

CryptoVerif: A Computationally-Sound Security Protocol Verifier

Bruno Blanchet
Inria Paris, France
`Bruno.Blanchet@inria.fr`

November 20, 2017

Abstract

This document presents the security protocol verifier CryptoVerif. In contrast to most previous provers, CryptoVerif does not rely on the Dolev-Yao model, but on the computational model. It can verify secrecy and correspondence properties (which include authentication). It produces proofs presented as sequences of games, like those manually written by cryptographers; these games are formalized in a probabilistic polynomial-time process calculus. CryptoVerif provides a generic method for specifying security properties of the cryptographic primitives. It produces proofs valid for a any number of sessions of the protocol, and provides an upper bound on the probability of success of an attack against the protocol as a function of the probability of breaking each primitive and of the number of sessions. It can work automatically, or the user can guide it with manual proof indications.

1 Introduction

There exist two main approaches for analyzing security protocols. In the computational model, messages are bitstrings, and the adversary is a probabilistic polynomial-time Turing machine. This model is close to the real execution of protocols, but the proofs are usually manual and informal. In contrast, in the symbolic, Dolev-Yao model, cryptographic primitives are considered as perfect blackboxes, modeled by function symbols in an algebra of terms, possibly with equations. The adversary can compute using these blackboxes. This abstract model makes it easier to build automatic verification tools, but the security proofs are in general not sound with respect to the computational model.

In contrast to most protocol verifiers, CryptoVerif works directly in the computational model, without considering the Dolev-Yao model. It produces proofs valid for any number of sessions of the protocol, in the presence of an active adversary. These proofs are presented as sequences of games, as used by cryptographers [3,9,10]: the initial game represents the protocol to prove; the goal is to bound the probability of breaking a certain security property in this game; intermediate games are obtained each from the previous one by transformations such that the difference of probability between consecutive games can easily be bounded; the final game is such that the desired probability is obviously bounded from the form of the game. (In general, it is simply 0 in that game.) The desired probability can then be easily bounded in the initial game.

We represent games in a process calculus. This calculus is inspired by the pi-calculus and by the calculi of [8] and of [7]. In this calculus, messages are bitstrings, and cryptographic primitives are functions from bitstrings to bitstrings. The calculus has a probabilistic semantics. The main tool for specifying security properties is indistinguishability: Q is indistinguishable from Q' up to probability p , $Q \approx_p Q'$, when the adversary has probability at most p of distinguishing Q from Q' . With respect to previous calculi mentioned above, our calculus introduces an important

novelty which is key for the automatic proof of security protocols: the values of all variables during the execution of a process are stored in arrays. For instance, $x[i]$ is the value of x in the i -th copy of the process that defines x . Arrays replace lists often used by cryptographers in their manual proofs of protocols. For example, consider the standard security assumption on a message authentication code (MAC). Informally, this definition says that the adversary has a negligible probability of forging a MAC, that is, that all correct MACs have been computed by calling the MAC oracle (*i.e.*, function). So, in cryptographic proofs, one defines a list containing the arguments of calls to the MAC oracle, and when checking a MAC of a message m , one can additionally check that m is in this list, with a negligible change in probability. In our calculus, the arguments of the MAC oracle are stored in arrays, and we perform a lookup in these arrays in order to find the message m . Arrays make it easier to automate proofs since they are always present in the calculus: one does not need to add explicit instructions to insert values in them, in contrast to the lists used in manual proofs. Therefore, many trivially sound but difficult to automate syntactic transformations disappear. Furthermore, relations between elements of arrays can easily be expressed by equalities, possibly involving computations on array indices.

CryptoVerif relies on a collection of game transformations, in order to transform the initial protocol into a game on which the desired security property is obvious. The most important kind of transformations exploits the security assumptions on cryptographic primitives in order to obtain a simpler game. These transformations can be specified in a generic way: we represent the security assumption of each cryptographic primitive by an observational equivalence $L \approx_p R$, where the processes L and R encode oracles: they input the arguments of the oracle and send its result back. Then, the prover can automatically transform a process Q that calls the oracles of L (more precisely, contains as subterms terms that perform the same computations as oracles of L) into a process Q' that calls the oracles of R instead. We have used this technique to specify several variants of shared-key and public-key encryption, signature, message authentication codes, hash functions, Diffie-Hellman key agreement, simply by giving the appropriate equivalence $L \approx_p R$ to the prover. Other game transformations are syntactic transformations, used in order to be able to apply an assumption on a cryptographic primitive, or to simplify the game obtained after applying such an assumption.

In order to prove protocols, these game transformations are organized using a proof strategy based on advice: when a transformation fails, it suggests other transformations that should be applied before, in order to enable the desired transformation. Thanks to this strategy, protocols can often be proved in a fully automatic way. For delicate cases, CryptoVerif has an interactive mode, in which the user can manually specify the transformations to apply. It is usually sufficient to specify a few transformations coming from the security assumptions of primitives, by indicating the concerned cryptographic primitive and the concerned secret key if any; the prover infers the intermediate syntactic transformations by the advice strategy. This mode is helpful for proving some public-key protocols, in which several security assumptions on primitives can be applied, but only one leads to a proof of the protocol. Importantly, CryptoVerif is always sound: whatever indications the user gives, when the prover shows a security property of the protocol, the property indeed holds assuming the given assumptions on the cryptographic primitives.

CryptoVerif has been implemented in Ocaml (29800 lines of code for version 1.12 of CryptoVerif) and is available at <http://cryptoverif.inria.fr/>.

Outline Currently, this document only presents the process calculus that CryptoVerif uses for representing games, with its syntax, type system and formal semantics, in the next section. We plan to add information on the game transformations, the proof strategy, and the algorithms for proving security properties in the future.

Notations We recall the following standard notations. We denote by $\{M_1/x_1, \dots, M_m/x_m\}$ the substitution that replaces x_j with M_j for each $j \leq m$. The cardinal of a set or multiset S is denoted by $|S|$. We use \uplus for multiset union. When S is a multiset, $S(x)$ is the number of elements of S equal to x . If S is a finite set, $x \stackrel{R}{\leftarrow} S$ chooses a random element uniformly in S and assigns it to x . If \mathcal{A} is a probabilistic algorithm, $x \leftarrow \mathcal{A}(x_1, \dots, x_m)$ denotes the experiment of choosing random coins r and assigning to x the result of running $\mathcal{A}(x_1, \dots, x_m)$ with coins r . Otherwise, $x \leftarrow M$ is a simple assignment statement. If D is a discrete probability distribution, we denote by $D(a)$ the probability that $X = a$, $\Pr[X = a]$, where X is a random variable with probability distribution D .

2 A Calculus for Cryptographic Games

2.1 Syntax and Informal Semantics

CryptoVerif represents games in the syntax of Figure 1. This calculus assumes a countable set of channel names, denoted by c . It uses parameters, denoted by n , which are integers that bound the number of executions of processes.

It also uses types, denoted by T , which are sets of values. A type is *fixed* when it is the set of all bitstrings of a certain length; a type is *bounded* when it is a finite set. Particular types are predefined: $bool = \{\text{true}, \text{false}\}$, where false is 0 and true is 1; $bitstring$ is the set of all bitstrings; $bitstring_{\perp} = bitstring \cup \{\perp\}$ where \perp is a special symbol; $[1, n]$ where n is a parameter. (We consider integers as bitstrings without leading zeroes.)

The calculus also uses function symbols f . Each function symbol comes with a type declaration $f : T_1 \times \dots \times T_m \rightarrow T$, and represents an efficiently computable, deterministic function that maps each tuple in $T_1 \times \dots \times T_m$ to an element of T . Particular functions are predefined, and some of them use the infix notation: $M = N$ for the equality test, $M \neq N$ for the inequality test (both taking two values of the same type T and returning a value of type $bool$), $M \vee N$ for the boolean or, $M \wedge N$ for the boolean and, $\neg M$ for the boolean negation (taking and returning values of type $bool$), tuples (M_1, \dots, M_m) (taking values of any types and returning values of type $bitstring$; tuples are assumed to provide unambiguous concatenation, with tags for the types of M_1, \dots, M_m so that tuples of different types are always different).

In this calculus, terms represent computations on bitstrings. The replication index i is an integer which serves in distinguishing different copies of a replicated process $!^{i \leq n}$. (Replication indices are typically used as array indices.) The variable access $x[M_1, \dots, M_m]$ returns the content of the cell of indices M_1, \dots, M_m of the m -dimensional array variable x . We use x, y, z, u as variable names. The function application $f(M_1, \dots, M_m)$ returns the result of applying function f to M_1, \dots, M_m . Terms contain additional constructs which are very similar to those also included in output processes and explained below. These constructs conclude by evaluating a term, instead of executing a process. The construct `event_abort` e executes event e (without argument) and aborts the game; it is in fact intended for use in the definition of cryptographic primitives.

The calculus distinguishes two kinds of processes: input processes Q are ready to receive a message on a channel; output processes P output a message on a channel after executing some internal computations. The input process 0 does nothing; $Q \mid Q'$ is the parallel composition of Q and Q' ; $!^{i \leq n} Q$ represents n copies of Q in parallel, each with a different value of $i \in [1, n]$; `newChannel` $c; Q$ creates a new private channel c and executes Q ; this construct is useful in proofs, but does not occur in games manipulated by CryptoVerif. The semantics of the input $c[M_1, \dots, M_i](p); P$ will be explained below together with the semantics of the output.

The output process `new` $x[\tilde{i}] : T; P$ chooses a new random value in T , stores it in $x[\tilde{i}]$, and executes P . The abbreviation \tilde{i} stands for a sequence of replication indices i_1, \dots, i_m . The random value is chosen according to the default distribution D_T for type T , which is determined

$M, N ::=$	i $x[M_1, \dots, M_m]$ $f(M_1, \dots, M_m)$ $\text{new } x[\tilde{i}] : T; N$ $\text{let } p = M \text{ in } N \text{ else } N'$ $\text{let } x[\tilde{i}] : T = M \text{ in } N$ $\text{if defined}(M_1, \dots, M_l) \wedge M \text{ then } N \text{ else } N'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M'_j \text{ then } N'_j) \text{ else } N'$ $\text{insert } Tbl(M_1, \dots, M_l); N$ $\text{get } Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } N \text{ else } N'$ $\text{event } e(M_1, \dots, M_l); N$ $\text{event_abort } e$	terms replication index variable access function application random number assignment (pattern-matching) assignment conditional conditional array lookup insert in table get from table event event e and abort
$p ::=$	$x[\tilde{i}] : T$ $f(p_1, \dots, p_m)$ $=M$	pattern variable function application comparison with a term
$Q ::=$	0 $Q \mid Q'$ $!^{i \leq n} Q$ $\text{newChannel } c; Q$ $c[M_1, \dots, M_l](p); P$	input process nil parallel composition replication n times channel restriction input
$P ::=$	$\overline{c[M_1, \dots, M_l]\langle N \rangle}; Q$ $\text{new } x[\tilde{i}] : T; P$ $\text{let } p = M \text{ in } P \text{ else } P'$ $\text{if defined}(M_1, \dots, M_l) \wedge M \text{ then } P \text{ else } P'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P'_j) \text{ else } P$ $\text{insert } Tbl(M_1, \dots, M_l); P$ $\text{get } Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } P \text{ else } P'$ $\text{event } e(M_1, \dots, M_l); P$ $\text{event_abort } e$ yield	output process output random number assignment conditional conditional array lookup insert in table get from table event event e and abort end

Figure 1: Syntax of the process calculus

as follows:

- When the type T is declared with option *nonuniform*, the default probability distribution D_T for type T may be non-uniform. It is left unspecified.
- Otherwise, if T is *fixed*, T consists of all bitstrings of a certain length, and the default distribution is the uniform distribution. The probability of each element of T is $1/|T|$.
- If T is *bounded* but not *fixed*, T is finite, and the default distribution is an approximately uniform distribution, such that its distance to the uniform distribution is at most ϵ_T . The distance between two probability distributions D_1 and D_2 for type T is

$$d(D_1, D_2) = \sum_{a \in T} |D_1(a) - D_2(a)|$$

Indeed, probabilistic Turing machines that run in bounded time cannot choose random elements exactly uniformly in sets whose cardinal is not a power of 2.

For example, a possible algorithm to obtain a random integer in $[0, m - 1]$ is to choose a random integer x' uniformly among $[0, 2^k - 1]$ for a certain k large enough and return $x' \bmod m$. By euclidian division, we have $2^k = qm + r$ with $r \in [0, m - 1]$. With this algorithm

$$D(a) = \begin{cases} \frac{q+1}{2^k} & \text{if } a \in [0, r - 1] \\ \frac{q}{2^k} & \text{if } a \in [r, m - 1] \end{cases}$$

so

$$\left| D(a) - \frac{1}{m} \right| = \begin{cases} \frac{q+1}{2^k} - \frac{1}{m} & \text{if } a \in [0, r - 1] \\ \frac{1}{m} - \frac{q}{2^k} & \text{if } a \in [r, m - 1] \end{cases}$$

Therefore

$$\begin{aligned} d(D_T, \text{uniform}) &= \sum_{a \in T} \left| D(a) - \frac{1}{m} \right| = r \left(\frac{q+1}{2^k} - \frac{1}{m} \right) - (m-r) \left(\frac{1}{m} - \frac{q}{2^k} \right) \\ &= \frac{2r(m-r)}{m \cdot 2^k} \leq \frac{m}{2^k} \end{aligned}$$

so we can take $\epsilon_T = \frac{m}{2^k}$. A given precision of $\epsilon_T = \frac{1}{2^{k'}}$ can be obtained by choosing $k = (k' + \text{number of bits of } m)$ random bits.

By default, CryptoVerif does not display ϵ_T in probability formulas, to make them more readable.

When T is not declared with any of the options *nonuniform*, *fixed*, or *bounded*, CryptoVerif rejects the construct $\text{new } x[\tilde{i}] : T; P$. Function symbols represent deterministic functions, so all random numbers must be chosen by $\text{new } x[\tilde{i}] : T$. Deterministic functions make automatic syntactic manipulations easier: we can duplicate a term without changing its value.

The process $\text{let } x[\tilde{i}] : T = M \text{ in } P$ stores the value of M (which must be in T) in $x[\tilde{i}]$ and executes P . Furthermore, we say that a function $f : T_1 \times \dots \times T_m \rightarrow T$ is *efficiently injective* when it is injective and its inverses are efficiently computable, that is, there exist functions $f_j^{-1} : T \rightarrow T_j$ ($1 \leq j \leq m$) such that $f_j^{-1}(f(x_1, \dots, x_m)) = x_j$ and f_j^{-1} is efficiently computable. When f is efficiently injective, we define a pattern matching construct $\text{let } f(x_1, \dots, x_m) = M \text{ in } P \text{ else } Q$ as an abbreviation for $\text{let } y : T = M \text{ in let } x'_1 : T_1 = f_1^{-1}(y) \text{ in } \dots \text{let } x'_m : T_m = f_m^{-1}(y) \text{ in if } f(x'_1, \dots, x'_m) = y \text{ then (let } x_1 : T_1 = x'_1 \text{ in } \dots \text{let } x_m : T_m = x'_m \text{ in } P) \text{ else } Q$ where y, x'_1, \dots, x'_m are fresh variables. (The variables x'_1, \dots, x'_m are introduced to make sure that none of the variables x_1, \dots, x_m is defined when the pattern-matching fails.) We naturally generalize this construct to $\text{let } p = M \text{ in } P \text{ else } Q$ where p is built from variables, efficiently

injective functions, and equality tests. When p is simply a variable, the pattern-matching always succeeds, so the `else` branch of the assignment is never executed and can be omitted.

The process event $e(M_1, \dots, M_l); P$ executes the event $e(M_1, \dots, M_l)$, then runs P . This event records that a certain program point has been reached with certain values of M_1, \dots, M_l , but otherwise does not affect the execution of the process.

The process event `abort` e executes event e (without argument) and aborts the game.

Next, we explain the process `find`[*unique?*] $(\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j}$ **suchthat** $\text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j$ **then** P_j) **else** P . The order and array indices on tuples are taken component-wise, so for instance, $u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j}$ can be further abbreviated $\tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j$. A simple example is the following: `find` $u = i \leq n$ **suchthat** $\text{defined}(x[i]) \wedge x[i] = a$ **then** P' **else** P tries to find an index i such that $x[i]$ is defined and $x[i] = a$, and when such an i is found, it stores it in u and executes P' with that value of u ; otherwise, it executes P . In other words, this `find` construct looks for the value a in the array x , and when a is found, it stores in u an index such that $x[u] = a$. Therefore, the `find` construct allows us to access arrays, which is key for our purpose. More generally, `find` $u_1[\tilde{i}] = i_1 \leq n_1, \dots, u_m[\tilde{i}] = i_m \leq n_m$ **suchthat** $\text{defined}(M_1, \dots, M_l) \wedge M$ **then** P' **else** P tries to find values of i_1, \dots, i_m for which M_1, \dots, M_l are defined and M is true. In case of success, it stores the obtained values in $u_1[\tilde{i}], \dots, u_m[\tilde{i}]$ and executes P' . In case of failure, it executes P . This is further generalized to m branches: `find` $(\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j}$ **suchthat** $\text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j$ **then** P_j) **else** P tries to find a branch j in $[1, m]$ such that there are values of i_{j1}, \dots, i_{jm_j} for which M_{j1}, \dots, M_{jl_j} are defined and M_j is true. In case of success, it stores them in $u_{j1}[\tilde{i}], \dots, u_{jm_j}[\tilde{i}]$ and executes P_j . In case of failure for all branches, it executes P . More formally, it evaluates the conditions $\text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j$ for each j and each value of i_{j1}, \dots, i_{jm_j} in $[1, n_{j1}] \times \dots \times [1, n_{jm_j}]$. If none of these conditions is true, it executes P . Otherwise, it chooses randomly one j and one value of i_{j1}, \dots, i_{jm_j} such that the corresponding condition is true, according to the distribution $D_{\text{find}}(S)$ where S is the set of possible solutions $j, i_{j1}, \dots, i_{jm_j}$, stores it in $u_{j1}[\tilde{i}], \dots, u_{jm_j}[\tilde{i}]$, and executes P_j . The distribution $D_{\text{find}}(S)$ is almost uniform: formally, the distance between $D_{\text{find}}(S)$ and the uniform distribution is at most ϵ_{find} , that is, $d(D_{\text{find}}(S), \text{uniform}) \leq \epsilon_{\text{find}}$. By default, CryptoVerif does not display ϵ_{find} in probability formulas, to make them more readable. We cannot take the first element found because the game transformations made by CryptoVerif may reorder the elements. For these transformations to preserve the behavior of the game, the distribution of the chosen element must be invariant by reordering, up to a small probability ϵ_{find} . In this definition, the variables i_{j1}, \dots, i_{jm_j} are considered as replication indices, while $u_{j1}[\tilde{i}], \dots, u_{jm_j}[\tilde{i}]$ are considered as array variables. The indication [*unique?*] stands for either [`uniquee`] or empty. The empty case has just been explained. When the `find` is marked [`uniquee`] and there are several solutions (that is, $k > 1$), we execute the event e and abort the game. When there is zero or one solution (that is, $k = 0$ or $k = 1$), the `find` is executed as when [*unique?*] is empty. This semantics allows us to perform game transformations that require the `find` to have a single solution.

The conditional `if` $\text{defined}(M_1, \dots, M_l) \wedge M$ **then** P **else** P' executes P if M_1, \dots, M_l are defined and M evaluates to true. Otherwise, it executes P' . This conditional is equivalent to `find` **suchthat** $\text{defined}(M_1, \dots, M_l) \wedge M$ **then** P **else** P' . The conjunct $\text{defined}(M_1, \dots, M_l)$ can be omitted when $l = 0$ and M can be omitted when it is true.

The constructs `insert` and `get` handle tables, used for instance to store the keys of the protocol participants. A table can be represented as a list of tuples; `insert` $Tbl(M_1, \dots, M_l); P$ inserts the element M_1, \dots, M_l in the table Tbl ; `get` $Tbl(x_1 : T_1, \dots, x_l : T_l)$ **suchthat** M **in** P **else** P' tries to retrieve an element (x_1, \dots, x_l) in the table Tbl such that M is true. When such an element is found, it executes P with x_1, \dots, x_l bound to that element. (When several such elements are found, one of them is chosen randomly according to distribution $D_{\text{get}}(\{1, \dots, |L|\})$ where L is the list of suitable elements, with $d(D_{\text{get}}(\{1, \dots, |L|\}), \text{uniform}) \leq \epsilon_{\text{find}}$.) When no such element is found, P' is executed. We can generalize this construct to patterns instead of

variables similarly to the `let` case. CryptoVerif internally translates the `insert` and `get` constructs into `find`.

Let us explain the output $\overline{c[M_1, \dots, M_l]}(N); Q$. A channel $c[M_1, \dots, M_l]$ consists of both a channel name c and a tuple of terms M_1, \dots, M_l . Channel names c can be declared private by `newChannel` c ; the adversary can never have access to channel $c[M_1, \dots, M_l]$ when c is private. (This is useful in the proofs, although all channels of protocols are often public.) Terms M_1, \dots, M_l are intuitively analogous to IP addresses and ports, which are numbers that the adversary may guess. A semantic configuration always consists of a single output process (the process currently being executed) and several input processes. When the output process executes $\overline{c[M_1, \dots, M_l]}(N); Q$, one looks for an input on channel $c[M'_1, \dots, M'_l]$, where M'_1, \dots, M'_l evaluate to the same bitstrings as M_1, \dots, M_l , in the available input processes. If no such input process is found, the process blocks. Otherwise, one such input process $c[M'_1, \dots, M'_l](x[\tilde{i}] : T); P$ is chosen randomly according to the probability distribution $D_{\text{in}}(S)$ where S is the multiset of suitable input processes. The communication is then executed: the output message N is evaluated and stored in $x[\tilde{i}]$ if it is in T (otherwise the process blocks). Finally, the output process P that follows the input is executed. The input process Q that follows the output is stored in the available input processes for future execution. The input construct can be generalized to patterns instead of variables similarly to the `let` case; when pattern-matching fails, the input process executes `yield`. The syntax requires an output to be followed by an input process, as in [7]. If one needs to output several messages consecutively, one can simply insert fictitious inputs between the outputs. The adversary can then schedule the outputs by sending messages to these inputs.

Using different channels for each input and output allows the adversary to control the network. For instance, we may write $!^{i \leq n} c[i](x[i] : T) \dots \overline{c'[i]}(M) \dots$. The adversary can then decide which copy of the replicated process receives its message, simply by sending it on $c[i]$ for the appropriate value of i .

