT}$ and $\Pr[\mathcal{CTS}_{j,i}] = p_i \cdot \Pr[\mathcal{CT}_j]$. Let us

denote $\mathcal{CTS}'_i \stackrel{\text{def}}{=} \bigcup_{j \leq m} \mathcal{CTS}_{j,i}$. Let us also denote $\mathfrak{C}_i^{\text{cs}'}$ the extension of the trace $\mathfrak{C}_i^{\text{cs}}$ until $\mathfrak{C}_i^{\text{cs}'}$, for $i \leq n'$. There is at least one new reduction step, so there are more reduction steps in $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_{n-1}^{\text{cs}}, \mathfrak{E}\mathfrak{T}_1^{\text{cs}'}, \dots, \mathfrak{E}\mathfrak{T}_{n'}^{\text{cs}'}$ than in $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_n^{\text{cs}}$. Let us prove that $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_{n-1}^{\text{cs}}, \mathfrak{E}\mathfrak{T}_1^{\text{cs}'}, \dots, \mathfrak{E}\mathfrak{T}_{n'}^{\text{cs}'} \equiv_{\text{t}} \mathcal{CTS}_1, \dots, \mathcal{CTS}_{n-1}, \mathcal{CTS}'_1, \dots, \mathcal{CTS}'_{n'}$. All traces $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_{n-1}^{\text{cs}}, \mathfrak{E}\mathfrak{T}_1^{\text{cs}'}, \dots, \mathfrak{E}\mathfrak{T}_{n'}^{\text{cs}'}$ start at $\mathfrak{C}_0(Q_0, \text{program}_0)$ and by Lemma G.2, none of these traces is a prefix of another of these traces, so Property 1 is verified. Similarly, by applying Lemma G.2 to each trace \mathcal{CT}_j for $j \leq m$, Property 2 is verified. By Property 3 on the initial traces, $\forall i \leq n, \Pr[\mathfrak{C}_i^{\text{cs}}] = \Pr[\mathcal{CTS}_i]$. We have that $\Pr[\mathcal{CTS}'_i] = \sum_{j \leq m} \Pr[\mathcal{CTS}_{j,i}]$ because all the sets $\mathcal{CTS}_{j,i}$ are disjoint. So,

$$\Pr[\mathcal{CTS}'_i] = \sum_{j \leq m} p_i \cdot \Pr[\mathcal{CT}_j] = p_i \cdot \Pr[\mathcal{CTS}_n], \text{ so}$$

$$\Pr[\mathfrak{C}_i^{\text{cs}'}] = p_i \cdot \Pr[\mathfrak{C}_n^{\text{cs}}] = \Pr[\mathcal{CTS}'_i],$$

Property 3 is verified. By Property 4 on the initial traces, $\sum_{i \leq n} \Pr[\mathfrak{C}_i^{\text{cs}}] = 1$. We have that

$$\sum_{i \leq n'} \Pr[\mathfrak{C}_i^{\text{cs}'}] = \sum_{i \leq n'} p_i \cdot \Pr[\mathfrak{C}_n^{\text{cs}}] = \Pr[\mathfrak{C}_n^{\text{cs}}], \text{ so}$$

$$\sum_{i \leq n-1} \Pr[\mathfrak{C}_i^{\text{cs}}] + \sum_{i \leq n'} \Pr[\mathfrak{C}_i^{\text{cs}'}] = \sum_{i \leq n} \Pr[\mathfrak{C}_i^{\text{cs}}] = 1.$$

So Property 4 is verified. We inherit Property 5 for the $n-1$ first elements. For all $i \leq n'$, for all traces $\mathcal{CT} \in \mathcal{CTS}'_i$, we have $\mathfrak{C}_i^{\text{cs}'} \equiv \mathcal{CT}$, and $\mathfrak{C}_i^{\text{cs}'}$ is the last configuration of $\mathfrak{E}\mathfrak{T}_i^{\text{cs}'}$. So Property 5(b) is verified for all the new elements. Hence Property 5 is verified. Therefore, $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_{n-1}^{\text{cs}}, \mathfrak{E}\mathfrak{T}_1^{\text{cs}'}, \dots, \mathfrak{E}\mathfrak{T}_{n'}^{\text{cs}'} \equiv_{\text{t}} \mathcal{CTS}_1, \dots, \mathcal{CTS}_{n-1}, \mathcal{CTS}'_1, \dots, \mathcal{CTS}'_{n'}$.