The `yield` construct is an abbreviation for $\overline{\text{yield}}(\langle \rangle)$. By performing an output, this construct returns control to the adversary, which is going to receive the message. An `else` branch of `find`, `if`, `get`, or `let` may be omitted when it is `else yield`. (Note that “else 0” would not be syntactically correct.) Similarly, `;``yield` may be omitted after `event`, `new`, or `insert` and `in yield` may be omitted after `let`. A trailing 0 after an output may be omitted.

The *current replication indices* at a certain program point in a process are the replication indices i_1, \dots, i_m bound by replications and `find` above that program point. The replication $!^{i \leq n} Q$ binds the replication index i in Q . The `find` construct $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } \dots) \text{ else } \dots$ binds the replication indices i_{j1}, \dots, i_{jm_j} in $\text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j$. We often abbreviate $x[i_1, \dots, i_m]$ by x when i_1, \dots, i_m are the current replication indices, but it should be kept in mind that this is only an abbreviation. Variables defined under a replication must be arrays: for example $!^{i_1 \leq n_1} \dots !^{i_m \leq n_m} \text{let } x[i_1, \dots, i_m] : T = M \text{ in } \dots$. More formally, we require the following invariant:

Invariant 1 (Single definition) The process Q_0 satisfies Invariant 1 if and only if

1. in every definition of $x[i_1, \dots, i_m]$ in Q_0 , the indices i_1, \dots, i_m of x are the current replication indices at that definition, and
2. two different definitions of the same variable x in Q_0 are in different branches of a `find` (or `if` or `let`) or `get`.

Invariant 1 guarantees that each variable is assigned at most once for each value of its indices. (Indeed, item 2 shows that only one definition of each variable can be executed for given indices in each trace.) A definition of $x[\tilde{i}]$ can be `new` $x[\tilde{i}] : T$, a `let`, `get`, or input that contains the pattern $x[\tilde{i}] : T$, or `find` $\dots x[\tilde{i}] = i \leq n \dots$

Invariant 2 (Defined variables) The process Q_0 satisfies Invariant 2 if and only if every occurrence of a variable access $x[M_1, \dots, M_m]$ in Q_0 is either

- syntactically under the definition of $x[M_1, \dots, M_m]$ (in which case M_1, \dots, M_m are in fact the current replication indices at the definition of x);
- or in a defined condition in a find process or term;
- or in M'_j in a process or term of the form $\text{find } (\bigoplus_{j=1}^{m''} \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(M'_{j_1}, \dots, M'_{j_{l_j}}) \wedge M'_j \text{ then } P_j) \text{ else } P$ where for some $k \leq l_j$, $x[M_1, \dots, M_m]$ is a subterm of M'_{j_k} .
- or in P_j in a process or term of the form $\text{find } (\bigoplus_{j=1}^{m''} \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(M'_{j_1}, \dots, M'_{j_{l_j}}) \wedge M'_j \text{ then } P_j) \text{ else } P$ where for some $k \leq l_j$, there is a subterm N of M'_{j_k} such that $N\{\tilde{u}_j[\tilde{i}]/\tilde{i}_j\} = x[M_1, \dots, M_m]$.

Invariant 2 guarantees that variables can be accessed only when they have been initialized. It checks that the definition of the variable access is either in scope (first item) or checked by a find (last two items). We recall that if is a particular case of find. The scope of variable definitions is defined as follows: $x[\tilde{i}]$ is syntactically under its definition when it is

- inside P in $\text{new } x[\tilde{i}] : T; P$;
- inside N in $\text{new } x[\tilde{i}] : T; N$;
- inside P in $\text{let } p = M \text{ in } P \text{ else } P'$ when $x[\tilde{i}] : T$ is bound in the pattern p ;
- inside N in $\text{let } p = M \text{ in } N \text{ else } N'$ when $x[\tilde{i}] : T$ is bound in the pattern p ;
- inside N in $\text{let } x[\tilde{i}] : T = M \text{ in } N$;
- inside P_j in $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j \text{ then } P_j) \text{ else } P$ when x is u_{jk} for some $k \leq m_j$;
- inside N_j in $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j \text{ then } N_j) \text{ else } N$ when x is u_{jk} for some $k \leq m_j$;
- inside M or P in $\text{get } Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } P \text{ else } P'$ when $x[\tilde{i}] : T$ is bound in one of the patterns p_1, \dots, p_l ;
- inside M or N in $\text{get } Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } N \text{ else } N'$ when $x[\tilde{i}] : T$ is bound in one of the patterns p_1, \dots, p_l ;
- inside P in $c[M_1, \dots, M_l](p); P$ when $x[\tilde{i}] : T$ is bound in the pattern p .

A variable access that does not correspond to the first item of Invariant 2 is called an *array access*. We furthermore require the following invariant.

Invariant 3 (Variables defined in find and get conditions) The process Q_0 satisfies Invariant 3 with public variables V if and only if the variables defined in conditions of find and the variables defined in patterns and in conditions of get have no array accesses and are not in the set of variables V .

These conditions are needed for variables of get, because they will be transformed into variables defined in conditions of find by the transformation of get into find.

Invariant 4 (Terms in find and get conditions) The process Q_0 satisfies Invariant 4 if and only if new, event, and insert do not occur in conditions of find and get.

Invariant 4 guarantees that evaluating the condition of a `find` or `get` does not change the state of the system.

Invariant 5 (Terms in input channels and defined conditions) The process Q_0 satisfies Invariant 5 if and only if only replication indices, variables, and function applications occur in input channels $c[M_1, \dots, M_l]$ and in conditions $\text{defined}(M_1, \dots, M_l)$ in `find` or `if`.

Other terms (`new`, `let`, `if`, `find`, `event`) are handled by expanding them into their corresponding processes. Invariant 5 is needed because terms in input channels cannot be expanded, as we need an output process to put the computations coming from expanded terms, and similarly terms in defined conditions cannot be expanded. By combining this invariant with Invariant 2, we see that the terms of all variable accesses $x[M_1, \dots, M_m]$ contain only replication indices, variables, and function applications.

The last 3 invariants did not appear in previous versions of the calculus because all terms were only replication indices, variables, and function applications.

Invariant 6 (Events) We distinguish three disjoint sets of events e , Shoup events, non-unique events, and other events. The process Q_0 satisfies Invariant 6 if and only if

- Shoup events occur only in processes of the form `event_abort` e in Q_0 , and
- non-unique events occur only in `find[uniquee]` in Q_0 .

The name “Shoup events” is used because these events are introduced when applying Shoup’s lemma [11]. The non-unique events are those triggered when a `find[uniquee]` actually has several solutions.

All these invariants are checked by the prover for the initial game and preserved by all game transformations.

We denote by $\text{var}(P)$ the set of variables that occur in P , $\text{vardef}(P)$ the set of variables defined in P ($\text{var}(P)$ may contain more variables than $\text{vardef}(P)$ in case some variables are read using `find` but never defined), and by $\text{fc}(P)$ the set of free channels of P . (We use similar notations for input processes.)

2.2 Example

Let us introduce two cryptographic primitives that we use below.

Definition 1 Let T_{mr} , T_{mk} , and T_{ms} be types that correspond intuitively to random seeds, keys, and message authentication codes, respectively; T_{mr} is a fixed-length type. A message authentication code scheme `MAC` [2] consists of three function symbols:

- `mkgen` : $T_{mr} \rightarrow T_{mk}$ is the key generation algorithm taking as argument a random bitstring and returning a key. (Usually, `mkgen` is a randomized algorithm; here, since we separate the choice of random numbers from computation, `mkgen` takes an additional argument representing the random coins.)
- `mac` : $\text{bitstring} \times T_{mk} \rightarrow T_{ms}$ is the MAC algorithm taking as arguments a message and a key, and returning the corresponding tag. (We assume here that `mac` is deterministic; we could easily encode a randomized `mac` by adding an additional argument as for `mkgen`.)
- `verify` : $\text{bitstring} \times T_{mk} \times T_{ms} \rightarrow \text{bool}$ is a verification algorithm such that $\text{verify}(m, k, t) = \text{true}$ if and only if t is a valid MAC of message m under key k . (Since `mac` is deterministic, $\text{verify}(m, k, t)$ is typically $\text{mac}(m, k) = t$.)

We have $\forall m \in \text{bitstring}, \forall r \in T_{mr}, \text{verify}(m, \text{mkgen}(r), \text{mac}(m, \text{mkgen}(r))) = \text{true}$.

The advantage of an adversary against unforgeability under chosen message attacks (UF-CMA) is

$$\text{Succ}_{\text{MAC}}^{\text{uf-cma}}(t, q_m, q_v, l) = \max_{\mathcal{A}} \Pr \left[\begin{array}{l} r \xleftarrow{R} T_{mr}; k \leftarrow \text{mkgen}(r); \\ (m, s) \leftarrow \mathcal{A}^{\text{mac}(\cdot, k), \text{verify}(\cdot, k, \cdot)} : \text{verify}(m, k, s) \\ \wedge m \text{ was never queried to the oracle } \text{mac}(\cdot, k) \end{array} \right]$$

where the adversary \mathcal{A} is any probabilistic Turing machine that runs in time at most t , calls $\text{mac}(\cdot, k)$ at most q_m times with messages of length at most l , and calls $\text{verify}(\cdot, k, \cdot)$ at most q_v times with messages of length at most l .

$\text{Succ}_{\text{MAC}}^{\text{uf-cma}}(t, q_m, q_v, l)$ is the probability that an adversary forges a MAC, that is, returns a pair (m, s) where s is a correct MAC for m , without having queried the MAC oracle $\text{mac}(\cdot, k)$ on m . Intuitively, when the MAC is secure, this probability is small: the adversary has little chance of forging a MAC. Hence, the MAC guarantees the integrity of the MACed message because one cannot compute the MAC without the secret key.

Two frameworks exist for expressing security properties. In the asymptotic framework, used in [4, 5], the length of keys is determined by a security parameter η , and a MAC is UF-CMA when $\text{Succ}_{\text{MAC}}^{\text{uf-cma}}(t, q_m, q_v, l)$ is a negligible function of η when t is polynomial in η . ($f(\eta)$ is *negligible* when for all polynomials q , there exists $\eta_o \in \mathbb{N}$ such that for all $\eta > \eta_o$, $f(\eta) \leq \frac{1}{q(\eta)}$.) The assumption that functions are efficiently computable means that they are computable in time polynomial in η and in the length of their arguments. The goal is to show that the probability of success of an attack against the protocol is negligible, assuming the parameters n are polynomial in η and the network messages are of length polynomial in η . In contrast, in the exact security framework, on which we focus in this report, one computes the probability of success of an attack against the protocol as a function of the probability of breaking the primitives such as $\text{Succ}_{\text{MAC}}^{\text{uf-cma}}(t, q_m, q_v, l)$, of the runtime of functions, of the parameters n , and of the length of messages, thus providing a more precise security result. Intuitively, the probability $\text{Succ}_{\text{MAC}}^{\text{uf-cma}}(t, q_m, q_v, l)$ is assumed to be small (otherwise, the computed probability of attack will be large), but no formal assumption on this probability is needed to establish the security theorem.

Definition 2 Let T_r and T'_r be fixed-length types representing random coins; let T_k and T_e be types for keys and ciphertexts respectively. A symmetric encryption scheme SE [2] consists of three function symbols:

- $\text{kgen} : T_r \rightarrow T_k$ is the key generation algorithm taking as argument random coins and returning a key,
- $\text{enc} : \text{bitstring} \times T_k \times T'_r \rightarrow T_e$ is the encryption algorithm taking as arguments the cleartext, the key, and random coins, and returning the ciphertext,
- $\text{dec} : T_e \times T_k \rightarrow \text{bitstring}_{\perp}$ is the decryption algorithm taking as arguments the ciphertext and the key, and returning either the cleartext when decryption succeeds or \perp when decryption fails,

such that $\forall m \in \text{bitstring}, \forall r \in T_r, \forall r' \in T'_r, \text{dec}(\text{enc}(m, \text{kgen}(r), r'), \text{kgen}(r)) = m$.

Let $LR(x, y, b) = x$ if $b = 0$ and $LR(x, y, b) = y$ if $b = 1$, defined only when x and y are bitstrings of the same length. The advantage of an adversary against indistinguishability under chosen plaintext attacks (IND-CPA) is

$$\text{Succ}_{\text{SE}}^{\text{ind-cpa}}(t, q_e, l) = \max_{\mathcal{A}} 2 \Pr \left[\begin{array}{l} b \xleftarrow{R} \{0, 1\}; r \xleftarrow{R} T_r; k \leftarrow \text{kgen}(r); \\ b' \leftarrow \mathcal{A}^{r' \xleftarrow{R} T'_r; \text{enc}(LR(\cdot, \cdot), k, r')} : b' = b \end{array} \right] - 1$$

where \mathcal{A} is any probabilistic Turing machine that runs in time at most t and calls $r' \xleftarrow{R} T'_r$; $\text{enc}(LR(\cdot, \cdot, b), k, r')$ at most q_e times on messages of length at most l .

Given two bitstrings a_0 and a_1 of the same length, the left-right encryption oracle $r' \xleftarrow{R} T'_r$; $\text{enc}(LR(\cdot, \cdot, b), k, r')$ returns $r' \xleftarrow{R} T'_r$; $\text{enc}(LR(a_0, a_1, b), k, r')$, that is, encrypts a_0 when $b = 0$ and a_1 when $b = 1$. $\text{Succ}_{\text{SE}}^{\text{ind-cpa}}(t, q_e, l)$ is the probability that the adversary distinguishes the encryption of the messages a_0 given as first arguments to the left-right encryption oracle from the encryption of the messages a_1 given as second arguments. Intuitively, when the encryption scheme is IND-CPA secure, this probability is small: the ciphertext gives almost no information what the cleartext is (one cannot determine whether it is a_0 or a_1 without having the secret key).

Example 1 Let us consider the following trivial protocol:

$$A \rightarrow B : e, \text{mac}(e, x_{mk}) \quad \text{where } e = \text{enc}(x'_k, x_k, x'_r) \\ \text{and } x'_r, x'_k \text{ are fresh random numbers}$$

A and B are assumed to share a key x_k for a symmetric encryption scheme and a key x_{mk} for a message authentication code. A creates a fresh key x'_k and sends it encrypted under x_k to B . A MAC is appended to the message, in order to guarantee integrity. In other words, the protocol sends the key x'_k encrypted using an encrypt-then-MAC scheme [2]. The goal of the protocol is that x'_k should be a secret key shared between A and B . This protocol can be modeled in our calculus by the following process Q_0 :

$$Q_0 = \text{start}(); \text{new } x_r : T_r; \text{let } x_k : T_k = \text{kgen}(x_r) \text{ in} \\ \quad \text{new } x_{mr} : T_{mr}; \text{let } x_{mk} : T_{mk} = \text{mkgen}(x_{mr}) \text{ in } \bar{c}\langle \rangle; (Q_A \mid Q_B) \\ Q_A = !^{i \leq n} c_A[i](); \text{new } x'_k : T_k; \text{new } x'_r : T'_r; \\ \quad \text{let } x_m : \text{bitstring} = \text{enc}(\text{k2b}(x'_k), x_k, x'_r) \text{ in } \overline{c_A[i]}\langle x_m, \text{mac}(x_m, x_{mk}) \rangle \\ Q_B = !^{i \leq n} c_B[i'](x'_m, x_{ma}); \text{if } \text{verify}(x'_m, x_{mk}, x_{ma}) \text{ then} \\ \quad \text{let } i_{\perp}(\text{k2b}(x''_k)) = \text{dec}(x'_m, x_k) \text{ in } \overline{c_B[i']}\langle \rangle$$

When Q_0 receives a message on channel start , it begins execution: it generates the keys x_k and x_{mk} by choosing random coins x_r and x_{mr} and applying the appropriate key generation algorithms. Then it yields control to the adversary, by outputting on channel c . After this output, n copies of processes for A and B are ready to be executed, when the adversary outputs on channels $c_A[i]$ or $c_B[i]$ respectively. In a session that runs as expected, the adversary first sends a message on $c_A[i]$. Then Q_A creates a fresh key x'_k (T_k is assumed to be a fixed-length type), encrypts it under x_k with random coins x'_r , computes the MAC under x_{mk} of the ciphertext, and sends the ciphertext and the MAC on $c_A[i]$. The function $\text{k2b} : T_k \rightarrow \text{bitstring}$ is the natural injection $\text{k2b}(x) = x$; it is needed only for type conversion. The adversary is then expected to forward this message on $c_B[i]$. When Q_B receives this message, it verifies the MAC, decrypts, and stores the obtained key in x''_k . (The function $i_{\perp} : \text{bitstring} \rightarrow \text{bitstring}_{\perp}$ is the natural injection; it is useful to check that decryption succeeded.) This key x''_k should be secret.

The adversary is responsible for forwarding messages from A to B . It can send messages in unexpected ways in order to mount an attack.

This very small example is sufficient to illustrate the main features of CryptoVerif.

2.3 Type System

We use a type system to check that bitstrings of the proper type are passed to each function and that array indices are used correctly.

To be able to type variable accesses used not under their definition (such accesses are guarded by a `find` construct), the type-checking algorithm proceeds in two passes. In the first pass, it builds a type environment \mathcal{E} , which maps variable names x to types $[1, n_1] \times \dots \times [1, n_m] \rightarrow T$, where the definition of $x[i_1, \dots, i_m]$ of type T occurs under replications or `find` that bind i_1, \dots, i_m with declaration $i_j \leq n_j$. (For instance, the definition of $x[i_1, \dots, i_m]$ occurs under $!^{i_1 \leq n_1}, \dots, !^{i_m \leq n_m}$ or it occurs in the condition of `find` $u_1 = i_1 \leq n_1, \dots, u_m = i_m \leq n_m$ under no replication. The type T is the one given in the definition of x in `new` $x[\tilde{i}] : T$ or in a pattern $x[\tilde{i}] : T$ in an assignment, an input, or a `get`. In the `find` construct, `find` $\dots x[\tilde{i}] = i \leq n$, the type T of x is $T = [1, n]$.) The tool checks that all definitions of the same variable x yield the same value of $\mathcal{E}(x)$, so that \mathcal{E} is properly defined.

In the second pass, the process is typechecked in the type environment \mathcal{E} using the rules of Figures 2 and 3. These figures defines four judgments:

- $\mathcal{E} \vdash M : T$ means that the term M has type T in environment \mathcal{E} .
- $\mathcal{E} \vdash p : T$ means that the pattern p has type T in environment \mathcal{E} .
- $\mathcal{E} \vdash P$ and $\mathcal{E} \vdash Q$ mean that the output process P and the input process Q are well-typed in environment \mathcal{E} , respectively.

In $x[M_1, \dots, M_m]$, M_1, \dots, M_m must be of the suitable interval type. When $f(M_1, \dots, M_m)$ is called and $f : T_1 \times \dots \times T_m \rightarrow T$, M_j must be of type T_j , and $f(M_1, \dots, M_m)$ is then of type T .

The term `new` $x[\tilde{i}] : T; N$ is accepted only when T is declared *fixed*, *bounded*, or *nonuniform*. We check that $x[\tilde{i}]$ is of type T (which is in fact always true when the construction of \mathcal{E} succeeds). N must well-typed, and its type is also the type of `new` $x[\tilde{i}] : T; N$.

In `let` $p = M$ in N `else` N' , p must have the same type as M , and N and N' must have the same type, which is also the type of `let` $p = M$ in N `else` N' . The typing rules for patterns p are found at the bottom of Figure 2. The pattern $x[\tilde{i}] : T$ has type T , provided $x[\tilde{i}]$ has type T (which is in fact always true when the construction of \mathcal{E} succeeds). The other typing rules for patterns are straightforward. The particular case `let` $x[\tilde{i}] : T = M$ in N is typed similarly, except that the `else` branch is omitted.

In

$$\text{find}[\text{unique?}] \left(\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat} \right. \\ \left. \text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } N_j \right) \text{ else } N$$

the replication indices i_{j1}, \dots, i_{jm_j} are bound in $M_{j1}, \dots, M_{jl_j}, M_j$, of types $[1, n_{j1}], \dots, [1, n_{jm_j}]$ respectively; M_j is of type *bool* for all $j \leq m$; N_j for all $j \leq m$ and N all have the same type, which is also the type of the `find` term.

In `insert` $Tbl(M_1, \dots, M_l); N$, M_1, \dots, M_l must be of the type declared for the elements of the table Tbl , and the type of N is the type of the `insert` term.

In `get` $Tbl(p_1, \dots, p_l)$ `suchthat` M in N `else` N' , p_1, \dots, p_l must be of the type declared for the elements of the table Tbl and M must be of type *bool*. The terms N and N' must have the same type, which is also the type of the `get` term.

In `event` $e(M_1, \dots, M_l); N$, M_1, \dots, M_l must be of the type declared for the arguments of event e , and the type of N is the type of the event term.

The term `event_abort` e can have any type (because it aborts the game); the event e must be declared without argument, which we denote by $e : ()$.

The type system for processes requires each subterm to be well-typed. In $!^{i \leq n} Q$, i is of type $[1, n]$ in Q . The processes `new`, `let`, `find`, `insert`, `get`, `event`, and `event_abort` are typed similarly to the corresponding terms.

We say that an occurrence of a term M in a process Q is of type T when $\mathcal{E} \vdash M : T$ where \mathcal{E} is the type environment of Q extended with $i \mapsto [1, n]$ for each replication $!^{i \leq n}$ above M

Typing rules for terms:

$$\begin{array}{c}
\frac{\mathcal{E}(i) = T}{\mathcal{E} \vdash i : T} \quad (\text{TIndex}) \\
\frac{\mathcal{E}(x) = T_1 \times \dots \times T_m \rightarrow T \quad \forall j \leq m, \mathcal{E} \vdash M_j : T_j}{\mathcal{E} \vdash x[M_1, \dots, M_m] : T} \quad (\text{TVar}) \\
\frac{f : T_1 \times \dots \times T_m \rightarrow T \quad \forall j \leq m, \mathcal{E} \vdash M_j : T_j}{\mathcal{E} \vdash f(M_1, \dots, M_m) : T} \quad (\text{TFun}) \\
\frac{T \text{ fixed, bounded, or nonuniform} \quad \mathcal{E} \vdash x[\tilde{i}] : T \quad \mathcal{E} \vdash N : T'}{\mathcal{E} \vdash \text{new } x[\tilde{i}] : T; N : T'} \quad (\text{TNewT}) \\
\frac{\mathcal{E} \vdash M : T \quad \mathcal{E} \vdash p : T \quad \mathcal{E} \vdash N : T' \quad \mathcal{E} \vdash N' : T'}{\mathcal{E} \vdash \text{let } p = M \text{ in } N \text{ else } N' : T'} \quad (\text{TLetT}) \\
\frac{\mathcal{E} \vdash M : T \quad \mathcal{E} \vdash x[\tilde{i}] : T \quad \mathcal{E} \vdash N : T'}{\mathcal{E} \vdash \text{let } x[\tilde{i}] : T = M \text{ in } N : T'} \quad (\text{TLetT2}) \\
\frac{\forall j \leq m, \forall k \leq m_j, \mathcal{E} \vdash u_{jk}[\tilde{i}] : [1, n_{jk}] \quad \forall j \leq m, \forall k \leq l_j, \mathcal{E}[i_{j1} \mapsto [1, n_{j1}], \dots, i_{jm_j} \mapsto [1, n_{jm_j}]] \vdash M_{jk} : T_{jk} \quad \forall j \leq m, \mathcal{E}[i_{j1} \mapsto [1, n_{j1}], \dots, i_{jm_j} \mapsto [1, n_{jm_j}]] \vdash M_j : \text{bool} \quad \forall j \leq m, \mathcal{E} \vdash N_j : T \quad \mathcal{E} \vdash N : T}{\mathcal{E} \vdash \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } N_j) \text{ else } N : T} \quad (\text{TFindT}) \\
\frac{Tbl : T_1 \times \dots \times T_l \quad \forall j \leq l, \mathcal{E} \vdash M_j : T_j \quad \mathcal{E} \vdash N : T}{\mathcal{E} \vdash \text{insert } Tbl(M_1, \dots, M_l); N : T} \quad (\text{TInsertT}) \\
\frac{Tbl : T_1 \times \dots \times T_l \quad \forall j \leq l, \mathcal{E} \vdash p_j : T_j \quad \mathcal{E} \vdash M : \text{bool} \quad \mathcal{E} \vdash N : T \quad \mathcal{E} \vdash N' : T}{\text{get } Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } N \text{ else } N' : T} \quad (\text{TGetT}) \\
\frac{e : T_1 \times \dots \times T_l \quad \forall j \leq l, \mathcal{E} \vdash M_j : T_j \quad \mathcal{E} \vdash N : T}{\mathcal{E} \vdash \text{event } e(M_1, \dots, M_l); N : T} \quad (\text{TEvent}) \\
\frac{e : ()}{\mathcal{E} \vdash \text{event.abort } e : T'} \quad (\text{TEventAbortT})
\end{array}$$

Typing rules for patterns:

$$\begin{array}{c}
\frac{\mathcal{E} \vdash x[\tilde{i}] : T}{\mathcal{E} \vdash (x[\tilde{i}] : T) : T} \quad (\text{TVarP}) \\
\frac{f : T_1 \times \dots \times T_m \rightarrow T \quad \forall j \leq m, \mathcal{E} \vdash p_j : T_j}{\mathcal{E} \vdash f(p_1, \dots, p_m) : T} \quad (\text{TFunP}) \\
\frac{\mathcal{E} \vdash M : T}{\mathcal{E} \vdash =M : T} \quad (\text{TEqP})
\end{array}$$

Figure 2: Typing rules (1)

Typing rules for input processes:

$$\begin{array}{c}
\mathcal{E} \vdash 0 \quad \text{(TNil)} \\
\frac{\mathcal{E} \vdash Q \quad \mathcal{E} \vdash Q'}{\mathcal{E} \vdash Q \mid Q'} \quad \text{(TPar)} \\
\frac{\mathcal{E}[i \mapsto [1, n]] \vdash Q}{\mathcal{E} \vdash !^{i \leq n} Q} \quad \text{(TRepl)} \\
\frac{\mathcal{E} \vdash Q}{\mathcal{E} \vdash \text{newChannel } c; Q} \quad \text{(TNewChannel)} \\
\frac{\forall j \leq l, \mathcal{E} \vdash M_j : T'_j \quad \mathcal{E} \vdash p : T \quad \mathcal{E} \vdash P}{\mathcal{E} \vdash c[M_1, \dots, M_l](p); P} \quad \text{(TIn)}
\end{array}$$

Typing rules for output processes:

$$\begin{array}{c}
\frac{\forall j \leq l, \mathcal{E} \vdash M_j : T'_j \quad \mathcal{E} \vdash N : T \quad \mathcal{E} \vdash Q}{\mathcal{E} \vdash c[M_1, \dots, M_l]\langle N \rangle; Q} \quad \text{(TOut)} \\
\frac{T \text{ fixed, bounded, or nonuniform} \quad \mathcal{E} \vdash x[\tilde{i}] : T \quad \mathcal{E} \vdash P}{\mathcal{E} \vdash \text{new } x[\tilde{i}] : T; P} \quad \text{(TNew)} \\
\frac{\mathcal{E} \vdash M : T \quad \mathcal{E} \vdash p : T \quad \mathcal{E} \vdash P \quad \mathcal{E} \vdash P'}{\mathcal{E} \vdash \text{let } p = M \text{ in } P \text{ else } P'} \quad \text{(TLet)} \\
\frac{\forall j \leq m, \forall k \leq m_j, \mathcal{E} \vdash u_{jk}[\tilde{i}] : [1, n_{jk}] \quad \forall j \leq m, \forall k \leq l_j, \mathcal{E}[i_{j1} \mapsto [1, n_{j1}], \dots, i_{jm_j} \mapsto [1, n_{jm_j}]] \vdash M_{jk} : T_{jk} \quad \forall j \leq m, \mathcal{E}[i_{j1} \mapsto [1, n_{j1}], \dots, i_{jm_j} \mapsto [1, n_{jm_j}]] \vdash M_j : \text{bool} \quad \forall j \leq m, \mathcal{E} \vdash P_j \quad \mathcal{E} \vdash P}{\mathcal{E} \vdash \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P} \quad \text{(TFind)} \\
\frac{Tbl : T_1 \times \dots \times T_l \quad \forall j \leq l, \mathcal{E} \vdash M_j : T_j \quad \mathcal{E} \vdash P}{\mathcal{E} \vdash \text{insert } Tbl(M_1, \dots, M_l); P} \quad \text{(TInsert)} \\
\frac{Tbl : T_1 \times \dots \times T_l \quad \forall j \leq l, \mathcal{E} \vdash p_j : T_j \quad \mathcal{E} \vdash M : \text{bool} \quad \mathcal{E} \vdash P \quad \mathcal{E} \vdash P'}{\text{get } Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } P \text{ else } P'} \quad \text{(TGet)} \\
\frac{e : T_1 \times \dots \times T_l \quad \forall j \leq l, \mathcal{E} \vdash M_j : T_j \quad \mathcal{E} \vdash P}{\mathcal{E} \vdash \text{event } e(M_1, \dots, M_l); P} \quad \text{(TEvent)} \\
\frac{e : ()}{\mathcal{E} \vdash \text{event_abort } e} \quad \text{(TEventAbort)} \\
\mathcal{E} \vdash \text{yield} \quad \text{(TYield)}
\end{array}$$

Figure 3: Typing rules (2)

in Q and with $i_{j_1} \mapsto [1, n_{j_1}], \dots, i_{j_{m_j}} \mapsto [1, n_{j_{m_j}}]$ for each $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j_1}[\tilde{i}] = i_{j_1} \leq n_{j_1}, \dots, u_{j_{m_j}}[\tilde{i}] = i_{j_{m_j}} \leq n_{j_{m_j}} \text{ suchthat defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j \text{ then } P_j) \text{ else } P$ such that the considered occurrence of M is in the condition $\text{defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j$.

Invariant 7 (Typing) The process Q_0 satisfies Invariant 7 if and only if the type environment \mathcal{E} for Q_0 is well-defined, and $\mathcal{E} \vdash Q_0$.

We require 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 values may appear at each point of the protocol.

2.4 Formal Semantics

2.4.1 Definition of the Semantics

The formal semantics of our calculus is presented in Figures 4, 5, 6, and 7. A semantic configuration is a sextuple $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$, where

- E is an environment mapping array cells to values.
- (σ, P) is the output process P currently scheduled, with the associated function σ which gives values of replication indices.
- \mathcal{Q} is the multiset of input processes running in parallel with P , with their associated substitutions giving values of replication indices.
- \mathcal{C} is the set of channels already created.
- \mathcal{T} defines the contents of tables. It is a list of $\text{Tbl}(a_1, \dots, a_m)$ indicating that table Tbl contains the element (a_1, \dots, a_m) .
- $\mathcal{E}v$ is the sequence of events $e(a_1, \dots, a_m)$ executed so far.

There is one exceptional configuration, of the form $\text{abort}, \mathcal{E}v$, corresponding to the situation in which the game has been aborted after executing the events in $\mathcal{E}v$.

The semantics is defined by reduction rules of the form $E, P, \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', P', \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$ meaning that $E, P, \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$ reduces to $E', P', \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$ with probability p . The index t just serves in distinguishing reductions that yield the same configuration with the same probability in different ways, so that the probability of a certain reduction can be computed correctly:

$$\Pr[E, P, \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \rightarrow E', P', \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'] = \sum_{E, P, \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', P', \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'} p$$

The probability of a trace $\text{Tr} = E_1, P_1, \mathcal{Q}_1, \mathcal{C}_1, \mathcal{T}_1, \mathcal{E}v_1 \xrightarrow{p_1}_{t_1} \dots \xrightarrow{p_{m-1}}_{t_{m-1}} E_m, P_m, \mathcal{Q}_m, \mathcal{C}_m, \mathcal{T}_m, \mathcal{E}v_m$ is $p_1 \times \dots \times p_{m-1}$. We define the semantics only for patterns $x[\tilde{i}] : T$, the other patterns can be encoded as outlined in Section 2.1.

In Figures 4 and 5, we define an auxiliary relation for evaluating terms: $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', \sigma, M', \mathcal{T}', \mathcal{E}v'$ means that the term M reduces to M' in environment E with the replication indices defined by σ , the table contents \mathcal{T} , and the sequence of events $\mathcal{E}v$, with probability p . Here, the terms M, M' can be values a in addition to the grammar to terms given in Figure 1. Rule (ReplIndex) evaluates replication indices using the function σ . Rule (Var) looks for the value of the variable in the environment E . Rule (Fun) evaluates the function call. Rule (NewT)

$$\begin{array}{c}
E, \sigma, i, \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, \sigma, \sigma(i), \mathcal{T}, \mathcal{E}v \quad (\text{ReplIndex}) \\
\frac{x[a_1, \dots, a_m] \in \text{Dom}(E)}{E, \sigma, x[a_1, \dots, a_m], \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, \sigma, E(x[a_1, \dots, a_m]), \mathcal{T}, \mathcal{E}v} \quad (\text{Var}) \\
\frac{f : T_1 \times \dots \times T_m \rightarrow T \quad \forall j \leq m, a_j \in T_j \quad f(a_1, \dots, a_m) = a}{E, \sigma, f(a_1, \dots, a_m), \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, \sigma, a, \mathcal{T}, \mathcal{E}v} \quad (\text{Fun}) \\
\frac{a \in T \quad E' = E[x[\sigma(\tilde{i})] \mapsto a]}{E, \sigma, \text{new } x[\tilde{i}] : T; N, \mathcal{T}, \mathcal{E}v \xrightarrow{D_T(a)} E', \sigma, N, \mathcal{T}, \mathcal{E}v} \quad (\text{NewT}) \\
\frac{a \in T \quad E' = E[x[\sigma(\tilde{i})] \mapsto a]}{E, \sigma, \text{let } x[\tilde{i}] : T = a \text{ in } N, \mathcal{T}, \mathcal{E}v \xrightarrow{1} E', \sigma, N, \mathcal{T}, \mathcal{E}v} \quad (\text{LetT}) \\
\begin{array}{l}
(v_k)_{1 \leq k \leq l} \text{ is the sequence of } (j, a_1, \dots, a_{m_j}) \text{ for } a_1 \in [1, n_{j1}], \dots, a_{m_j} \in [1, n_{jm_j}] \\
\text{ordered in increasing lexicographic order} \\
\exists l_0 \leq l, \forall k \in [1, l_0], E, \sigma[i_{j1} \mapsto a_1, \dots, i_{jm_j} \mapsto a_{m_j}], D_j \wedge M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{p_k^*}_{t_k} E'', \sigma', r_k, \mathcal{T}, \mathcal{E}v \\
\text{where } v_k = (j, a_1, \dots, a_{m_j}) \text{ and for } k < l_0, r_k \text{ is a value, } r_{l_0} = \text{event_abort } e
\end{array} \\
\hline
\begin{array}{l}
E, \sigma, \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat} \\
D_j \wedge M_j \text{ then } N_j) \text{ else } N, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_{l_0}}_{t_1 \dots t_{l_0}} E, \sigma, \text{event_abort } e, \mathcal{T}, \mathcal{E}v
\end{array} \quad (\text{FindTE}) \\
\begin{array}{l}
(v_k)_{1 \leq k \leq l} \text{ is the sequence of } (j, a_1, \dots, a_{m_j}) \text{ for } a_1 \in [1, n_{j1}], \dots, a_{m_j} \in [1, n_{jm_j}] \\
\text{ordered in increasing lexicographic order} \\
\forall k \in [1, l], E, \sigma[i_{j1} \mapsto a_1, \dots, i_{jm_j} \mapsto a_{m_j}], D_j \wedge M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{p_k^*}_{t_k} E'', \sigma', r_k, \mathcal{T}, \mathcal{E}v \\
\text{where } v_k = (j, a_1, \dots, a_{m_j}) \text{ and } r_k \text{ is a value} \\
S = \{v_k \mid r_k = \text{true}\} \quad |S| = 1 \text{ or } [\text{unique?}] \text{ is empty} \\
v_0 = (j', a'_1, \dots, a'_{m_{j'}}) \in S \quad E' = E[u_{j'1}[\sigma(\tilde{i})] \mapsto a'_1, \dots, u_{j'm_{j'}}[\sigma(\tilde{i})] \mapsto a'_{m_{j'}}]
\end{array} \\
\hline
\begin{array}{l}
E, \sigma, \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat} \\
D_j \wedge M_j \text{ then } N_j) \text{ else } N, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_l D_{\text{find}}(S)(v_0)}_{t_1 \dots t_l F1(v_0)} E', \sigma, N_{j'}, \mathcal{T}, \mathcal{E}v
\end{array} \quad (\text{FindT1}) \\
\hline
\begin{array}{l}
\text{First four lines as in (FindT1)} \quad S = \{v_k \mid r_k = \text{true}\} = \emptyset
\end{array} \quad (\text{FindT2}) \\
\hline
\begin{array}{l}
E, \sigma, \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat} \\
D_j \wedge M_j \text{ then } N_j) \text{ else } N, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_l}_{t_1 \dots t_l F2} E, \sigma, N, \mathcal{T}, \mathcal{E}v
\end{array} \\
\hline
\begin{array}{l}
\text{First four lines as in (FindT1)} \quad S = \{v_k \mid r_k = \text{true}\} \quad |S| > 1
\end{array} \quad (\text{FindT3}) \\
\hline
\begin{array}{l}
E, \sigma, \text{find}[\text{unique}_e] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat} \\
D_j \wedge M_j \text{ then } N_j) \text{ else } N, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_l}_{t_1 \dots t_l F3} E, \sigma, \text{event_abort } e, \mathcal{T}, \mathcal{E}v
\end{array}
\end{array}$$

Figure 4: Semantics (1): terms, first part

$$\begin{array}{c}
\frac{E, \sigma, \text{insert } Tbl(a_1, \dots, a_l); N, \mathcal{T}, \mathcal{E}v}{E, \sigma, N, (\mathcal{T}, Tbl(a_1, \dots, a_l)), \mathcal{E}v} \quad (\text{InsertT}) \\
\\
\frac{\begin{array}{c} [v_1, \dots, v_m] = [x \in \mathcal{T} \mid \exists a_1, \dots, \exists a_l, x = Tbl(a_1, \dots, a_l)] \\ \exists m_0 \leq m, \forall k \in [1, l_0], E[x_1[\sigma(\tilde{i})] \mapsto a_1, \dots, x_l[\sigma(\tilde{i})] \mapsto a_l], \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{p_k}_{t_k}^* E'', \sigma, r_k, \mathcal{T}, \mathcal{E}v \\ \text{where } v_k = Tbl(a_1, \dots, a_l) \text{ and for } k < m_0, r_k \text{ is a value, } r_{m_0} = \text{event_abort } e \end{array}}{E, \sigma, \text{get } Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } N \text{ else } N', \mathcal{T}, \mathcal{E}v} \\
\frac{\quad}{\xrightarrow{p_1 \dots p_{m_0}}_{t_1 \dots t_{m_0}}} E, \sigma, \text{event_abort } e, \mathcal{T}, \mathcal{E}v \quad (\text{GetTE}) \\
\\
\frac{\begin{array}{c} [v_1, \dots, v_m] = [x \in \mathcal{T} \mid \exists a_1, \dots, \exists a_l, x = Tbl(a_1, \dots, a_l)] \\ \forall k \in [1, m], E[x_1[\sigma(\tilde{i})] \mapsto a_1, \dots, x_l[\sigma(\tilde{i})] \mapsto a_l], \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{p_k}_{t_k}^* E'', \sigma, r_k, \mathcal{T}, \mathcal{E}v \\ \text{where } v_k = Tbl(a_1, \dots, a_l) \text{ and } r_k \text{ is a value} \\ L = [v_k \mid k \in [1, m], r_k = \text{true}] \\ Tbl(a_1, \dots, a_l) = \text{nth}(L, j) \quad E' = E[x_1[\sigma(\tilde{i})] \mapsto a_1, \dots, x_l[\sigma(\tilde{i})] \mapsto a_l] \end{array}}{E, \sigma, \text{get } Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } N \text{ else } N', \mathcal{T}, \mathcal{E}v} \quad (\text{GetT1}) \\
\frac{\quad}{\xrightarrow{p_1 \dots p_m D_{\text{get}}(\{1, \dots, L\})(j)}_{t_1 \dots t_m G1(j)}} E', \sigma, N, \mathcal{T}, \mathcal{E}v \\
\\
\frac{\text{First four lines as in (GetT1)} \quad L = \emptyset}{E, \sigma, \text{get } Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } N \text{ else } N', \mathcal{T}, \mathcal{E}v} \quad (\text{GetT2}) \\
\frac{\quad}{\xrightarrow{p_1 \dots p_m}_{t_1 \dots t_m G2}} E, \sigma, N', \mathcal{T}, \mathcal{E}v \\
\\
E, \sigma, \text{event } e(a_1, \dots, a_l); N, \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, \sigma, N, \mathcal{T}, (\mathcal{E}v, e(a_1, \dots, a_l)) \quad (\text{EventT}) \\
\\
\frac{E, \sigma, N, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', \sigma', N', \mathcal{T}', \mathcal{E}v'}{E, \sigma, C[N], \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E, \sigma', C[N'], \mathcal{T}', \mathcal{E}v'} \quad (\text{CtxT}) \\
\\
E, \sigma, C[\text{event_abort } e], \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, \sigma, \text{event_abort } e, \mathcal{T}, \mathcal{E}v \quad (\text{CtxEventT}) \\
\\
\frac{\neg \forall j \leq l, \exists a_j, E, \sigma, M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{1}^* E, \sigma, a_j, \mathcal{T}, \mathcal{E}v}{E, \sigma, \text{defined}(M_1, \dots, M_l) \wedge M, \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, \sigma, \text{false}, \mathcal{T}, \mathcal{E}v} \quad (\text{DefinedNo}) \\
\\
\frac{\forall j \leq l, \exists a_j, E, \sigma, M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{1}^* E, \sigma, a_j, \mathcal{T}, \mathcal{E}v}{E, \sigma, \text{defined}(M_1, \dots, M_l) \wedge M, \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, \sigma, M, \mathcal{T}, \mathcal{E}v} \quad (\text{DefinedYes})
\end{array}$$

Figure 5: Semantics (2): terms, second part, and defined conditions

chooses a random $a \in T$ according to distribution D_T , and stores it in $x[\sigma(\tilde{i})]$ by extending the environment E accordingly. Similarly, Rule (LetT) extends the environment E with the value of $x[\sigma(\tilde{i})]$.

Rules (FindTE) to (FindT3) define the semantics of `find`. First, they all evaluate the conditions for all branches j and all values of the indices i_{j1}, \dots, i_{jm_j} . If one of these evaluations executes an event (which can happen in case the condition contains a `find[uniquee]`), the whole `find` executes the same event (Rule (FindTE)). Otherwise, the branch and indices for which the condition is true are collected in a set S . If S is empty, the `else` branch of the `find` is executed (Rule (FindT2)). When S is not empty, two cases can happen. Either the `find` is not marked `[uniquee]`, and we choose an element $v_0 = (j', a'_1, \dots, a'_{m_{j'}}$ of S randomly according to the distribution $D_{\text{find}}(S)$, store the corresponding indices $a'_1, \dots, a'_{m_{j'}}$ in $u_{j'1}[\sigma(\tilde{i})], \dots, u_{j'm_{j'}}[\sigma(\tilde{i})]$ by extending the environment accordingly, and we continue with the selected branch N'_j . If the `find` is marked `[uniquee]` and S has a single element, we do the same. If the `find` is marked `[uniquee]` and S has several elements, we execute the event e (Rule (FindT3)). We recall that $D_{\text{find}}(S)(v_0)$ denotes the probability of choosing v_0 in the distribution $D_{\text{find}}(S)$. The terms in conditions of `find` may define variables, included in the environment E'' ; we ignore these additional variables and compute the final environment from the initial environment E , because these variables have no array accesses by Invariant 3, so the values of these variables are not used after the evaluation of the condition. The conditions of `find`, $D \wedge M$, are evaluated using Rules (DefinedNo) and (DefinedYes). An element of the `defined` condition D is not defined, the condition is false (Rule (DefinedNo)); when all elements of the `defined` condition are defined, we evaluate M (Rule (DefinedYes)). Since terms in conditions of `find` do not contain `insert` nor `event` (Invariant 4), the table contents and the sequence of events are left unchanged by the evaluation of the condition of `find`.

Rule (InsertT) inserts the new table element in \mathcal{T} . Rules (GetTE), (GetT1) and (GetT2) define the semantics of `get`. We denote by $[x \in L \mid f(x)]$ the list of all elements x of the list L that satisfy $f(x)$, in the same order as in L . We denote by $|L|$ the length of list L . We denote by $\text{nth}(L, j)$ the j -th element of the list L . Rule (GetTE) executes `event_abort e` when the evaluation of the condition M executes `event_abort e` for some element of the table Tbl . Rules (Get1) and (Get2) compute the list L of elements of \mathcal{T} that are in table Tbl and satisfy condition M . When L is not empty, one of its elements is chosen randomly according to distribution $D_{\text{get}}(\{1, \dots, |L|\})$, we store this element in $x_1[\sigma(\tilde{i})], \dots, x_l[\sigma(\tilde{i})]$ by extending the environment E , and continue by executing N (Rule (Get1)). If L is empty, we execute N' (Rule (Get2)). Since terms in conditions of `get` do not contain `insert` nor `event` (Invariant 4), the table contents and the sequence of events are left unchanged by the evaluation of M . The modified environment E'' obtained after evaluating a condition M can be ignored because there are no array accesses to the variables defined in conditions of `get`, by Invariant 3, so the values of these variables are not used after the evaluation of the condition.