- Otherwise, each trace $\mathcal{CT} \in \mathcal{CTS}_n$ is complete, $\mathfrak{C}_n^{\text{cs}} \rightarrow^* \mathfrak{C}_1^{\text{cs}}$, $\mathfrak{C}_1^{\text{cs}}$ cannot reduce, and the event list \mathcal{E} of $\mathfrak{C}_1^{\text{cs}}$ and the event list *events* of the last configuration of \mathcal{CT} satisfy *events* = $\mathbb{G}_{\text{ev}}(\mathcal{E})$. Let $\mathfrak{E}\mathfrak{T}_n^{\text{cs}'}$ be the extension of the trace $\mathfrak{E}\mathfrak{T}_n^{\text{cs}}$ until $\mathfrak{C}_1^{\text{cs}}$. The trace $\mathfrak{E}\mathfrak{T}_n^{\text{cs}'}$ contains more steps than $\mathfrak{E}\mathfrak{T}_n^{\text{cs}}$, so there are more reduction steps in $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_{n-1}^{\text{cs}}, \mathfrak{E}\mathfrak{T}_n^{\text{cs}'}$ than in $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_n^{\text{cs}}$. Let us prove that $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_{n-1}^{\text{cs}}, \mathfrak{E}\mathfrak{T}_n^{\text{cs}'} \equiv_{\text{t}} \mathcal{CTS}_1, \dots, \mathcal{CTS}_n$. By Lemma G.2, no traces in $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_{n-1}^{\text{cs}}, \mathfrak{E}\mathfrak{T}_n^{\text{cs}'}$ are prefixes of one another, so Property 1 is verified. Property 2 is inherited. We have that $\Pr[\mathfrak{C}_n^{\text{cs}'}] = \Pr[\mathfrak{C}_n^{\text{cs}}]$, so Properties 3 and 4 are verified. The trace $\mathfrak{E}\mathfrak{T}_n^{\text{cs}'}$ is complete, every trace $\mathcal{CT} \in \mathcal{CTS}_n$ is complete, and the event list *events* of the last configuration of traces in \mathcal{CTS}_n and the event list \mathcal{E} of $\mathfrak{C}_1^{\text{cs}}$ verify *events* = $\mathbb{G}_{\text{ev}}(\mathcal{E})$, so Property 5(a) holds for the last elements. Other elements inherit Property 5, so Property 5 holds. Therefore, $\mathfrak{E}\mathfrak{T}_1^{\text{cs}}, \dots, \mathfrak{E}\mathfrak{T}_{n-1}^{\text{cs}}, \mathfrak{E}\mathfrak{T}_n^{\text{cs}'} \equiv_{\text{t}} \mathcal{CTS}_1, \dots, \mathcal{CTS}_n$. \square

Proof (of Proposition 6.30) By Lemma 6.28, we have a trace $\mathfrak{C}\mathfrak{T}_0 = \mathfrak{C}_0(Q_0, program_0) \rightsquigarrow^* \mathfrak{C}^{cs}$ where $\mathfrak{C}^{cs} \equiv \mathcal{CT}_0$ and $\mathcal{CT}_0 = \mathcal{C}_0(Q_0, program_0)$. We prove easily that $\mathfrak{C}\mathfrak{T}_0 \equiv_t \{\mathcal{CT}_0\}$. The number of steps in complete traces from configuration $\mathfrak{C}_0(Q_0, program_0)$ is finite. Let us consider traces such that $\mathfrak{C}\mathfrak{T}_1^{cs}, \dots, \mathfrak{C}\mathfrak{T}_n^{cs} \equiv_t \mathcal{CTS}_1, \dots, \mathcal{CTS}_n$ with the maximum number of reduction steps. By Lemma G.3, the traces $\mathfrak{C}\mathfrak{T}_1^{cs}, \dots, \mathfrak{C}\mathfrak{T}_n^{cs}$ are complete. (Otherwise, we could extend them.) Since the sum of their probabilities is 1, these are all complete traces starting at $\mathfrak{C}_0(Q_0, program_0)$. The proposition follows. \square