Remark 1 Another way of defining the semantics of tables would be to consider two distinct calculi, one with tables but without `find` and `defined` conditions (used for the initial game), and one with `find` and `defined` conditions but without tables (used for the other games). The semantics of the calculus with `find` and `defined` conditions but without tables can be defined without the component \mathcal{T} . The semantics of the calculus with tables but without `find` and `defined` conditions can be defined with one environment for each process, instead of a global environment, and no arrays (only the value at the current replication indices is stored). We then need to relate the two semantics.

Rule EventT adds the executed event to $\mathcal{E}v$.

Rules (CtxT) and (CtxEventT) allow evaluating terms under a context. In these rules, C

$$\begin{array}{c}
E, \{(\sigma, 0)\} \uplus \mathcal{Q}, \mathcal{C} \rightsquigarrow E, \mathcal{Q}, \mathcal{C} \quad (\text{Nil}) \\
E, \{(\sigma, Q_1 \mid Q_2)\} \uplus \mathcal{Q}, \mathcal{C} \rightsquigarrow E, \{(\sigma, Q_1), (\sigma, Q_2)\} \uplus \mathcal{Q}, \mathcal{C} \quad (\text{Par}) \\
E, \{(\sigma, !^{i \leq n} Q)\} \uplus \mathcal{Q}, \mathcal{C} \rightsquigarrow E, \{(\sigma[i \mapsto a], Q) \mid a \in [1, n]\} \uplus \mathcal{Q}, \mathcal{C} \quad (\text{Repl}) \\
\frac{c' \notin \mathcal{C}}{E, \{(\sigma, \text{newChannel } c; Q)\} \uplus \mathcal{Q}, \mathcal{C} \rightsquigarrow E, \{(\sigma, Q\{c'/c\})\} \uplus \mathcal{Q}, \mathcal{C} \cup \{c'\}} \quad (\text{NewChannel}) \\
\frac{\forall j \leq l, E, \sigma, M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{1}^* E, \sigma, a_j, \mathcal{T}, \mathcal{E}v}{E, \{(\sigma, c[M_1, \dots, M_l](x[\tilde{i}] : T); P)\} \uplus \mathcal{Q}, \mathcal{C} \rightsquigarrow E, \{(\sigma, c[a_1, \dots, a_l](x[\tilde{i}] : T); P)\} \uplus \mathcal{Q}, \mathcal{C}} \quad (\text{Input}) \\
\text{reduce}(E, \mathcal{Q}, \mathcal{C}) \text{ is the normal form of } E, \mathcal{Q}, \mathcal{C} \text{ by } \rightsquigarrow
\end{array}$$

Figure 6: Semantics (3): input processes

is an elementary context, of the form:

$$\begin{array}{l}
C ::= x[a_1, \dots, a_k, [], M_{k+2}, \dots, M_m] \\
\quad f(a_1, \dots, a_k, [], M_{k+2}, \dots, M_m) \\
\quad \text{let } x[\tilde{i}] : T = [] \text{ in } N \\
\quad \text{event } e(a_1, \dots, a_k, [], M_{k+2}, \dots, M_l); N \\
\quad \text{insert } Tbl(a_1, \dots, a_k, [], M_{k+2}, \dots, M_l); N
\end{array}$$

When the term N reduces to some other term N' , Rule (CtxT) allows one to reduce it in the same way under a context C . When the term N is an event, $C[N]$ also executes the same event by Rule (CtxEventT).

These rules define a small-step semantics for terms. We consider the reflexive and transitive closure \xrightarrow{p}_t^* of the relation \xrightarrow{p}_t to reach directly the normal form of the term, which can be either a value a or an event `event_abort` e . We have $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{1}^* E, \sigma, M, \mathcal{T}, \mathcal{E}v$ and, if $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mathcal{E}v'$ and $E', \sigma', M', \mathcal{T}', \mathcal{E}v' \xrightarrow{p'}_{t'}^* E'', \sigma'', M'', \mathcal{T}'', \mathcal{E}v''$, then $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{p \times p'}_{t, t'}^* E'', \sigma'', M'', \mathcal{T}'', \mathcal{E}v''$: we take the product of the probabilities to have the probability of a sequence of reductions, and we specify which sequence was taken by a list of indices t, t' .

Figure 6 defines the semantics of input processes. We use an auxiliary reduction relation \rightsquigarrow , for reducing input processes. This relation transforms configurations of the form $E, \mathcal{Q}, \mathcal{C}$. Rule (Nil) removes nil processes. Rules (Par) and (Repl) expand parallel compositions and replications, respectively. Rule (NewChannel) creates a new channel and adds it to \mathcal{C} . Semantic configurations are considered equivalent modulo renaming of channels in \mathcal{C} , so that a single semantic configuration is obtained after applying (NewChannel). Rule (Input) evaluates the terms in the input channel. The input itself is not executed: the communication is done by the (Output) rule. In the (Input) rule, the terms M_1, \dots, M_l contain only replication indices, variables, and function applications by Invariant 5, so their evaluation is deterministic (the unique result is obtained with probability 1), and the environment E , the contents of tables \mathcal{T} , and the sequence of events $\mathcal{E}v$ are unchanged, that is why we can write $E, \sigma, M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{1}^* E, \sigma, a_j, \mathcal{T}, \mathcal{E}v$. The relation \rightsquigarrow is convergent (confluent and terminating), so it has normal forms. Processes in \mathcal{Q} in configurations $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$ are always in normal form by \rightsquigarrow , so they always start with an input.

Finally, Figure 7 defines the semantics of output processes. The rules (New) to (Event) are very similar to those for terms: they just use processes instead of terms as continuations, and include a whole semantic configuration. Rule (Output) performs communications: it selects an

$$\begin{array}{c}
\frac{a \in T \quad E' = E[x[\sigma(\tilde{i})] \mapsto a]}{E, (\sigma, \text{new } x[\tilde{i}] : T; P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{D_T(a)}_{N(a)} E', (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v} \quad \text{(New)} \\
\frac{a \in T \quad E' = E[x[\sigma(\tilde{i})] \mapsto a]}{E, (\sigma, \text{let } x[\tilde{i}] : T = a \text{ in } P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{1} E', (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v} \quad \text{(Let)} \\
\frac{\text{Same assumption as in (FindTE)}}{E, (\sigma, \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat } D_j \wedge M_j \text{ then } P_j) \text{ else } P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_{l_0}}_{t_1 \dots t_{l_0}} \text{abort}, (\mathcal{E}v, e)} \quad \text{(FindE)} \\
\frac{\text{First four lines as in (FindT1)} \quad S = \{v_k \mid r_k = \text{true}\} \quad |S| = 1 \text{ or } [\text{unique?}] \text{ is empty} \\ v_0 = (j', a'_1, \dots, a'_{m_{j'}}) \in S \quad E' = E[u_{j'1}[\sigma(\tilde{i})] \mapsto a'_1, \dots, u_{j'm_{j'}}[\sigma(\tilde{i})] \mapsto a'_{m_{j'}}]}{E, (\sigma, \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat } D_j \wedge M_j \text{ then } P_j) \text{ else } P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_l D_{\text{find}}(S)(v_0)}_{t_1 \dots t_l F1(v_0)} E', (\sigma, P_{j'}), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v} \quad \text{(Find1)} \\
\frac{\text{First four lines as in (FindT1)} \quad S = \{v_k \mid r_k = \text{true}\} = \emptyset}{E, (\sigma, \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat } D_j \wedge M_j \text{ then } P_j) \text{ else } P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_l}_{t_1 \dots t_l F2} E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v} \quad \text{(Find2)} \\
\frac{\text{First four lines as in (FindT1)} \quad S = \{v_k \mid r_k = \text{true}\} \quad |S| > 1}{E, (\sigma, \text{find}[\text{unique}_e] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat } D_j \wedge M_j \text{ then } P_j) \text{ else } P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_l}_{t_1 \dots t_l F3} \text{abort}, (\mathcal{E}v, e)} \quad \text{(Find3)} \\
\frac{S = \{(\sigma', Q) \in \mathcal{Q} \mid Q = c[a_1, \dots, a_l](x'[\tilde{i}] : T').P' \text{ and } b \in T' \text{ for some } x', \tilde{a}'', T', P'\} \\ (\sigma', Q_0) \in S \quad Q_0 = c[a_1, \dots, a_l](x[\tilde{i}] : T).P}{E, \mathcal{Q}', \mathcal{C}' = \text{reduce}(E, \{(\sigma, Q'')\}, \mathcal{C}) \\ E, (\sigma, \overline{c[a_1, \dots, a_l]}(b).Q''), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{S(Q_0) \times D_{\text{in}}(S)(Q_0)}_{O(Q_0)} \\ E[x[\sigma'(\tilde{i})] \mapsto b], P, \mathcal{Q} \uplus \mathcal{Q}' \setminus \{(\sigma', Q_0)\}, \mathcal{C}', \mathcal{T}, \mathcal{E}v} \quad \text{(Output)} \\
E, (\sigma, \text{insert } Tbl(a_1, \dots, a_l); P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, (\sigma, P), \mathcal{Q}, \mathcal{C}, (\mathcal{T}, Tbl(a_1, \dots, a_l)), \mathcal{E}v \quad \text{(Insert)} \\
\frac{\text{Same assumption as in (GetTE)}}{E, (\sigma, \text{get } Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } P \text{ else } P'), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_{m_0}}_{t_1 \dots t_{m_0}} \text{abort}, (\mathcal{E}v, e)} \quad \text{(GetE)} \\
\frac{\text{First four lines as in (GetT1)} \\ Tbl(a_1, \dots, a_l) = \text{nth}(L, j) \quad E' = E[x_1[\sigma(\tilde{i})] \mapsto a_1, \dots, x_l[\sigma(\tilde{i})] \mapsto a_l]}{E, (\sigma, \text{get } Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } P \text{ else } P'), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_m D_{\text{get}}(\{1, \dots, |L|\})(j)}_{t_1 \dots t_m G1(j)} E', (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v} \quad \text{(Get1)} \\
\frac{\text{First four lines as in (GetT1)} \quad L = \emptyset}{E, (\sigma, \text{get } Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } P \text{ else } P'), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p_1 \dots p_m}_{t_1 \dots t_m G2} E, (\sigma, P'), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v} \quad \text{(Get2)} \\
E, (\sigma, \text{event } e(a_1, \dots, a_l); P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, (\mathcal{E}v, e(a_1, \dots, a_l)) \quad \text{(Event)} \\
E, (\sigma, \text{event_abort } e), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{1} \text{abort}, (\mathcal{E}v, e) \quad \text{(EventAbort)} \\
\frac{E, \sigma, N, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', \sigma', N', \mathcal{T}', \mathcal{E}v'}{E, (\sigma, C[N]), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E, (\sigma', C[N']), \mathcal{Q}, \mathcal{C}, \mathcal{T}', \mathcal{E}v'} \quad \text{(Ctx)} \\
E, (\sigma, C[\text{event_abort } e]), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{1} \text{abort}, (\mathcal{E}v, e) \quad \text{(CtxEvent)}
\end{array}$$

input on the desired channel randomly, and immediately executes the communication. (The process blocks if no suitable input is available.) The scheduled process after this rule is the receiving process. The input processes that follow the output are stored in the available input processes, after reducing them by rules of Figure 6. In this rule, S is a multiset. When we take probabilities over multisets, we consider that $D_{\text{in}}(S)(Q_0)$ is the probability of choosing *one* of the elements equal to Q_0 in S according to the distribution $D_{\text{in}}(S)$, so that the probability of choosing any element equal to Q_0 is in fact $S(Q_0) \times D_{\text{in}}(S)(Q_0)$.

Rule EventAbort executes event e and aborts the game, by reducing to the special configuration `abort`, $(\mathcal{E}v, e)$.

Similarly to the case of terms, Rules (Ctx) and (CtxEvent) allow evaluating terms under a context inside a process. In these rules, C is an elementary context, of the form:

$$\begin{aligned}
C ::= & \text{let } x[\tilde{i}] : T = [] \text{ in } P \\
& \overline{c[a_1, \dots, a_k, [], M_{k+2}, \dots, M_l] \langle N \rangle}.Q \\
& \overline{c[a_1, \dots, a_l] \langle [] \rangle}.Q \\
& \text{insert } Tbl(a_1, \dots, a_k, [], M_{k+2}, \dots, M_l); P \\
& \text{event } e(a_1, \dots, a_k, [], M_{k+2}, \dots, M_l); P
\end{aligned}$$

After finishing execution of a process, the system produces two results: the sequence of executed events $\mathcal{E}v$, and the information whether we aborted the game ($a = \text{abort}$) or it terminated normally ($a = 0$). These events and result can be used to distinguish games, so we introduce an additional algorithm, a *distinguisher* D that takes as input the sequence of events $\mathcal{E}v$ and the result a , and returns true or false.

An example of distinguisher is D_e defined by $D_e(\mathcal{E}v, a) = \text{true}$ if and only if $e \in \mathcal{E}v$: this distinguisher detects the execution of event e . We will denote the distinguisher D_e simply by e . More generally, distinguishers can detect various properties of the sequence of events $\mathcal{E}v$ executed by the game and of its result a . We denote by $D \vee D'$, $D \wedge D'$, and $\neg D$ the distinguishers such that $(D \vee D')(\mathcal{E}v, a) = D(\mathcal{E}v, a) \vee D'(\mathcal{E}v, a)$, $(D \wedge D')(\mathcal{E}v, a) = D(\mathcal{E}v, a) \wedge D'(\mathcal{E}v, a)$, and $(\neg D)(\mathcal{E}v, a) = \neg D(\mathcal{E}v, a)$. We denote by $\text{Pr}[Q : D]$ the probability that Q executes a sequence of events $\mathcal{E}v$ and returns a result a , such that $D(\mathcal{E}v, a) = \text{true}$. This is formally defined as follows.

Definition 3 The initial configuration for running process Q is $\text{initConfig}(Q) = \emptyset, \overline{\text{start}} \langle \rangle, \mathcal{Q}, \mathcal{C}, \emptyset, \emptyset$ where $\emptyset, \mathcal{Q}, \mathcal{C} = \text{reduce}(\emptyset, \{(\sigma_0, Q)\}, \text{fc}(Q))$ and σ_0 is the empty function.

Let $\mathcal{T}r$ be the set of traces $Tr = \text{initConfig}(Q) \xrightarrow{p_1}_{t_1} \dots \xrightarrow{p_{m-1}}_{t_{m-1}} \text{Conf}_m$ such that Conf_m cannot be reduced.

For such traces, we define $\text{Pr}[Tr] = p_1 \times \dots \times p_m$, and $D(Tr) = D(\mathcal{E}v_m, \text{abort})$ if $\text{Conf}_m = \text{abort}, \mathcal{E}v_m$ and $D(Tr) = D(\mathcal{E}v_m, 0)$ if $\text{Conf}_m = E_m, P_m, \mathcal{Q}_m, \mathcal{C}_m, \mathcal{T}_m, \mathcal{E}v_m$.

We have $\text{Pr}[Q : D] = \sum_{Tr \in \mathcal{T}r, D(Tr) = \text{true}} \text{Pr}[Tr]$.

2.4.2 Each Variable is Defined at Most Once

In this section, we show that Invariant 1 implies that each array cell is assigned at most once during the execution of a process.

When S and S' are multisets, $\max(S, S')$ is the multiset such that $\max(S, S')(x) = \max(S(x), S'(x))$. We define the multiset of variable accesses that may be defined by a term or a process (given the replication indices fixed by a function σ) as follows:

$$\begin{aligned}
\text{Defined}(\sigma, i) &= \text{Defined}(\sigma, a) = \emptyset \\
\text{Defined}(\sigma, x[M_1, \dots, M_m]) &= \biguplus_{j=1}^m \text{Defined}(\sigma, M_j)
\end{aligned}$$

$$\text{Defined}(\sigma, f(M_1, \dots, M_m)) = \bigoplus_{j=1}^m \text{Defined}(\sigma, M_j)$$

$$\text{Defined}(\sigma, \text{new } x[\tilde{i}] : T; N) = \{x[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, N)$$

$$\text{Defined}(\sigma, \text{let } x[\tilde{i}] : T = M \text{ in } N) = \{x[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, M) \uplus \text{Defined}(\sigma, N)$$

$$\text{Defined}(\sigma, \text{find } (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j \text{ then } N_j \text{ else } N) = \\ \max(\max_{j=1}^m \{\tilde{u}_j[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, N_j), \text{Defined}(\sigma, N))$$

$$\text{Defined}(\sigma, \text{insert } \text{Tbl}(M_1, \dots, M_l); N) = \bigoplus_{j=1}^l \text{Defined}(\sigma, M_j) \uplus \text{Defined}(\sigma, N)$$

$$\text{Defined}(\sigma, \text{get } \text{Tbl}(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } N \text{ else } N') = \\ \max(\{x_j[\sigma(\tilde{i})] \mid j \leq l\} \uplus \text{Defined}(\sigma, N), \text{Defined}(\sigma, N'))$$

$$\text{Defined}(\sigma, \text{event } e(M_1, \dots, M_l); N) = \bigoplus_{j=1}^l \text{Defined}(\sigma, M_j) \uplus \text{Defined}(\sigma, N)$$

$$\text{Defined}(\sigma, \text{event_abort } e) = \emptyset$$

$$\text{Defined}(\sigma, 0) = \emptyset$$

$$\text{Defined}(\sigma, Q_1 \mid Q_2) = \text{Defined}(\sigma, Q_1) \uplus \text{Defined}(\sigma, Q_2)$$

$$\text{Defined}(\sigma, !^{i \leq n} Q) = \bigoplus_{a \in [1, n]} \text{Defined}(\sigma [i \mapsto a], Q)$$

$$\text{Defined}(\sigma, \text{newChannel } c; Q) = \text{Defined}(\sigma, Q)$$

$$\text{Defined}(\sigma, c[M_1, \dots, M_l](x[\tilde{i}] : T); P) = \{x[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, P)$$

$$\text{Defined}(\sigma, \overline{c[M_1, \dots, M_l]} \langle N \rangle; Q) = \bigoplus_{j=1}^l \text{Defined}(\sigma, M_j) \uplus \text{Defined}(\sigma, N) \uplus \text{Defined}(\sigma, Q)$$

$$\text{Defined}(\sigma, \text{new } x[\tilde{i}] : T; P) = \{x[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, P)$$

$$\text{Defined}(\sigma, \text{let } x[\tilde{i}] : T = M \text{ in } P) = \{x[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, M) \uplus \text{Defined}(\sigma, P)$$

$$\text{Defined}(\sigma, \text{find } (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j \text{ then } P_j) \text{ else } P) = \\ \max(\max_{j=1}^m \{\tilde{u}_j[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, P_j), \text{Defined}(\sigma, P))$$

$$\text{Defined}(\sigma, \text{insert } \text{Tbl}(M_1, \dots, M_l); P) = \bigoplus_{j=1}^l \text{Defined}(\sigma, M_j) \uplus \text{Defined}(\sigma, P)$$

$$\text{Defined}(\sigma, \text{get } \text{Tbl}(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } P \text{ else } P') = \\ \max(\{x_j[\sigma(\tilde{i})] \mid j \leq l\} \uplus \text{Defined}(\sigma, P), \text{Defined}(\sigma, P'))$$

$$\text{Defined}(\sigma, \text{event } e(M_1, \dots, M_l); P) = \bigoplus_{j=1}^l \text{Defined}(\sigma, M_j) \uplus \text{Defined}(\sigma, P)$$

$$\text{Defined}(\sigma, \text{event_abort } e) = \emptyset$$

Notice that, by Invariant 5, the terms M_j in channels of inputs and the terms M_{j_k} in defined conditions of find do not define any variable. By Invariant 3, the variables defined in conditions of find and get can be ignored. We define $\text{Defined}(E) = \text{Dom}(E)$, $\text{Defined}(E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v) = \text{Defined}(E) \uplus \text{Defined}(\sigma, P) \uplus \bigoplus_{(\sigma, Q) \in \mathcal{Q}} \text{Defined}(\sigma, Q)$.

Invariant 8 (Single definition, for executing games) The semantic configuration $E, (\sigma,$

$P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$ satisfies Invariant 8 if and only if $\text{Defined}(E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v)$ does not contain duplicate elements.

Lemma 1 *If Q_0 satisfies Invariant 1, then $\text{initConfig}(Q_0)$ satisfies Invariant 8.*

Lemma 2 *If $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{P}_t E', P', \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$ and $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$ satisfies Invariant 8, then so does $E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$.*

Proof sketch We show by induction following the definition of \xrightarrow{P}_t that

- if $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{P}_t E', \sigma', M', \mathcal{T}', \mathcal{E}v'$ then $\text{Defined}(E) \uplus \text{Defined}(\sigma, M) \supseteq \text{Defined}(E') \uplus \text{Defined}(\sigma', M')$.
- if $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{P}_t E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$ then $\text{Defined}(E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v) \supseteq \text{Defined}(E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v')$.

The result follows. □

Therefore, if Q_0 satisfies Invariant 1, then each variable is defined at most once for each value of its array indices in a trace of Q_0 . Indeed, by Invariant 8, just before executing a definition of $x[\tilde{a}]$, $\text{Defined}(E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v)$ does not contain duplicate elements, so $x[\tilde{a}] \notin \text{Dom}(E)$ since $x[\tilde{a}] \in \text{Defined}(\sigma, P)$.

2.4.3 Variables are Defined Before Being Used

In this section, we show that Invariant 2 implies that all variables are defined before being used. In order to show this property, we use the following invariant:

Invariant 9 (Defined variables, for executing games) The semantic configuration $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$ satisfies Invariant 9 if and only if every occurrence of a variable access $x[M_1, \dots, M_m]$ in (σ, P) or \mathcal{Q} is either

1. present in $\text{Dom}(E)$: if $x[M_1, \dots, M_m]$ occurs in a process P' for $(\sigma', P') \in \{(\sigma, P)\} \cup \mathcal{Q}$, then for all $j \leq m$, $E, \sigma', M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma', a_j, \mathcal{T}, \mathcal{E}v$ and $x[a_1, \dots, a_m] \in \text{Dom}(E)$;
2. or syntactically under the definition of $x[M_1, \dots, M_m]$ (in which case for all $j \leq m$, M_j is a constant or variable replication index);
3. or in a defined condition in a find process or term;
4. or in M'_j in a process or term of the form $\text{find } (\bigoplus_{j=1}^{m''} \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(M'_{j_1}, \dots, M'_{j_{l_j}}) \wedge M'_j \text{ then } P_j) \text{ else } P$ where for some $k \leq l_j$, $x[M_1, \dots, M_m]$ is a subterm of M'_{j_k} .
5. or in P_j in a process or term of the form $\text{find } (\bigoplus_{j=1}^{m''} \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(M'_{j_1}, \dots, M'_{j_{l_j}}) \wedge M'_j \text{ then } P_j) \text{ else } P$ where for some $k \leq l_j$, there is a subterm N of M'_{j_k} such that $N\{\tilde{u}_j[\tilde{i}]/\tilde{i}_j\} = x[M_1, \dots, M_m]$.

Similarly, $E, \sigma, M, \mathcal{T}, \mathcal{E}v$ satisfies Invariant 9 if and only if every occurrence of a variable access $x[M_1, \dots, M_m]$ in M either is present in $\text{Dom}(E)$ (for all $j \leq m$, $E, \sigma, M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma, a_j, \mathcal{T}, \mathcal{E}v$ and $x[a_1, \dots, a_m] \in \text{Dom}(E)$) or satisfies one of the last four conditions above.

$E, \sigma, \text{defined}(M'_1, \dots, M'_l) \wedge M, \mathcal{T}, \mathcal{E}v$ satisfies Invariant 9 if and only if every occurrence of a variable access $x[M_1, \dots, M_m]$ in M either is a subterm of M'_1, \dots, M'_l , or is present in $\text{Dom}(E)$ (for all $j \leq m$, $E, \sigma, M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma, a_j, \mathcal{T}, \mathcal{E}v$ and $x[a_1, \dots, a_m] \in \text{Dom}(E)$) or satisfies one of the last four conditions above.

Recall that, by Invariants 2 and 5, the terms of all variable accesses $x[M_1, \dots, M_m]$ contain only replication indices, variables, and function applications. That is why we can evaluate them by $E, \sigma', M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma', a_j, \mathcal{T}, \mathcal{E}v$.

Lemma 3 *If Q_0 satisfies Invariant 2, then $\text{initConfig}(Q_0)$ satisfies Invariant 9.*

Lemma 4 *Let M be a term that contains only replication indices, variables, and functions. If $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma, a, \mathcal{T}, \mathcal{E}v$, then for all subterms $x[M_1, \dots, M_m]$ of M , for all $j' \leq m$, $E, \sigma, M_{j'}, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma, a_{j'}, \mathcal{T}, \mathcal{E}v$ and $x[a_1, \dots, a_m]$ is in $\text{Dom}(E)$.*

Proof sketch By induction on M . □

Lemma 5 *Let N, M be terms that contain only replication indices, variables, and functions. If $E, \sigma[i \mapsto a'], N, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma[i \mapsto a'], a, \mathcal{T}, \mathcal{E}v$ and $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma, a', \mathcal{T}, \mathcal{E}v$, then $E, \sigma, N\{M/i\}, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma, a, \mathcal{T}, \mathcal{E}v$.*

Proof sketch By induction on N . □

Lemma 6 *If $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{P_t} E', \sigma', M', \mathcal{T}', \mathcal{E}v'$ and $E, \sigma, M, \mathcal{T}, \mathcal{E}v$ satisfies Invariant 9, then so does $E', \sigma', M', \mathcal{T}', \mathcal{E}v'$.*

If $E, \sigma, \text{defined}(M_1, \dots, M_m) \wedge M, \mathcal{T}, \mathcal{E}v \xrightarrow{P_t} E', \sigma', M', \mathcal{T}', \mathcal{E}v'$ and $E, \sigma, \text{defined}(M_1, \dots, M_m) \wedge M, \mathcal{T}, \mathcal{E}v$ satisfies Invariant 9, then so does $E', \sigma', M', \mathcal{T}', \mathcal{E}v'$.

If $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{P_t} E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$ and $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$ satisfies Invariant 9, then so does $E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$.

Moreover, if the rules that define $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{P_t} E', \sigma', M', \mathcal{T}', \mathcal{E}v'$ (resp. $E, \sigma, \text{defined}(M_1, \dots, M_m) \wedge M, \mathcal{T}, \mathcal{E}v \xrightarrow{P_t} E', \sigma', M', \mathcal{T}', \mathcal{E}v'$ or $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{P_t} E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$) require as assumption $E'', \sigma'', M'', \mathcal{T}'', \mathcal{E}v'' \xrightarrow{P_t} \dots$ or $E'', \sigma'', \text{defined}(M''_1, \dots, M''_m) \wedge M'', \mathcal{T}'', \mathcal{E}v'' \xrightarrow{P_t} \dots$, and the initial configuration $E, \sigma, M, \mathcal{T}, \mathcal{E}v$ (resp. $E, \sigma, \text{defined}(M_1, \dots, M_m) \wedge M, \mathcal{T}, \mathcal{E}v$ or $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$) satisfies Invariant 9, then so does the initial configuration of the assumption, $E'', \sigma'', M'', \mathcal{T}'', \mathcal{E}v''$ or $E'', \sigma'', \text{defined}(M''_1, \dots, M''_m) \wedge M'', \mathcal{T}'', \mathcal{E}v''$.

Proof sketch The proof proceeds by induction following the definition of $\xrightarrow{P_t}$. We just sketch the main arguments.

If $x[M_1, \dots, M_m]$ is in the second case of Invariant 9, and we execute the definition of $x[M_1, \dots, M_m]$, then for all $j \leq m$, M_j is a variable replication index and $x[\sigma(M_1), \dots, \sigma(M_m)]$ is added to $\text{Dom}(E)$ by rules (NewT), (LetT), (FindT1), (GetT1), (New), (Let), (Find1), (Output), or (Get1) so it moves to the first case of Invariant 9.

If $x[M_1, \dots, M_m]$ is in the third case of Invariant 9, and we execute the corresponding find, this access to x simply disappears.

If $x[M_1, \dots, M_m]$ is in the fourth case of Invariant 9, and we execute the find, then $x[M_1, \dots, M_m]$ is a subterm of M'_{jk} for some $j \leq m''$ and $k \leq l_j$. Therefore, the initial configuration of the assumption $E, \sigma[\tilde{i}_j \mapsto \tilde{a}], D_j \wedge M'_j, \mathcal{T}, \mathcal{E}v \xrightarrow{P_{k'}^*} E'', \sigma', r_{k'}, \mathcal{T}, \mathcal{E}v$ with $D_j \wedge M'_j = \text{defined}(M'_{j_1}, \dots, M'_{j_{l_j}}) \wedge M'_j$ and $\sigma' = \sigma[\tilde{i}_j \mapsto \tilde{a}]$ also satisfies Invariant 9. In case this assumption is reduced by (DefinedYes), we have $E, \sigma', M'_{jk}, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma', a_{jk}, \mathcal{T}, \mathcal{E}v$. Therefore, by Lemma 4, for all $j' \leq m$, $E, \sigma', M_{j'}, \mathcal{T}, \mathcal{E}v \xrightarrow{1^*} E, \sigma', a_{j'}, \mathcal{T}, \mathcal{E}v$ and $x[a_1, \dots, a_m]$ is in $\text{Dom}(E)$. So $x[M_1, \dots, M_m]$ moves to the first case of Invariant 9 in $E, \sigma', M'_j, \mathcal{T}, \mathcal{E}v$ after reduction by (DefinedYes).

If $x[M_1, \dots, M_m]$ is in the last case of Invariant 9, and we execute the find selecting branch j by (FindT1) or (Find1), then there is a subterm N of M'_{jk} for some $k \leq l_j$ such that $N\{\tilde{u}_j[\tilde{i}]/\tilde{i}_j\} = x[M_1, \dots, M_m]$. By hypothesis of (FindT1) or (Find1), we have $E, \sigma[\tilde{i}_j \mapsto$

$\tilde{a}']$, $D_j \wedge M'_{j'} \mathcal{T}, \mathcal{E}v \xrightarrow{pk'}^* E, \sigma', r_{k'}, \mathcal{T}, \mathcal{E}v$ where $r_{k'} = \text{true}$, $v_0 = (j, \tilde{a}') \in S$, $D_j \wedge M'_{j'} = \text{defined}(M'_{j_1}, \dots, M'_{j_{l_j}}) \wedge M'_{j'}$, and $\sigma' = \sigma[\tilde{i}_j \mapsto \tilde{a}']$. This assumption cannot reduce by (DefinedNo) because the result is true, so it reduces by (DefinedYes). Therefore, we have $E, \sigma', M'_{j_k}, \mathcal{T}, \mathcal{E}v \xrightarrow{1}^* E, \sigma', a, \mathcal{T}, \mathcal{E}v$ for some a . The term $N = x[N_1, \dots, N_m]$ is a subterm of M'_{j_k} . Therefore, by Lemma 4, for all $j' \leq m$, $E, \sigma', N_{j'}, \mathcal{T}, \mathcal{E}v \xrightarrow{1}^* E, \sigma', a_{j'}, \mathcal{T}, \mathcal{E}v$ and $x[a_1, \dots, a_m]$ is in $\text{Dom}(E)$. Moreover, the resulting environment E' is an extension of E , so a fortiori for all $j' \leq m$, $E', \sigma', N_{j'}, \mathcal{T}, \mathcal{E}v \xrightarrow{1}^* E', \sigma', a_{j'}, \mathcal{T}, \mathcal{E}v$ and $x[a_1, \dots, a_m]$ is in $\text{Dom}(E')$. We have for all $j' \leq m$, $M_{j'} = N_{j'}\{\tilde{u}_j[\tilde{i}]/\tilde{i}_j\}$, $E'(\tilde{u}_j[\tilde{i}]) = \tilde{a}'$, and $\sigma'(\tilde{i}_j) = \tilde{a}'$, so by Lemma 5, for all $j' \leq m$, $E', \sigma, M_{j'}, \mathcal{T}, \mathcal{E}v \xrightarrow{1}^* E', \sigma, a_{j'}, \mathcal{T}, \mathcal{E}v$ and $x[a_1, \dots, a_m]$ is in $\text{Dom}(E')$. So $x[M_1, \dots, M_m]$ also moves to the first case of Invariant 9.

In all other cases, the situation remains unchanged. For context rules, this is because, in the allowed contexts, the hole is never under a **defined** condition. \square

Therefore, if Q_0 satisfies Invariant 2, then in traces of Q_0 , the test $x[a_1, \dots, a_m] \in \text{Dom}(E)$ in rule (Var) always succeeds, except when the considered term occurs in a **defined** condition of a **find**.

Indeed, consider an application of rule (Var), where the array access $x[M_1, \dots, M_m]$ is not in a **defined** condition of a **find**. Then, this array access is not under any variable definition or **find**, so it is present in $\text{Dom}(E)$: for all $j \leq m$, $E, \sigma, M_j, \mathcal{T}, \mathcal{E}v \xrightarrow{1} E, \sigma, a_j, \mathcal{T}, \mathcal{E}v$ and $x[a_1, \dots, a_m] \in \text{Dom}(E)$. Hence, the test $x[a_1, \dots, a_m] \in \text{Dom}(E)$ succeeds.

2.4.4 Typing

In this section, we show that our type system is compatible with the semantics of the calculus, that is, we define a notion of typing for semantic configurations and show that typing is preserved by reduction (subject reduction). Finally, the property that semantic configurations are well-typed shows that certain conditions in the semantics always hold.

We use the following definitions:

- $\mathcal{E} \vdash E$ if and only if $E(x[a_1, \dots, a_m]) = a$ implies $\mathcal{E}(x) = T_1 \times \dots \times T_m \rightarrow T$ with for all $j \leq m$, $a_j \in T_j$ and $a \in T$.
- We define $\mathcal{E} \vdash P$, $\mathcal{E} \vdash Q$, and $\mathcal{E} \vdash M : T$ as in Section 2.3, with the additional rule $\mathcal{E} \vdash a : T$ if and only if $a \in T$. (This rule is useful to type evaluated terms.)
- $\mathcal{E} \vdash (\sigma, P)$ if and only if $\mathcal{E}[i_1 \mapsto [1, n_1], \dots, i_m \mapsto [1, n_m]] \vdash P$ and for all $j \leq m$, $\sigma(i_j) \in [1, n_j]$ for some n_1, \dots, n_m , where $\text{Dom}(\sigma) = \{i_1, \dots, i_m\}$. The judgments $\mathcal{E} \vdash (\sigma, Q)$ and $\mathcal{E} \vdash (\sigma, M) : T$ are defined in the same way.
- $\mathcal{E} \vdash \mathcal{T}$ if and only if $Tbl(a_1, \dots, a_m) \in \mathcal{T}$ implies $Tbl : T_1 \times \dots \times T_m$ with for all $j \leq m$, $a_j \in T_j$.
- $\mathcal{E} \vdash \mathcal{E}v$ if and only if $e(a_1, \dots, a_m) \in \mathcal{E}v$ implies $e : T_1 \times \dots \times T_m$ with for all $j \leq m$, $a_j \in T_j$.
- $\mathcal{E} \vdash E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$ if and only if $\mathcal{E} \vdash E$, $\mathcal{E} \vdash (\sigma, P)$, $\mathcal{E} \vdash \mathcal{T}$, $\mathcal{E} \vdash \mathcal{E}v$, and for all $(\sigma', Q) \in \mathcal{Q}$, $\mathcal{E} \vdash (\sigma', Q)$.
- $\mathcal{E} \vdash E, \mathcal{Q}, \mathcal{C}$ if and only if $\mathcal{E} \vdash E$ and for all $(\sigma', Q) \in \mathcal{Q}$, $\mathcal{E} \vdash (\sigma', Q)$.
- $\mathcal{E} \vdash E, \sigma, M : T, \mathcal{T}, \mathcal{E}v$ if and only if $\mathcal{E} \vdash E$, $\mathcal{E} \vdash (\sigma, M) : T$, $\mathcal{E} \vdash \mathcal{T}$, and $\mathcal{E} \vdash \mathcal{E}v$.

Lemma 7 If $\mathcal{E} \vdash E, \sigma, M : T, \mathcal{T}, \mathcal{E}v$ and $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mathcal{E}v'$, then $\mathcal{E} \vdash E', \sigma', M' : T, \mathcal{T}', \mathcal{E}v'$.

So, $\mathcal{E} \vdash E, \sigma, M : T, \mathcal{T}, \mathcal{E}v$ and $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t^* E', \sigma', a, \mathcal{T}', \mathcal{E}v'$, then $\mathcal{E} \vdash E', \sigma', a : T, \mathcal{T}', \mathcal{E}v'$.

Proof sketch By induction on the derivation of $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mathcal{E}v'$. \square

Lemma 8 If $\mathcal{E} \vdash E, \mathcal{Q}, \mathcal{C}$ and $E, \mathcal{Q}, \mathcal{C} \rightsquigarrow E', \mathcal{Q}', \mathcal{C}'$, then $\mathcal{E} \vdash E', \mathcal{Q}', \mathcal{C}'$.

So, if $\mathcal{E} \vdash E, \mathcal{Q}, \mathcal{C}$, then $\mathcal{E} \vdash \text{reduce}(E, \mathcal{Q}, \mathcal{C})$.

Proof sketch By cases on the derivation of $E, \mathcal{Q}, \mathcal{C} \rightsquigarrow E', \mathcal{Q}', \mathcal{C}'$. In the case of the replication, we have $\mathcal{E} \vdash (\sigma, !^{i \leq n} Q)$, so $\mathcal{E}[i_1 \mapsto [1, n_1], \dots, i_m \mapsto [1, n_m]] \vdash !^{i \leq n} Q$ and for all $j \leq m$, $\sigma(i_j) \in [1, n_j]$ for some n_1, \dots, n_m , where $\text{Dom}(\sigma) = \{i_1, \dots, i_m\}$. By (TRepl), $\mathcal{E}[i_1 \mapsto [1, n_1], \dots, i_m \mapsto [1, n_m], i \mapsto [1, n]] \vdash Q$, so $\mathcal{E} \vdash (\sigma[i \mapsto a], Q)$ for $a \in [1, n]$. In the case of the input, we use Lemma 7. \square

Lemma 9 If $\mathcal{E} \vdash Q_0$, then $\mathcal{E} \vdash \text{initConfig}(Q_0)$.

Proof sketch By Lemma 8 and the previous definitions. \square

Lemma 10 (Subject reduction) If $\mathcal{E} \vdash E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$ and $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$, then $\mathcal{E} \vdash E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$.

Proof sketch By cases on the derivation of $E, P, \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$, using Lemmas 7 and 8. \square

Moreover, if the rules that define $E, \sigma, M, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mathcal{E}v'$ (resp. $E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v \xrightarrow{p}_t E', (\sigma', P'), \mathcal{Q}', \mathcal{C}', \mathcal{T}', \mathcal{E}v'$) require as assumption $E'', \sigma'', M'', \mathcal{T}'', \mathcal{E}v'' \xrightarrow{p}_t \dots$ and the initial configuration is well-typed $\mathcal{E} \vdash E, \sigma, M : T, \mathcal{T}, \mathcal{E}v$ (resp. $\mathcal{E} \vdash E, (\sigma, P), \mathcal{Q}, \mathcal{C}, \mathcal{T}, \mathcal{E}v$) then so is the initial configuration of the assumption, that is, there exists T'' such that $\mathcal{E} \vdash E'', \sigma'', M'' : T'', \mathcal{T}'', \mathcal{E}v''$.

As an immediate consequence of Lemmas 9, 10, and 7 and the observation above, we obtain: if Q_0 satisfies Invariant 7, then in traces of Q_0 , the tests $a \in T$ in rules (LetT) and (Let) and $\forall j \leq m, a_j \in T_j$ in rule (Fun) always succeed. Moreover, in rules (NewT) and (New), we always have that T is *fixed*, *bounded*, or *nonuniform*. In the rules for *find*, we have $r_k \in \{\text{false}, \text{true}\}$ when r_k is a value (not an event). In the rules (InsertT), (GetTE), (GetT1), (GetT2), (Insert), (GetE), (Get1), and (Get2), we have $a_j \in T_j$ for $j \leq l$, where $Tbl : T_1 \times \dots \times T_l$. In the rules (EventT) and (Event), we have $a_j \in T_j$ for $j \leq l$, where $e : T_1 \times \dots \times T_l$.

2.5 Subset for the Initial Game

What is described in this section is not implemented yet, it is a project. Currently, process macros with arguments are not implemented. CryptoVerif additionally requires that variables in V and in defined conditions have a single definition.

The variables are always defined with the current replication indices \tilde{i} , so we omit them, writing x for $x[\tilde{i}]$; they are implicitly added by CryptoVerif. When a variable is used with the current replication indices, we can also omit the indices.

Along similar lines, the channels c are used without indices, and the current replication indices are implicitly added by CryptoVerif. This allows the adversary the select to which copy of processes it sends messages. The construct `newChannel` cannot occur in games manipulated by CryptoVerif. It is used only inside proofs. The grammar of the resulting calculus is summarized in Figure 8.

$M, N ::=$ i $x[M_1, \dots, M_m]$ $f(M_1, \dots, M_m)$ $\text{new } x : T; N$ $\text{let } p = M \text{ in } N \text{ else } N'$ $\text{let } x : T = M \text{ in } N$ $\text{if defined}(M_1, \dots, M_l) \wedge M \text{ then } N \text{ else } N'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j \text{ then } N_j) \text{ else } N'$ $\text{insert } Tbl(M_1, \dots, M_l); N$ $\text{get } Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } N \text{ else } N'$ $\text{event } e(M_1, \dots, M_l); N$ $\text{event_abort } e$	terms replication index variable access function application random number assignment (pattern-matching) assignment conditional array lookup insert in table get from table event event e and abort
$p ::=$ $x : T$ $f(p_1, \dots, p_m)$ $=M$	pattern variable function application comparison with a term
$Q ::=$ 0 $Q \mid Q'$ $!^{i \leq n} Q$ $c(p); P$	input process nil parallel composition replication n times input
$P ::=$ $\bar{c}\langle N \rangle; Q$ $\text{new } x : T; P$ $\text{let } p = M \text{ in } P \text{ else } P'$ $\text{if defined}(M_1, \dots, M_l) \wedge M \text{ then } P \text{ else } P'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j \text{ then } P_j) \text{ else } P$ $\text{insert } Tbl(M_1, \dots, M_l); P$ $\text{get } Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } P \text{ else } P'$ $\text{event } e(M_1, \dots, M_l); P$ $\text{event_abort } e$ yield	output process output random number assignment conditional array lookup insert in table get from table event event e and abort end

Figure 8: Subset of the calculus for the initial game

We recommend using the constructs `get` and `insert` to manage key tables, instead of `find` or `if` with defined conditions. When no `find` nor `if` with defined conditions occurs in the game, by Invariant 2, all accesses to variable x are of the form $x[\tilde{i}]$ where \tilde{i} are the current replication indices at the definition of x . Such accesses are simply abbreviated as x . Variables can then be considered as ordinary variables instead of arrays, since we only access the array cell at the current replication indices. This choice has several other advantages:

- Tables with `get/insert` are closer to lists usually used by cryptographers than `find`, they should be easier to understand for the user.
- Tables are supported by the symbolic protocol verifier ProVerif while `find` is not. Similarly, ProVerif does not support channels with indices. So avoiding `find` and channels with indices allows us to have a language compatible with ProVerif.
- Our compiler that translates CryptoVerif specifications into OCaml implementations [6] does not support `find`, because tables with `get/insert` are also much easier to implement than `find`.

We can define processes by macros: `let pid(x1 : T1, ..., xm : Tm) = P` or `let qid(x1 : T1, ..., xm : Tm) = Q`. If a process `pid(M1, ..., Mm)` occurs in the initial game, CryptoVerif verifies that M_1, \dots, M_m are of types T_1, \dots, T_m respectively, and replaces `pid(M1, ..., Mm)` with the expansion $P\{M_1/x_1, \dots, M_m/x_m\}$.

We can also define functions by macros: `letfun f(x1 : T1, ..., xm : Tm) = M`. If a term `f(M1, ..., Mm)` occurs in the initial game, CryptoVerif verifies that M_1, \dots, M_m are of types T_1, \dots, T_m respectively, and replaces `f(M1, ..., Mm)` with the expansion $M\{M_1/x_1, \dots, M_m/x_m\}$.

In the initial game, all bound variables that do not occur in V nor in defined conditions of `find` or `if` are renamed to distinct names, so that Invariant 1 is satisfied for these variables. The condition on input channels in Invariant 5 is always satisfied by definition of the language. CryptoVerif checks the rest of the invariants.

2.6 Subsets used inside the Sequence of Games

During the computation of the sequence of games, several properties are used by CryptoVerif, either required by some game transformations or guaranteed by others. We summarize them in this section.

Property 1 No function takes as argument or returns values of interval types. The types of the elements of tables, of the arguments of events, of values chosen by `new x[\tilde{i}] : T` are not interval types. The type T of the sent message in the (TOut) rule and of the receiving pattern in the (TIn) rule are not interval types.

This property is satisfied by all games manipulated by CryptoVerif, but not by processes that model the adversary. Combined with Invariants 7, 2, and 5, it implies that the terms of variable accesses $x[M_1, \dots, M_m]$ contain only replication indices and variables.

For processes that model security assumptions on primitives, the receiving variable can be of an interval type. (This is used for instance to specify the computational Diffie-Hellman assumption.)

Property 2 The `newChannel c` construct does not appear in games.

Property 3 The indices of channels are always the current replication indices.

These properties are also satisfied by all games manipulated by CryptoVerif, but not by processes that model the adversary.

$M, N ::=$ i $x[M_1, \dots, M_m]$ $f(M_1, \dots, M_m)$	terms replication index variable access function application
$FC ::=$ M $\text{let } p = M \text{ in } FC \text{ else } FC'$ $\text{let } x[\tilde{i}] : T = M \text{ in } N$ $\text{if defined}(M_1, \dots, M_l) \wedge M \text{ then } FC \text{ else } FC'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge FC'_j \text{ then } FC_j \text{ else } FC''$	find condition term assignment (pattern-matching) assignment conditional conditional array lookup
$p ::=$ $x[\tilde{i}] : T$ $f(p_1, \dots, p_m)$ $=M$	pattern variable function application comparison with a term
$Q ::=$ 0 $Q \mid Q'$ $!^{i \leq n} Q$ $c[M_1, \dots, M_l](p); P$	input process nil parallel composition replication n times input
$P ::=$ $\overline{c[M_1, \dots, M_l]} \langle N \rangle; Q$ $\text{new } x[\tilde{i}] : T; P$ $\text{let } p = M \text{ in } P \text{ else } P'$ $\text{if defined}(M_1, \dots, M_l) \wedge M \text{ then } P \text{ else } P'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge FC'_j \text{ then } P_j \text{ else } P$ $\text{event } e(M_1, \dots, M_l); P$ $\text{event_abort } e$ yield	output process output random number assignment conditional conditional array lookup array lookup event event e and abort end

Figure 9: Subset after game expansion

Property 4 The constructs `insert` and `get` do not occur in the game.

This property is not valid in the initial game, but it is in all other games of the sequence produced by CryptoVerif. The very first game transformation applied by CryptoVerif, **ExpandTables**, encodes `insert` and `get` using `find`. The constructs `insert` and `get` are never introduced by subsequent game transformations, so this property remains valid in the rest of the sequence.

Property 5 The variables defined in conditions of `find` have pairwise distinct names.

This property is enforced by the transformation **AutoSARename** by renaming variables defined in conditions of `find` to distinct names. (This is easy since these variables do not have array accesses by Invariant 3.) Property 5 is required as a precondition by many game transformations, and may be broken by game transformations that duplicate code. Therefore, we apply **AutoSARename** after these game transformations.

Property 6 The terms M contain only replication indices, variables, and function applications except for conditions of `find`.

The grammar of the language taking into account this property as well as Properties 2 and 4 is shown in Figure 9. By Invariant 4, `new` and `event` do not occur in conditions of `find`, so `new` and `event` never occur as terms. Property 6 is enforced by the transformation **Expand** by converting other terms into processes. This transformation is applied on the initial game after **ExpandTables**. Property 6 is broken by the cryptographic transformation that relies on assumptions on primitives, so **Expand** is called again after this transformation. Most game transformations require Property 6 as a precondition.

2.7 Security Properties, Indistinguishability

A context is a process containing a hole $[\]$. An evaluation context C is a context built from $[\]$, `newChannel` $c; C$, $Q \mid C$, and $C \mid Q$. We use an evaluation context to represent the adversary. We denote by $C[Q]$ the process obtained by replacing the hole $[\]$ in the context C with the process Q .

Definition 4 (Indistinguishability) Let Q and Q' be two processes and V a set of variables. Assume that Q and Q' satisfy Invariants 1 to 7 with public variables V , and the variables of V are defined in Q and Q' , with the same types.

An evaluation context C is said to be *acceptable* for Q with public variables V if and only if $\text{var}(C) \cap \text{var}(Q) \subseteq V$, $\text{vardef}(C) \cap V = \emptyset$, C and Q do not use any common table, and $C[Q]$ satisfies Invariants 1 to 7 with public variables V .

We write $Q \approx_p^V Q'$ when, for all evaluation contexts C acceptable for Q and Q' with public variables V and all distinguishers D that run in time at most t_D , $|\Pr[C[Q] : D] - \Pr[C[Q'] : D]| \leq p(C, t_D)$.

This definition formalizes that the probability that algorithms C and D distinguish the games Q and Q' is at most $p(C, t_D)$. The probability p typically depends on the runtime of C and D , but may also depend on other parameters, such as the number of queries to each oracle made by C . That is why p takes as arguments the whole algorithm C and the runtime of D .

The unusual requirement on variables of C comes from the presence of arrays and of the associated `find` construct which gives C direct access to variables of Q and Q' : the context C is allowed to access variables of Q and Q' only when they are in V . (In more standard settings, the calculus does not have constructs that allow the context to access variables of Q and Q' .) When V is empty, we omit it and write $Q \approx_p Q'$.

When C is acceptable for Q with public variables V , and we transform Q into Q' , we can rename the fresh variables of Q' (introduced by the game transformation) so that they do not

occur in C . Then C is also acceptable for Q' with public variables V . (To establish this property, we use that the variables of V are defined in Q and Q' , with the same types, so that, if $C[Q]$ is well-typed, then so is $C[Q']$.)

When C is acceptable for Q with public variables V , we have that $\text{vardef}(C) \cap \text{var}(Q) = \emptyset$, because $\text{vardef}(C) \cap \text{var}(Q) = \text{vardef}(C) \cap \text{var}(C) \cap \text{var}(Q) \subseteq \text{vardef}(C) \cap V = \emptyset$.

The following lemma is a straightforward consequence of Definition 4:

Lemma 11 1. *Reflexivity:* $Q \approx_0^V Q$.

2. *Symmetry:* If $Q \approx_p^V Q'$, then $Q' \approx_p^V Q$.

3. *Transitivity:* If $Q \approx_p^V Q'$ and $Q' \approx_{p'}^V Q''$, then $Q \approx_{p+p'}^V Q''$.

4. If $Q \approx_p^V Q'$ and C is an evaluation context acceptable for Q and Q' with public variables V , then $C[Q] \approx_{p'}^{V'} C[Q']$, where $p'(C, t_D) = p(C'[C[]], t_D)$ and $V' \subseteq V \cup \text{var}(C)$.

Definition 5 (Indistinguishability with introduction of events) Let Q and Q' be two processes and V a set of variables. Assume that Q and Q' satisfy Invariants 1 to 7 with public variables V , and the variables of V are defined in Q and Q' , with the same types.

Let $\text{NonUnique}_Q = \bigvee \{e \mid \text{find}[\text{unique}_e] \text{ occurs in } Q\}$.

We write $Q, \text{EvUsed} \xrightarrow{V}_{p, D^+} Q', \text{EvUsed}'$ when $D^+ = e_1 \vee \dots \vee e_m$ where e_1, \dots, e_m are Shoup events, the events that occur in Q are in EvUsed , $\text{EvUsed} \subseteq \text{EvUsed}'$, the events e_1, \dots, e_m and the events that occur in Q' but not in Q are in EvUsed' but not in EvUsed , and, for all evaluation contexts C acceptable for Q and Q' with public variables V that do not contain non-unique events that occur in Q nor events in $\text{EvUsed}' \setminus \text{EvUsed}$, and all distinguishers D that run in time at most t_D ,

$$\begin{aligned} \Pr[C[Q] : D \wedge \neg \text{NonUnique}_Q] &\leq \Pr[C[Q'] : (D \vee D^+) \wedge \neg \text{NonUnique}_{Q'}] + p(C, t_D) \\ \Pr[C[Q] : D \vee \text{NonUnique}_Q] &\geq \Pr[C[Q'] : D \vee \text{NonUnique}_{Q'}] - p(C, t_D). \end{aligned}$$

Intuitively, the events EvUsed are those used by CryptoVerif in the sequence of games until the game Q included, while the events EvUsed' are those used until Q' . Hence, EvUsed contains the events that occur in Q ; EvUsed' contains EvUsed and the events that occur in Q' . The formula D^+ uses fresh Shoup events introduced during the transformation of Q into Q' ; hence these events are in EvUsed' but not in EvUsed . More generally, all events that occur in Q' but not in Q are fresh events introduced in the transformation of Q into Q' , so they are in EvUsed' but not in EvUsed .

We need to introduced distinct events for $\text{find}[\text{unique}_e]$ because we need to distinguish the non-unique events that occur in Q from those that occur in the context C . The context C must not contain the fresh events introduced in the transformation of Q into Q' (which can be guaranteed by choosing these fresh events appropriately).

Lemma 12 Let $D_{\text{false}}(\mathcal{E}v, a) = \text{false}$ for all $\mathcal{E}v, a$.

1. *Link with indistinguishability:* Suppose that Q and Q' do not contain non-unique events, the events that occur in Q are in EvUsed , and the events that occur in Q' also occur in Q . Then $Q, \text{EvUsed} \xrightarrow{V}_{p, D_{\text{false}}} Q', \text{EvUsed}$ if and only if $Q \approx_p^V Q'$.

2. *Reflexivity:* if the events that occur in Q are in EvUsed , then $Q, \text{EvUsed} \xrightarrow{V}_{0, D_{\text{false}}} Q, \text{EvUsed}$.

3. *Transitivity:* If $Q, \text{EvUsed} \xrightarrow{V}_{p, D_1^+} Q', \text{EvUsed}'$ and $Q', \text{EvUsed}' \xrightarrow{V}_{p', D_2^+} Q'', \text{EvUsed}''$, then $Q, \text{EvUsed} \xrightarrow{V}_{p'', D_1^+ \vee D_2^+} Q'', \text{EvUsed}''$, where $p''(C, t_D) = p(C, t_D) + p'(C, t_D + t_{D_1^+})$.

4. If $Q, EvUsed \xrightarrow{V}_{p, D^+} Q', EvUsed'$ and C is a context acceptable for Q and Q' with public variables V such that C does not contain the non-unique events in Q and the events that occur in C are in $EvUsed$, then $C[Q], EvUsed \xrightarrow{V'}_{p', D^+} C[Q'], EvUsed'$, where $p'(C', t_D) = p(C'[C[\cdot]], t_D)$ and $V' \subseteq V \cup \text{var}(C)$.
5. Renaming: if $Q, EvUsed \xrightarrow{V}_{p, D^+} Q', EvUsed'$ and $EvUsed^+$ is a set of events disjoint from $EvUsed$, then $Q, EvUsed \cup EvUsed^+ \xrightarrow{V}_{p', \sigma D^+} \sigma Q', \sigma EvUsed' \cup EvUsed^+$ where σ is a renaming of the events in $EvUsed' \setminus EvUsed$ to events not in $EvUsed \cup EvUsed^+$, $p'(C, t_D) = p(\sigma^{-1}C, t_D)$, and the distinguisher $D^+ \circ \sigma^{-1}$ is defined by $(D^+ \circ \sigma^{-1})(\mathcal{E}v, a) = D^+(\sigma^{-1}\mathcal{E}v, a)$.

In the last property, the formula that expresses the probability p is often independent of the names of events; in this case, we have $p' = p$. If $D^+ = e_1 \vee \dots \vee e_m$, then $D^+ \circ \sigma^{-1} = \sigma e_1 \vee \dots \vee \sigma e_m$.

Proof Property 1: Since Q and Q' do not contain non-unique events, NonUnique_Q and $\text{NonUnique}_{Q'}$ are always false, so given the hypothesis on $EvUsed$, $Q, EvUsed \xrightarrow{V}_{p, D^{\text{false}}} Q', EvUsed$ reduces to: for all evaluation contexts C acceptable for Q and Q' with public variables V and all distinguishers D that run in time at most t_D ,

$$\begin{aligned} \Pr[C[Q] : D] &\leq \Pr[C[Q'] : D] + p(C, t_D) \\ \Pr[C[Q] : D] &\geq \Pr[C[Q'] : D] - p(C, t_D). \end{aligned}$$

which is exactly $Q \approx_p^V Q'$.

Property 2: Obvious.

Property 3: The events in Q are in $EvUsed$. We have $EvUsed \subseteq EvUsed' \subseteq EvUsed''$. The distinguisher $D_1^+ \vee D_2^+$ is a disjunction of Shoup events that occur in D_1^+ or in D_2^+ ; in the former case, they are in they are in $EvUsed' \setminus EvUsed$; in the latter case, they are in they are in $EvUsed'' \setminus EvUsed'$; so in both cases they are in $EvUsed'' \setminus EvUsed$. The events that occur in Q'' but not in Q occur either in Q'' but not in Q' or in Q' but not in Q ; in the former case, they are in $EvUsed'' \setminus EvUsed'$; in the latter case, they are in $EvUsed' \setminus EvUsed$; so in both cases they are in $EvUsed'' \setminus EvUsed$.

Let C be any evaluation context acceptable for Q and Q'' with public variables V that does not contain non-unique events that occur in Q nor events in $EvUsed'' \setminus EvUsed$. After renaming the fresh variables of Q' that do not occur in Q and Q'' and the tables of Q' that do not occur in Q and Q'' so that they do not occur in C , C is also acceptable for Q' with public variables V . Since $EvUsed \subseteq EvUsed' \subseteq EvUsed''$, C does not contain events in $EvUsed'' \setminus EvUsed'$ nor in $EvUsed' \setminus EvUsed$. Moreover, the non-unique events in Q' either occur in Q , or they occur in Q' but not Q , so they are in $EvUsed' \setminus EvUsed$, so in both cases, they do not occur in C . Let D be any distinguisher D that runs in time at most t_D .

Then we have:

$$\begin{aligned} &\Pr[C[Q] : D \wedge \neg \text{NonUnique}_Q] \\ &\leq \Pr[C[Q'] : (D \vee D_1^+) \wedge \neg \text{NonUnique}_{Q'}] + p(C, t_D) \quad \text{since } Q, EvUsed \xrightarrow{V}_{p, D_1^+} Q', EvUsed' \\ &\leq \Pr[C[Q''] : ((D \vee D_1^+) \vee D_2^+) \wedge \neg \text{NonUnique}_{Q''}] + p'(C, t_{D \vee D_1^+}) + p(C, t_D) \\ &\quad \text{since } Q', EvUsed' \xrightarrow{V}_{p', D_2^+} Q'', EvUsed'' \\ &\leq \Pr[C[Q''] : (D \vee (D_1^+ \vee D_2^+)) \wedge \neg \text{NonUnique}_{Q''}] + p(C, t_D) + p'(C, t_D + t_{D_1^+}) \\ &\quad \text{since } t_{D \vee D_1^+} \leq t_D + t_{D_1^+} \\ &\leq \Pr[C[Q''] : (D \vee (D_1^+ \vee D_2^+)) \wedge \neg \text{NonUnique}_{Q''}] + p''(C, t_D) \quad \text{by definition of } p'' \end{aligned}$$

Similarly:

$$\begin{aligned}
& \Pr[C[Q] : D \vee \text{NonUnique}_Q] \\
& \geq \Pr[C[Q'] : D \vee \text{NonUnique}_{Q'}] - p(C, t_D) && \text{since } Q, \text{EvUsed} \xrightarrow{V}_{p, D_1^+} Q', \text{EvUsed}' \\
& \geq \Pr[C[Q''] : D \vee \text{NonUnique}_{Q''}] - p'(C, t_D) - p(C, t_D) && \text{since } Q', \text{EvUsed}' \xrightarrow{V}_{p', D_2^+} Q'', \text{EvUsed}'' \\
& \geq \Pr[C[Q''] : D \vee \text{NonUnique}_{Q''}] - p''(C, t_D) && \text{by definition of } p'', \text{ since } t_D \leq t_D + t_{D_1^+}
\end{aligned}$$

Therefore, we have $Q, \text{EvUsed} \xrightarrow{V}_{p'', D_1^+ \vee D_2^+} Q'', \text{EvUsed}''$.

Property 4: The events that occur in $C[Q]$ are either in C or in Q ; the former case, they are in EvUsed by hypothesis; in the latter case, they are also in EvUsed since $Q, \text{EvUsed} \xrightarrow{V}_{p, D^+} Q', \text{EvUsed}'$. We have $\text{EvUsed} \subseteq \text{EvUsed}'$ since $Q, \text{EvUsed} \xrightarrow{V}_{p, D^+} Q', \text{EvUsed}'$. The events that occur in $C[Q']$ but not in $C[Q]$ are in Q' but not in Q , so they are in $\text{EvUsed}' \setminus \text{EvUsed}$. The distinguisher D^+ is a disjunction of Shoup events that are in $\text{EvUsed}' \setminus \text{EvUsed}$, since $Q, \text{EvUsed} \xrightarrow{V}_{p, D^+} Q', \text{EvUsed}'$.

Let C' be any evaluation context acceptable for $C[Q]$ and $C[Q']$ with public variables V' that does not contain non-unique events that occur in $C[Q]$ nor events in $\text{EvUsed}' \setminus \text{EvUsed}$. We rename the variables of C' not in V' so that they are not in V ; this renaming does not change the probabilities. Then $C'[C[]]$ is an evaluation context acceptable for Q and Q' with public variables V . Indeed,

$$\begin{aligned}
\text{var}(C'[C[]]) \cap \text{var}(Q) &= (\text{var}(C') \cup \text{var}(C)) \cap \text{var}(Q) \\
&\subseteq (V' \cup \text{var}(C)) \cap \text{var}(Q) \\
&\quad \text{since } \text{var}(C') \cap \text{var}(Q) \subseteq \text{var}(C') \cap \text{var}(C[Q]) \subseteq V' \\
&\subseteq (V \cup \text{var}(C)) \cap \text{var}(Q) && \text{since } V' \subseteq V \cup \text{var}(C) \\
&\subseteq V && \text{since } \text{var}(C) \cap \text{var}(Q) \subseteq V
\end{aligned}$$

We have similarly $\text{var}(C'[C[]]) \cap \text{var}(Q') \subseteq V$. We also have $\text{vardef}(C'[C[]]) \cap V = (\text{vardef}(C') \cap V) \cup (\text{vardef}(C) \cap V) = \emptyset$ since $\text{vardef}(C) \cap V = \emptyset$ because C is an acceptable evaluation context for Q with public variables V and $\text{vardef}(C') \cap V \subseteq \text{vardef}(C') \cap V' = \emptyset$ because we have renamed the variables of C' not in V' so that they are not in V and C' is an acceptable evaluation context for $C[Q]$ and with public variables V' . Moreover, C and Q do not use any common table, and C' and $C[Q]$ do not use any common table so a fortiori C' and Q do not use any common table. Therefore, $C'[C[]]$ and Q do not use any common table. Similarly, $C'[C[]]$ and Q' do not use any common table. $C'[C[]]$ does not contain non-unique events that occur in Q , since C' does not contain non-unique events that occur in $C[Q]$ so a fortiori in Q , and C does not contain non-unique events that occur in Q . $C'[C[]]$ does not contain events in $\text{EvUsed}' \setminus \text{EvUsed}$, since C' and C do not. By using the property $Q, \text{EvUsed} \xrightarrow{V}_{p, D^+} Q', \text{EvUsed}'$ with the context $C'[C[]]$, we get immediately the desired probability bounds.

Property 5: The events that occur in Q are in EvUsed , so a fortiori in $\text{EvUsed} \cup \text{EvUsed}^+$. We have $\text{EvUsed} \cup \text{EvUsed}^+ \subseteq \sigma \text{EvUsed}' \cup \text{EvUsed}^+$, since σ leaves the events of EvUsed unchanged, so $\text{EvUsed} = \sigma \text{EvUsed} \subseteq \sigma \text{EvUsed}'$. The distinguisher $D^+ \circ \sigma^{-1}$ is a disjunction of Shoup events that are in $\sigma(\text{EvUsed}' \setminus \text{EvUsed}) = \sigma \text{EvUsed}' \setminus \text{EvUsed} = (\sigma \text{EvUsed}' \cup \text{EvUsed}^+) \setminus (\text{EvUsed} \cup \text{EvUsed}^+)$. The events that occur in $\sigma Q'$ but not in $Q = \sigma Q$ are also in $\sigma(\text{EvUsed}' \setminus \text{EvUsed}) = (\sigma \text{EvUsed}' \cup \text{EvUsed}^+) \setminus (\text{EvUsed} \cup \text{EvUsed}^+)$.

Let C be any evaluation context acceptable for Q and $\sigma Q'$ with public variables V that do not contain non-unique events that occur in Q nor events in $\sigma \text{EvUsed}' \setminus \text{EvUsed}$. Then $\sigma^{-1}C$ is an evaluation context acceptable for Q and Q' with public variables V that do not contain non-unique events that occur in Q nor events in $\text{EvUsed}' \setminus \text{EvUsed}$. Let D be a distinguisher

that runs in time at most t_D . Then $D \circ \sigma$ also runs in time at most t_D . By applying the property $Q, EvUsed \xrightarrow{V}_{p, D^+} Q', EvUsed'$ with the context $\sigma^{-1}C$ and the distinguisher $D \circ \sigma$, we obtain:

$$\begin{aligned} \Pr[(\sigma^{-1}C)[Q] : (D \circ \sigma) \wedge \neg \text{NonUnique}_Q] &\leq \Pr[(\sigma^{-1}C)[Q'] : ((D \circ \sigma) \vee D^+) \wedge \neg \text{NonUnique}_{Q'}] \\ &\quad + p(\sigma^{-1}C, t_D) \\ \Pr[(\sigma^{-1}C)[Q] : (D \circ \sigma) \vee \text{NonUnique}_Q] &\geq \Pr[\sigma^{-1}C[Q'] : (D \circ \sigma) \vee \text{NonUnique}_{Q'}] - p(\sigma^{-1}C, t_D). \end{aligned}$$

Since renaming events does not change the probabilities of traces, by applying σ , we get:

$$\begin{aligned} \Pr[C[Q] : D \wedge \neg \text{NonUnique}_Q] &\leq \Pr[C[\sigma Q'] : (D \vee (D^+ \circ \sigma^{-1})) \wedge \neg \text{NonUnique}_{\sigma Q'}] + p(\sigma^{-1}C, t_D) \\ \Pr[C[Q] : D \vee \text{NonUnique}_Q] &\geq \Pr[C[\sigma Q'] : D \vee \text{NonUnique}_{\sigma Q'}] - p(\sigma^{-1}C, t_D). \end{aligned}$$

which yields the desired result. \square

When CryptoVerif transforms a game G into a game G' , in most cases, we have $G \approx_p^V G'$, where p is the probability difference coming from the transformation, and computed by CryptoVerif. However, there are two exceptions to this situation:

- transformations that exploit the uniqueness of $\text{find}[\text{unique}_e]$, which are valid only when event e is not executed. These events are taken into account by NonUnique_Q .
- transformations that insert events using Shoup's lemma.

That is why, in general, when CryptoVerif transforms a game G into a game G' , we have $G, EvUsed \xrightarrow{V}_{p, e_1 \vee \dots \vee e_m} G', EvUsed'$, where e_1, \dots, e_m are the events introduced by Shoup's lemma during the transformation of G into G' .

2.7.1 Secrecy

Let us now define the secrecy properties that are proved by CryptoVerif.

Definition 6 (One-session secrecy) Let C be an evaluation context acceptable for $Q \mid Q_x$ with public variables V ($x \notin V$) that does not contain S nor \bar{S} .

The advantage of the adversary C against the *one-session secrecy* of x in process Q is

$$\text{Adv}_Q^{1\text{-ses. secrecy}(x)}(C) = \Pr[C[Q \mid Q_x] : S] - \Pr[C[Q \mid Q_x] : \bar{S}]$$

where

$$\begin{aligned} Q_x &= c_{s0}(); \text{new } b : \text{bool}; \overline{c_{s0}}(); \\ &\quad (c_s(u_1 : [1, n_1], \dots, u_m : [1, n_m])); \text{if defined}(x[u_1, \dots, u_m]) \text{ then} \\ &\quad \text{if } b \text{ then } \overline{c_s}\langle x[u_1, \dots, u_m] \rangle \text{ else new } y : T; \overline{c_s}\langle y \rangle \\ &\quad | c'_s(b' : \text{bool}); \text{if } b = b' \text{ then event_abort } S \text{ else event_abort } \bar{S} \end{aligned}$$

$c_{s0}, c_s, c'_s \notin \text{fc}(Q)$, $u_1, \dots, u_m, y, b, b' \notin \text{var}(Q)$, S, \bar{S} do not occur in Q , and $\mathcal{E}(x) = [1, n_1] \times \dots \times [1, n_m] \rightarrow T$.

The process Q preserves the *one-session secrecy* of x with public variables V ($x \notin V$) up to probability p when, for all evaluation contexts C acceptable for $Q \mid Q_x$ with public variables V that do not contain S nor \bar{S} , $\text{Adv}_Q^{1\text{-ses. secrecy}(x)}(C) \leq p(C)$.

Intuitively, the adversary cannot guess the random bit b , that is, it cannot distinguish whether the process outputs the value of the secret ($b = \text{true}$) or outputs a random number ($b = \text{false}$). The adversary performs a single test query, modeled by Q_x . In more detail, in Q_x , we choose

a random bit b ; the adversary sends the indices (u_1, \dots, u_m) on channel c_s to perform a test query on $x[u_1, \dots, u_m]$: if $b = \text{true}$, the test query sends back $x[u_1, \dots, u_m]$; if $b = \text{false}$, it sends back a random value y . Finally, the adversary should guess the bit b : it sends its guess b' on channel c'_s and, if the guess is correct, then event S is executed, and otherwise, event \bar{S} is executed. The probability of getting some information on the secret is the difference between the probability of S and the probability of \bar{S} . (When the adversary always sends a guess on channel c'_s , we have $\Pr[C[Q \mid Q_x] : \bar{S}] = 1 - \Pr[C[Q \mid Q_x] : S]$, so the advantage of the adversary is $\text{Adv}_Q^{1-\text{ses.secretcy}(x)}(C) = \Pr[C[Q \mid Q_x] : S] - \Pr[C[Q \mid Q_x] : \bar{S}] = 2\Pr[C[Q \mid Q_x] : S] - 1$, which is a more standard formula. By flipping a coin, the adversary can execute events S and \bar{S} with the same probability, that is why the probability that the adversary really guesses b is the difference between the probability of these two events. We need not take the absolute value of $\Pr[C[Q \mid Q_x] : S] - \Pr[C[Q \mid Q_x] : \bar{S}]$ because, when it is negative, we can obtain the opposite, positive value by considering an adversary that sends the guess $1 - b'$ instead of b' .)

Definition 7 (Secrecy) Let C be an evaluation context acceptable for $Q \mid R_x$ with public variables V ($x \notin V$) that does not contain S nor \bar{S} .

The advantage of the adversary C against the *secrecy* of x in process Q is

$$\text{Adv}_Q^{\text{Secrecy}(x)}(C) = \Pr[C[Q \mid R_x] : S] - \Pr[C[Q \mid R_x] : \bar{S}]$$

where

$$\begin{aligned} R_x &= c_{s0}(); \text{ new } b : \text{bool}; \overline{c_{s0}}(); \\ & \quad ({}^{i_s \leq n_s} c_s[i_s](u_1 : [1, n_1], \dots, u_m : [1, n_m])); \text{ if defined}(x[u_1, \dots, u_m]) \text{ then} \\ & \quad \text{if } b \text{ then } \overline{c_s[i_s]}(x[u_1, \dots, u_m]) \text{ else} \\ & \quad \text{find } u'_s = i'_s \leq n_s \text{ suchthat defined}(y[i'_s], u_1[i'_s], \dots, u_m[i'_s]) \wedge \\ & \quad \quad u_1[i'_s] = u_1 \wedge \dots \wedge u_m[i'_s] = u_m \\ & \quad \text{then } \overline{c_s[i_s]}(y[u'_s]) \\ & \quad \text{else new } y : T; \overline{c_s[i_s]}(y) \\ & \quad | c'_s(b' : \text{bool}); \text{ if } b = b' \text{ then event_abort } S \text{ else event_abort } \bar{S} \end{aligned}$$

$c_{s0}, c_s, c'_s \notin \text{fc}(Q)$, $u_1, \dots, u_m, u'_s, y, b, b' \notin \text{var}(Q)$, S, \bar{S} do not occur in Q , $\mathcal{E}(x) = [1, n_1] \times \dots \times [1, n_m] \rightarrow T$.

The process Q *preserves the secrecy of x* with public variables V ($x \notin V$) up to probability p when, for all evaluation contexts C acceptable for $Q \mid R_x$ with public variables V that do not contain S nor \bar{S} , $\text{Adv}_Q^{\text{Secrecy}(x)}(C) \leq p(C)$.

The replication bound n_s is chosen large enough so that it does not prevent communications that would otherwise occur.

Intuitively, the adversary cannot guess b , that is, it cannot distinguish whether the process outputs the value of the secret for several indices ($b = \text{true}$) or outputs independent random numbers ($b = \text{false}$). In this definition, the adversary can perform several test queries, modeled by R_x . This corresponds to the “real-or-random” definition of security [1]. (As shown in [1], this notion is stronger than the more standard approach in which the adversary can perform a single test query and some reveal queries, which always reveal $x[u_1, \dots, u_m]$.)

Lemma 13 *If Q preserves the secrecy of x with public variables V up to probability p and C is an acceptable evaluation context for Q with public variables V , then for all $V' \subseteq V \cup \text{var}(C)$, $C[Q]$ preserves the secrecy of x with public variables V' up to probability p' such that $p'(C') = p(C'[C])$.*

If $Q \approx_p^{V \cup \{x\}} Q'$ and Q preserves the secrecy of x with public variables V up to probability p' , then Q' preserves the secrecy of x with public variables V up to probability p'' such that $p''(C) = p'(C) + 2 \times p(C[[\] \mid R_x], t_S)$.

Proof Suppose that Q preserves the secrecy of x with public variables V ($x \notin V$) and C is an acceptable evaluation context for Q with public variables V . Let $V' \subseteq V \cup \text{var}(C)$. Choose channels c_{s0}, c_s, c'_s , variables $u_1, \dots, u_m, u'_s, y, b, b'$, and events S, \bar{S} such that they do not occur in $C[Q]$. Let C' be an acceptable evaluation context for $C[Q] \mid R_x$ with public variables V' that does not contain S nor \bar{S} . Then we have

$$\begin{aligned} \text{Adv}_{C[Q]}^{\text{Secrecy}(x)}(C') &= \Pr[C'[C[Q] \mid R_x] : S] - \Pr[C'[C[Q] \mid R_x] : \bar{S}] \\ &= \Pr[C'[C[Q \mid R_x]] : S] - \Pr[C'[C[Q \mid R_x]] : \bar{S}] \\ &\leq p(C'[C]) \end{aligned}$$

We can commute the context C with the parallel composition with R_x because the context C does not bind the channels of R_x . The context $C'[C]$ is an acceptable evaluation context for $Q \mid R_x$ with public variables V that does not contain S nor \bar{S} : there is no common table between C and Q , and between C' and $C[Q] \mid R_x$, so a fortiori between C' and Q and R_x does not use tables, so there is no common table between $C'[C]$ and $Q \mid R_x$; moreover

$$\begin{aligned} \text{var}(C'[C]) \cap \text{var}(Q \mid R_x) &= ((\text{var}(C') \cap \text{var}(Q \mid R_x)) \cup \text{var}(C)) \cap \text{var}(Q \mid R_x) \\ &\subseteq (V' \cup \text{var}(C)) \cap \text{var}(Q \mid R_x) \quad \text{since } \text{var}(C') \cap \text{var}(C[Q] \mid R_x) \subseteq V' \\ &\subseteq (V \cup \text{var}(C)) \cap \text{var}(Q \mid R_x) \quad \text{since } V' \subseteq V \cup \text{var}(C) \\ &\subseteq V \quad \text{since } \text{var}(C) \cap \text{var}(Q) \subseteq V \text{ and } \text{var}(C) \cap \text{var}(R_x) = \emptyset \end{aligned}$$

Suppose that $Q \approx_p^{V \cup \{x\}} Q'$ and Q preserves the secrecy of x with public variables V up to probability p' . Let C be an acceptable evaluation context for $Q' \mid R_x$ with public variables V that does not contain S nor \bar{S} .

$$\begin{aligned} \text{Adv}_{Q'}^{\text{Secrecy}(x)}(C) &= \Pr[C[Q' \mid R_x] : S] - \Pr[C[Q' \mid R_x] : \bar{S}] \\ &\leq \Pr[C[Q \mid R_x] : S] - \Pr[C[Q \mid R_x] : \bar{S}] + \\ &\quad |\Pr[C[Q' \mid R_x] : S] - \Pr[C[Q \mid R_x] : S]| + \\ &\quad |\Pr[C[Q \mid R_x] : \bar{S}] - \Pr[C[Q' \mid R_x] : \bar{S}]| \\ &\leq p'(C) + 2 \times p(C[[\] \mid R_x], t_S) \end{aligned}$$

since $t_S = t_{\bar{S}}$. Indeed, by renaming the variables and tables of Q that do not appear in Q' to variables and tables that also do not occur in C , C is also an acceptable evaluation context for $Q \mid R_x$ with public variables V . \square

2.7.2 Correspondences

In this section, we define non-injective and injective correspondences.

Non-injective Correspondences A non-injective correspondence is a property of the form “if some events have been executed, then some other events have been executed at least once”. Here, we generalize these correspondences to implications between logical formulae $\psi \Rightarrow \phi$, which may contain events. We use the following logical formulae:

$\phi ::=$	formula
M	term
$\text{event}(e(M_1, \dots, M_m))$	event
$\phi_1 \wedge \phi_2$	conjunction
$\phi_1 \vee \phi_2$	disjunction

Terms M, M_1, \dots, M_m in formulae must contain only variables x without array indices and function applications, and their variables are assumed to be distinct from variables of processes. The formula M holds when M evaluates to true. The formula $\text{event}(e(M_1, \dots, M_n))$ holds when the event $e(M_1, \dots, M_n)$ has been executed. The conjunction and disjunction are defined as usual. More formally, we write $\rho, \mathcal{E}v \vdash \phi$ when the sequence of events $\mathcal{E}v$ satisfies the formula ϕ , in the environment ρ that maps variables to bitstrings. We define $\rho, \mathcal{E}v \vdash \phi$ as follows:

$\rho, \mathcal{E}v \vdash M$ if and only if $\rho, \emptyset, M, \emptyset, \emptyset \xrightarrow{1^*} \rho, \emptyset, \text{true}, \emptyset, \emptyset$
 $\rho, \mathcal{E}v \vdash \text{event}(e(M_1, \dots, M_m))$ if and only if
for all $j \leq m$, $\rho, \emptyset, M_j, \emptyset, \emptyset \xrightarrow{1^*} \rho, \emptyset, a_j, \emptyset, \emptyset$ and $e(a_1, \dots, a_m) \in \mathcal{E}v$
 $\rho, \mathcal{E}v \vdash \phi_1 \wedge \phi_2$ if and only if $\rho, \mathcal{E}v \vdash \phi_1$ and $\rho, \mathcal{E}v \vdash \phi_2$
 $\rho, \mathcal{E}v \vdash \phi_1 \vee \phi_2$ if and only if $\rho, \mathcal{E}v \vdash \phi_1$ or $\rho, \mathcal{E}v \vdash \phi_2$

Formulae denoted by ψ are conjunctions of events.

Definition 8 The sequence of events $\mathcal{E}v$ satisfies the correspondence $\psi \Rightarrow \phi$, written $\mathcal{E}v \vdash \psi \Rightarrow \phi$, if and only if for all ρ defined on $\text{var}(\psi)$ such that $\rho, \mathcal{E}v \vdash \psi$, there exists an extension ρ' of ρ to $\text{var}(\phi)$ such that $\rho', \mathcal{E}v \vdash \phi$.

Intuitively, a sequence of events $\mathcal{E}v$ satisfies $\psi \Rightarrow \phi$ when, if $\mathcal{E}v$ satisfies ψ , then $\mathcal{E}v$ satisfies ϕ . The variables of ψ are universally quantified; those of ϕ that do not occur in ψ are existentially quantified.

Definition 9 We define a distinguisher $D(\mathcal{E}v, a) = \text{true}$ if and only if $\mathcal{E}v \vdash \psi \Rightarrow \phi$, and we denote this distinguisher D simply by $\psi \Rightarrow \phi$.

The advantage of the adversary C against the correspondence $\psi \Rightarrow \phi$ in process Q is $\text{Adv}_Q^{\psi \Rightarrow \phi}(C) = \Pr[C[Q] : \neg(\psi \Rightarrow \phi)]$, where C is an evaluation context acceptable for Q with any public variables that does not contain events used by $\psi \Rightarrow \phi$.

The process Q satisfies the correspondence $\psi \Rightarrow \phi$ with public variables V up to probability p if and only if for all evaluation contexts C acceptable for Q with public variables V that do not contain events used by $\psi \Rightarrow \phi$, $\text{Adv}_Q^{\psi \Rightarrow \phi}(C) \leq p(C)$.

A process satisfies $\psi \Rightarrow \phi$ up to probability p when the probability that it generates a sequence of events $\mathcal{E}v$ that does not satisfy $\psi \Rightarrow \phi$ is at most $p(C)$, in the presence of an adversary represented by the context C . The events used by $\psi \Rightarrow \phi$ are the events that occur in the formula $\psi \Rightarrow \phi$.

Example 2 The correspondence

$$\text{event}(e_B(x, y, z)) \Rightarrow \text{event}(e_A(x, y, z)) \quad (1)$$

means that, with overwhelming probability, for all x, y, z , if $e_B(x, y, z)$ has been executed, then $e_A(x, y, z)$ has been executed.

The correspondence

$$\begin{aligned} &\text{event}(e_1(x)) \wedge \text{event}(e_2(x)) \Rightarrow \\ &\text{event}(e_3(x)) \vee (\text{event}(e_4(x, y)) \wedge \text{event}(e_5(y, z))) \end{aligned}$$

means that, with overwhelming probability, for all x , if $e_1(x)$ and $e_2(x)$ have been executed, then $e_3(x)$ has been executed or there exists y such that both $e_4(x, y)$ and $e_5(x, y)$ have been executed.

Injective Correspondences Injective correspondences are properties of the form “if some event has been executed n times, then some other events have been executed at least n times”. In order to model them in our logical formulae, we extend the grammar of formulae ϕ with injective events $\text{inj-event}(e(M_1, \dots, M_m))$. The formula ψ is a conjunction of (injective or non-injective) events. The conditions on the number of executions of events apply only to injective events.

The definition of formula satisfaction is also extended: we indicate at which step each injective event has been executed, by a “pseudo-formula” ϕ^τ obtained from the formula ϕ by replacing terms and non-injective events with \perp and injective events with the step τ at which they have been executed (that is, their index τ in the sequence of events $\mathcal{E}v$) or \perp when their execution is not required. For example, if $\phi = \text{inj-event}(e_1(x)) \wedge (\text{inj-event}(e_2(x)) \vee \text{inj-event}(e_3(x)))$, then ϕ^τ is of the form $\tau_1 \wedge (\tau_2 \vee \tau_3)$ where τ_1 is the execution step of $e_1(x)$ and either τ_2 is the execution step of $e_2(x)$ or τ_3 is the execution step of $e_3(x)$. (One of the steps τ_2 and τ_3 may be \perp , but not both.) We define formula satisfaction $\rho, \mathcal{E}v \vdash^{\phi^\tau} \phi$ as follows:

$$\begin{aligned} \rho, \mathcal{E}v \vdash^\perp M & \text{ if and only if } \rho, \emptyset, M, \emptyset, \emptyset \xrightarrow{1^*} \rho, \emptyset, \text{true}, \emptyset, \emptyset \\ \rho, \mathcal{E}v \vdash^\perp \text{event}(e(M_1, \dots, M_m)) & \text{ if and only if} \\ & \text{for all } j \leq m, \rho, \emptyset, M_j, \emptyset, \emptyset \xrightarrow{1^*} \rho, \emptyset, a_j, \emptyset, \emptyset \text{ and } e(a_1, \dots, a_m) \in \mathcal{E}v \\ \rho, \mathcal{E}v \vdash^\tau \text{inj-event}(e(M_1, \dots, M_m)) & \text{ if and only if } \tau \neq \perp, \\ & \text{for all } j \leq m, \rho, \emptyset, M_j, \emptyset, \emptyset \xrightarrow{1^*} \rho, \emptyset, a_j, \emptyset, \emptyset, \text{ and } e(a_1, \dots, a_m) = \mathcal{E}v(\tau) \\ \rho, \mathcal{E}v \vdash^{\phi_1^\tau \wedge \phi_2^\tau} \phi_1 \wedge \phi_2 & \text{ if and only if } \rho, \mathcal{E}v \vdash^{\phi_1^\tau} \phi_1 \text{ and } \rho, \mathcal{E}v \vdash^{\phi_2^\tau} \phi_2 \\ \rho, \mathcal{E}v \vdash^{\phi_1^\tau \vee \phi_2^\tau} \phi_1 \vee \phi_2 & \text{ if and only if } \rho, \mathcal{E}v \vdash^{\phi_1^\tau} \phi_1 \text{ or } \rho, \mathcal{E}v \vdash^{\phi_2^\tau} \phi_2 \end{aligned}$$

This definition differs from the case of non-injective correspondences in that we propagate the pseudo-formula ϕ^τ and, in the case of injective events, we make sure that the event has been executed at step τ by requiring that $\tau \neq \perp$ and $e(a_1, \dots, a_m) = \mathcal{E}v(\tau)$.

Given a function \mathbb{F} that maps ψ^τ to ϕ^τ , the *projection* f of \mathbb{F} to the leaf at occurrence o of ϕ is such that $f(\psi^\tau)$ is the leaf at occurrence o of $\mathbb{F}(\psi^\tau)$. For example, if \mathbb{F} maps ψ^τ to ϕ^τ of the form $\tau_1 \wedge (\tau_2 \vee \tau_3)$, then \mathbb{F} has three projections, which map ψ^τ to τ_1 , τ_2 , and τ_3 respectively. We say that \mathbb{F} is *component-wise injective* when each projection f of \mathbb{F} is such that $f(\psi_1^\tau) = f(\psi_2^\tau) \neq \perp$ implies $\psi_1^\tau = \psi_2^\tau$. (Ignoring the result \perp , f is injective.)

Definition 10 The sequence of events $\mathcal{E}v$ satisfies the correspondence $\psi \Rightarrow \phi$, written $\mathcal{E}v \vdash \psi \Rightarrow \phi$, if and only if there exists a component-wise injective \mathbb{F} such that for all ρ defined on (ψ) , for all ψ^τ such that $\rho, \mathcal{E}v \vdash^{\psi^\tau} \psi$, there exists an extension ρ' of ρ to (ϕ) such that $\rho', \mathcal{E}v \vdash^{\mathbb{F}(\psi^\tau)} \phi$.

Intuitively, a sequence of events $\mathcal{E}v$ satisfies $\psi \Rightarrow \phi$ when, if $\mathcal{E}v$ satisfies ψ with execution steps defined by ψ^τ , then $\mathcal{E}v$ satisfies ϕ with execution steps defined by $\mathbb{F}(\psi^\tau)$. The injectivity is guaranteed because \mathbb{F} is component-wise injective. Definition 9 is unchanged for injective correspondences.

Example 3 The correspondence

$$\text{inj-event}(e_B(x, y, z)) \Rightarrow \text{inj-event}(e_A(x, y, z)) \quad (2)$$

means that, with overwhelming probability, each execution of $e_B(x, y, z)$ corresponds to a distinct execution of $e_A(x, y, z)$. In this case, ψ^τ is simply the execution step of $e_B(x, y, z)$ and ϕ^τ the execution step of $e_A(x, y, z)$. The function \mathbb{F} is an injective function that maps the execution step of $e_B(x, y, z)$ to the execution step of $e_A(x, y, z)$. (This step is never \perp .)

The correspondence

$$\begin{aligned} \text{event}(e_1(x)) \wedge \text{inj-event}(e_2(x)) \Rightarrow \text{inj-event}(e_3(x)) \vee \\ (\text{inj-event}(e_4(x, y)) \wedge \text{inj-event}(e_5(x, y))) \end{aligned}$$

means that, with overwhelming probability, for all x , if $e_1(x)$ has been executed, then each execution of $e_2(x)$ corresponds to distinct executions of $e_3(x)$ or to distinct executions of $e_4(x, y)$ and $e_5(x, y)$. The function \mathbb{F} maps $\perp \wedge \tau_2$ to $\tau_3 \vee (\tau_4 \wedge \tau_5)$, where $\tau_2, \tau_3, \tau_4, \tau_5$ are the execution steps of $e_2(x), e_3(x), e_4(x, y), e_5(x, y)$ respectively (either τ_3 or τ_4 and τ_5 may be \perp). The projections of \mathbb{F} map $\perp \wedge \tau_2$ to $\tau_3, \tau_4,$ and τ_5 respectively.

When no injective event occurs in $\psi \Rightarrow \phi$, Definition 10 reduces to the definition of non-injective correspondences.

Property

Lemma 14 *If Q satisfies a correspondence $corr$ with public variables V up to probability p and C is an acceptable evaluation context for Q with public variables V that does not contain events used in $corr$, then for all $V' \subseteq V \cup \text{var}(C)$, $C[Q]$ satisfies a correspondence $corr$ with public variables V' up to probability p' such that $p'(C') = p(C'[C])$.*

If $Q \approx_p^V Q'$ and Q satisfies a correspondence $corr$ with public variables V up to probability p' , then Q' satisfies $corr$ with public variables V up to probability p'' such that $p''(C) = p'(C) + p(C, t_{corr})$.

Proof Suppose that Q satisfies a correspondence $corr$ with public variables V and C is an acceptable evaluation context for Q with public variables V that does not contain events used in $corr$. Let $V' \subseteq V \cup \text{var}(C)$. Let C' be an evaluation context acceptable for $C[Q]$ with public variables V' that does not contain events used by $corr$. We rename the variables of C' not in V' so that they are not in V ; this renaming does not change the probabilities. We have

$$\text{Adv}_{C[Q]}^{corr}(C') = \Pr[C'[C[Q]] : \neg corr] \leq p(C'[C])$$

because $C'[C]$ is an evaluation context acceptable for Q with public variables V : there is no common table between C and Q , and between C' and $C[Q]$, so a fortiori between C' and Q , so there is no common table between $C'[C]$ and Q ; moreover

$$\begin{aligned} \text{var}(C'[C]) \cap \text{var}(Q) &= ((\text{var}(C') \cap \text{var}(Q)) \cup \text{var}(C)) \cap \text{var}(Q) \\ &\subseteq (V' \cup \text{var}(C)) \cap \text{var}(Q) && \text{since } \text{var}(C') \cap \text{var}(C[Q]) \subseteq V' \\ &\subseteq (V \cup \text{var}(C)) \cap \text{var}(Q) && \text{since } V' \subseteq V \cup \text{var}(C) \\ &\subseteq V && \text{since } \text{var}(C) \cap \text{var}(Q) \subseteq V \end{aligned}$$

We also have $\text{vardef}(C'[C[Q]]) \cap V = (\text{vardef}(C') \cap V) \cup (\text{vardef}(C) \cap V) = \emptyset$ since $\text{vardef}(C) \cap V = \emptyset$ because C is an acceptable evaluation context for Q with public variables V and $\text{vardef}(C') \cap V \subseteq \text{vardef}(C') \cap V' = \emptyset$ because we have renamed the variables of C' not in V' so that they are not in V and C' is an acceptable evaluation context for $C[Q]$ and with public variables V' .

Suppose that $Q \approx_p^V Q'$ and Q satisfies a correspondence $corr$ with public variables V up to probability p' . Let C be an evaluation context acceptable for Q' with public variables V that does not contain events used by $corr$. We have

$$\begin{aligned} \text{Adv}_{Q'}^{corr}(C) &= \Pr[C[Q'] : \neg corr] \\ &\leq \Pr[C[Q] : \neg corr] + |\Pr[C[Q'] : \neg corr] - \Pr[C[Q] : \neg corr]| \\ &\leq p'(C) + p(C, t_{corr}) \end{aligned}$$

Indeed, by renaming the variables and tables of Q that do not appear in Q' to variables and tables that also do not occur in C , C is also an acceptable evaluation context for Q with public variables V . \square

2.7.3 Computation of Advantages

Definition 11 Let C be an evaluation context acceptable for Q with any public variables. We define

$$\text{Adv}_Q(C, D) = \begin{cases} \Pr[C[Q] : D \wedge \neg \text{NonUnique}_Q] & \text{if } S \text{ does not occur in } D \\ \Pr[C[Q] : D \wedge \neg \text{NonUnique}_Q] - \Pr[C[Q] : \bar{S} \vee \text{NonUnique}_Q] & \text{if } D = S \vee D' \end{cases}$$

We assume that D is written as a logical formula (for instance, one of the correspondence formulas defined previously). The phrase “ S does not occur in D ” means that S occurs in this formula. We consider \vee as commutative and associative, so that $D = S \vee D'$ means $D = D_1 \vee \dots \vee D_l$ and $D_j = S$ for some $j \leq l$. In the following lemmas, the events used by D are the events that occur in this formula.

Lemma 15 1. *In the initial game, if C is an evaluation context acceptable for Q with public variables V that does not contain events used by D , and S does not occur in D , $\Pr[C[Q] : D] = \text{Adv}_Q(C, D)$.*

In the initial game, if C is an evaluation context acceptable for $Q \mid R_x$ with public variables V ($x \notin V$) that does not contain S nor \bar{S} , then $\text{Adv}_Q^{\text{Secrecy}(x)}(C) = \text{Adv}_Q(C[[\] \mid R_x], S)$.

In the initial game, if C is an evaluation context acceptable for $Q \mid Q_x$ with public variables V ($x \notin V$) that does not contain S nor \bar{S} , then $\text{Adv}_Q^{1-\text{ses.secrecy}(x)}(C) = \text{Adv}_Q(C[[\] \mid Q_x], S)$.

2. *If $Q, \text{EvUsed} \xrightarrow{V}_{p, D^+} Q', \text{EvUsed}'$, S and \bar{S} are in EvUsed , and C is an evaluation context acceptable for Q and Q' with public variables V that does not contain non-unique events of Q nor events in $\text{EvUsed}' \setminus \text{EvUsed}$, then*

- *if S does not occur in D , $\text{Adv}_Q(C, D) \leq p(C, t_D) + \text{Adv}_{Q'}(C, D \vee D^+)$;*
- *if $D = S \vee D'$, $\text{Adv}_Q(C, D) \leq 2p(C, t_D) + \text{Adv}_{Q'}(C, D \vee D^+)$.*

3. *If C is an evaluation context acceptable for Q with any public variables, then $\text{Adv}_Q(C, D \vee D') \leq \text{Adv}_Q(C, D) + \text{Adv}_Q(C, D')$, when S does not occur both in D and D' .*

Proof Property 1: the initial game does not contain $\text{find}[\text{unique}_e]$ (since it does not contain find), so NonUnique_Q is always false, and $D \wedge \neg \text{NonUnique}_Q = D$. Therefore, we have:

- In case S does not occur in D , $\Pr[C[Q] : D] = \text{Adv}_Q(C, D)$.
- If C is an evaluation context acceptable for $Q \mid R_x$ with public variables V ($x \notin V$) that does not contain S nor \bar{S} , then $\text{Adv}_Q^{\text{Secrecy}(x)}(C) = \Pr[C[Q \mid R_x] : S] - \Pr[C[Q \mid R_x] : \bar{S}] = \text{Adv}_Q(C[[\] \mid R_x], S)$.
- The case of one-session secrecy is similar to the case of secrecy.

Property 2, case S does not occur in D : since $Q, \text{EvUsed} \xrightarrow{V}_{p, D^+} Q', \text{EvUsed}'$ and C is acceptable for Q and Q' with public variables V and does not contain non-unique events in Q nor events in $\text{EvUsed}' \setminus \text{EvUsed}$, we have

$$\begin{aligned} \text{Adv}_Q(C, D) &= \Pr[C[Q] : D \wedge \neg \text{NonUnique}_Q] \\ &\leq p(C, t_D) + \Pr[C[Q'] : (D \vee D^+) \wedge \neg \text{NonUnique}_{Q'}] \\ &\leq p(C, t_D) + \text{Adv}_{Q'}(C, D \vee D^+). \end{aligned}$$

since S does not occur in D^+ , because $S \in \text{EvUsed}$.

Property 2, case $D = S \vee D'$: we have

$$\begin{aligned}
\text{Adv}_Q(C, D) &= \Pr[C[Q] : D \wedge \neg \text{NonUnique}_Q] - \Pr[C[Q] \mid R_x : \bar{S} \vee \text{NonUnique}_Q] \\
&\leq \Pr[C[Q'] : (D \vee D^+) \wedge \neg \text{NonUnique}_Q] + p(C, t_D) \\
&\quad - \Pr[C[Q'] : \bar{S} \vee \text{NonUnique}_Q] + p(C, t_{\bar{S}}) \\
&\leq \text{Adv}_{Q'}(C, D \vee D^+) + 2p(C, t_D)
\end{aligned}$$

since $t_{\bar{S}} \leq t_D$.

Property 3, case S does not occur in D nor D' :

$$\begin{aligned}
\text{Adv}_Q(C, D \vee D') &= \Pr[C[Q] : (D \vee D') \wedge \neg \text{NonUnique}_Q] \\
&= \Pr[C[Q] : (D \wedge \neg \text{NonUnique}_Q) \vee (D' \wedge \neg \text{NonUnique}_Q)] \\
&\leq \Pr[C[Q] : D \wedge \neg \text{NonUnique}_Q] + \Pr[C[Q] : D' \wedge \neg \text{NonUnique}_Q] \\
&\leq \text{Adv}_Q(C, D) + \text{Adv}_Q(C, D').
\end{aligned}$$

Property 3, case $D = S \vee D''$ and S does not occur in D' (the other case is symmetric):

$$\begin{aligned}
\text{Adv}_Q(C, D \vee D') &= \Pr[C[Q] : (D \vee D') \wedge \neg \text{NonUnique}_Q] - \Pr[C[Q] : \bar{S} \vee \text{NonUnique}_Q] \\
&\leq \Pr[C[Q] : D \wedge \neg \text{NonUnique}_Q] - \Pr[C[Q] : \bar{S} \vee \text{NonUnique}_Q] \\
&\quad + \Pr[C[Q] : D' \wedge \neg \text{NonUnique}_Q] \\
&\leq \text{Adv}_Q(C, D) + \text{Adv}_Q(C, D')
\end{aligned}$$

□

This lemma allows one to bound the advantage of the adversary against secrecy and correspondences. Property 1 is used in the initial game, to express the desired probability as $\text{Adv}_Q(C, D)$. Property 2 is used when a game Q is transformed into a game Q' during the proof. It allows one to bound the probability in Q from a bound in Q' . Property 3 is useful when distinct sequences of games are used for bounding the probabilities of the disjuncts D and D' . We bound these two probabilities $\text{Adv}_Q(C, D)$ and $\text{Adv}_Q(C, D')$ separately, then obtain a bound on $\text{Adv}_Q(C, D \vee D')$ by taking the sum.

Notice that, in Lemma 15, Property 1, if S does not occur in D , the context C can be any evaluation context acceptable for Q with public variables V that does not contain events used by D (since, in the initial game Q , Q contains no non-unique event at all). In subsequent game transformations, the introduced events and variables can be renamed so that they do not occur in C , by Lemma 12, Property 5. Therefore, Lemma 15 allows one to bound $\Pr[C[Q] : D]$ for any context C allowed in Definition 9. Similarly, for secrecy, C can be any evaluation context acceptable for $Q \mid R_x$ with public variables V ($x \notin V$) that does not contain S nor \bar{S} , so Lemma 15 allows one to bound $\text{Adv}_Q^{\text{Secrecy}(x)}(C)$ for any context C allowed in Definition 7. For one-session secrecy, C can be any evaluation context acceptable for $Q \mid Q_x$ with public variables V ($x \notin V$) that does not contain S nor \bar{S} , so Lemma 15 allows one to bound $\text{Adv}_Q^{1\text{-ses.}\text{Secrecy}(x)}(C)$ for any context C allowed in Definition 6.

More formally, consider the following cases:

- If we want to prove that Q_0 satisfies the correspondence $\psi \Rightarrow \phi$ with public variables V , we want to bound the probability $\Pr[C[Q_0] : \neg(\psi \Rightarrow \phi)]$ for any evaluation context C acceptable for Q_0 with public variables V that does not contain the events used by $\neg(\psi \Rightarrow \phi)$. We let $D_0 = \neg(\psi \Rightarrow \phi)$ and $C' = C$ be such a context. The special event S used for testing secrecy is supposed not to occur in D_0 . By Lemma 15, Property 1, we have $\text{Adv}_{Q_0}(C', D_0) = \Pr[C[Q_0] : \neg(\psi \Rightarrow \phi)]$.

- If we want prove that Q_0 preserves the secrecy of x with public variables V ($x \notin V$), let C be an evaluation context acceptable for $Q_0 \mid R_x$ with public variables V that does not contain the events S, \bar{S} . We want to bound the probability $\text{Adv}_{Q_0}^{\text{Secrecy}(x)}(C)$ that C breaks the secrecy of x . We let $D_0 = S$, $V = \{x\}$, and $C' = C[[\] \mid R_x]$. By Lemma 15, Property 1, we have $\text{Adv}_{Q_0}(C', D_0) = \text{Adv}_Q^{\text{Secrecy}(x)}(C)$.
- The situation for one-session secrecy is similar to the case of secrecy, using Q_x instead of R_x .

The proof produced by CryptoVerif can be represented as a tree whose nodes are labeled by triples (D, Q, EvUsed) and whose edges are labeled by triples (D', p, D^+) , where D^+ is a disjunction of Shoup events, and D and D' are disjunctions of Shoup events and possibly the initial formula D_0 . The root of the tree is labeled by $(D_0, Q_0, \text{EvUsed}_0)$ such that EvUsed_0 is the set containing S, \bar{S} , and the events used by D_0 or that occur in Q_0 . When a node labeled by (D, Q, EvUsed) has sons labeled by $(D_1, Q_1, \text{EvUsed}_1), \dots, (D_l, Q_l, \text{EvUsed}_l)$ linked with edges labeled respectively $(D'_1, p_1, D_1^+), \dots, (D'_l, p_l, D_l^+)$, then $D = D'_1 \vee \dots \vee D'_l$ (D is a disjunction of the form $e_1 \vee \dots \vee e_m$ or $D_0 \vee e_1 \vee \dots \vee e_m$, D'_1, \dots, D'_l are disjunctions that form a partition of the disjuncts of D), $D_j = D'_j \vee D_j^+$, and $Q, \text{EvUsed} \xrightarrow{V}_{p_j, D_j^+} Q_j, \text{EvUsed}_j$ for all $j \leq l$. When the proof is a basic sequence of games, each node has one son, which is the next game in the sequence, except the last game of the sequence which has no son. However, it may happen that distinct sequences of games are used to bound several events occurring in the game; in this case, there is a branching in the proof and a node has several sons. Examples of proof trees can be found in Figure 10; they are explained below.

By Lemma 12, Property 5, we build a similar tree such that, additionally, the events that occur in C' are in EvUsed_0 (by induction from the root to the leaves). We also rename the fresh variables introduced during the game transformations such that they do not occur in C' . Then we show by an easy induction that all nodes of this tree are labeled by (D, Q, EvUsed) such that EvUsed contains S, \bar{S} , and the events that occur in C' , and C' is an evaluation context acceptable for Q with public variables V that does not contain non-unique events that occur in Q . We associate to each node labeled by (D, Q, EvUsed) the advantage $\text{Adv}_Q(C', D)$, and we show how to bound these advantages for all nodes in the tree. Suppose that, in this tree, a node labeled by (D, Q, EvUsed) has sons labeled by $(D_1, Q_1, \text{EvUsed}_1), \dots, (D_l, Q_l, \text{EvUsed}_l)$ linked with edges labeled respectively $(D'_1, p_1, D_1^+), \dots, (D'_l, p_l, D_l^+)$. Then we have

$$\text{Adv}_Q(C', D) \leq \sum_{j=1}^l \text{Adv}_Q(C', D'_j)$$

by Lemma 15, Property 3, since $D = D'_1 \vee \dots \vee D'_l$. Moreover, if S does not occur in D , for all $j \leq l$,

$$\text{Adv}_Q(C', D'_j) \leq p_j(C', t_{D'_j}) + \text{Adv}_{Q_j}(C', D'_j \vee D_j^+) = p_j(C', t_{D'_j}) + \text{Adv}_{Q_j}(C', D_j)$$

by Lemma 15, Property 2, since $Q, \text{EvUsed} \xrightarrow{V}_{p_j, D_j^+} Q_j, \text{EvUsed}_j$. Therefore,

$$\text{Adv}_Q(C', D) \leq \sum_{j=1}^l \left(p_j(C', t_{D'_j}) + \text{Adv}_{Q_j}(C', D_j) \right)$$

If $D = S \vee D'$, then the event S occurs in exactly one formula D'_1, \dots, D'_l , say in D'_1 . We have

$$\begin{aligned} \text{Adv}_Q(C', D'_1) &\leq 2p_1(C', t_{D'_1}) + \text{Adv}_{Q_1}(C', D'_1 \vee D_1^+) \leq 2p_1(C', t_{D'_1}) + \text{Adv}_{Q_1}(C', D_1) \\ \text{Adv}_Q(C', D'_j) &\leq p_j(C', t_{D'_j}) + \text{Adv}_{Q_j}(C', D'_j \vee D_j^+) = p_j(C', t_{D'_j}) + \text{Adv}_{Q_j}(C', D_j) \text{ for } j \geq 2 \end{aligned}$$

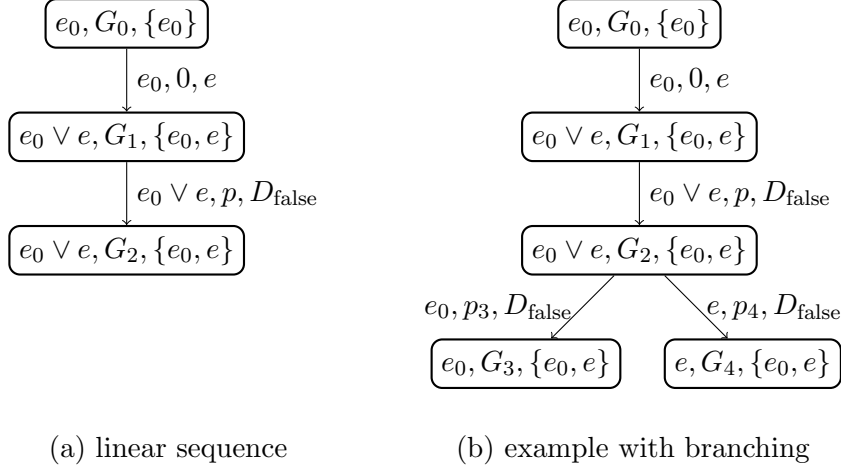


Figure 10: Examples of proof trees

by Lemma 15, Property 2, since $Q, EvUsed \xrightarrow{V}_{p_j, D_j^+} Q_j, EvUsed_j$. Therefore,

$$\text{Adv}_Q(C', D) \leq 2p_1(C', t_{D'_1}) + \sum_{j=2}^l p_j(C', t_{D'_j}) + \sum_{j=1}^l \text{Adv}_{Q_j}(C', D_j)$$

In all cases, we can then bound the advantage associated to a node from bounds on the advantages associated to its sons. Assuming that the advantage associated to the leaves of the tree is bounded, we can then bound the advantage associated to all nodes of the tree, by induction from the leaves to the root. At the root, we obtain a bound on $\text{Adv}_{Q_0}(C', D_0)$ which yields the desired result.

Lemma 15 allows us to obtain more precise probability bounds than the standard computation of probabilities generally done by cryptographers, when we use Shoup's lemma [11]. By Shoup's lemma, if G' is obtained from G by inserting an event e and modifying the code executed after e , the probability of distinguishing G' from G is bounded by the probability of executing e : for all contexts C acceptable for G and G' (with any public variables) and all distinguishers D , $|\Pr[C[G] : D] - \Pr[C[G'] : D]| \leq \Pr[C[G'] : e]$. Hence,

$$\Pr[C[G] : D] \leq \Pr[C[G'] : e] + \Pr[C[G'] : D].$$

We improve over this computation of probabilities by considering e and D simultaneously instead of making the sum of the two probabilities:

$$\Pr[C[G] : D] \leq \Pr[C[G'] : D \vee e].$$

For example, suppose that we want to bound the probability of event e_0 in G_0 ; we transform G_0 into G_1 using Shoup's lemma, so that G_1 differs from G_0 only when G_1 executes event e , and we have $G_0, EvUsed_0 \rightarrow_{0,e} G_1, EvUsed_1$; then we transform G_1 into G_2 , so that $G_1 \approx_p G_2$, and $G_1, EvUsed_1 \rightarrow_{p, D_{\text{false}}} G_2, EvUsed_1$; and G_2 executes neither e_0 nor e . The corresponding proof tree is given in Figure 10(a). Let C be an evaluation context acceptable for G_0 without public variables that does not contain event e_0 . Lemma 15 yields

$$\begin{aligned} \Pr[C[G_0] : e_0] &= \text{Adv}_{G_0}(C, e_0) \\ &\leq \text{Adv}_{G_1}(C, e_0 \vee e) \\ &\leq p(C, t_{e_0 \vee e}) + \text{Adv}_{G_2}(C, e_0 \vee e) = p(C, t_{e_0 \vee e}) \end{aligned}$$

If we suppose for simplicity that no `find[unique]` occurs, so that NonUnique_{G_i} is always false, we have $\text{Adv}_{G_i}(C, D) = \Pr[C[G_i] : D]$, so we can write the previous formulas simply using probabilities:

$$\Pr[C[G_0] : e_0] \leq \Pr[C[G_1] : e_0 \vee e] \leq p(C, t_{e_0 \vee e}) + \Pr[C[G_2] : e_0 \vee e] = p(C, t_{e_0 \vee e}).$$

In contrast, the standard computation of probabilities yields

$$\Pr[C[G_0] : e_0] \leq \Pr[C[G_1] : e_0] + \Pr[C[G_1] : e] \leq p(C, t_{e_0}) + p(C, t_e).$$

The runtime t_D of D is essentially the same for e_0 , e , and $e_0 \vee e$, so $\Pr[C[G_0] : e_0] \leq p(C, t_D)$ by Lemma 15, while $\Pr[C[G_0] : e_0] \leq 2p(C, t_D)$ by the standard computation, so we have gained a factor 2. The probability that comes from the transformation of G_1 into G_2 is counted once (for distinguisher $e_0 \vee e$) instead of counting it twice (once for e_0 and once for e).

The standard computation of probabilities corresponds to applying point 3 of Lemma 15 to bound each probability separately and compute the sum, as soon as the considered distinguisher D has several disjuncts. Instead, we use point 3 of Lemma 15 only when the proof uses different sequences of games to bound the probabilities of the events, as in Figure 10(b).

2.8 Turing Machine Adversary

In `CryptoVerif`, the adversary is modeled as an evaluation context. However, usually, in cryptographic results, an adversary is a bounded-time probabilistic Turing machine. In this section, we explain how any bounded-time probabilistic Turing machine that communicates on channels can be represented as a `CryptoVerif` evaluation context.

Let Q_0 be the initial game that interacts with an adversary. Let c_1, \dots, c_k be the channels used in Q_0 . Let T_{all} be the union of all types that occur in Q_0 . Let T'_{all} be the type of pairs containing the encoding a channel as first component and an element of T_{all} as second component. The encoding of a channel is either the constant *yield* or a tuple of integers $(j, i_1, \dots, i_{k'})$ with $1 \leq j \leq k$. (We assume that unambiguous tuples can be encoded as `CryptoVerif` values, and that the constant *yield* is different from a tuple.) Let d_0, d_1 , and d_2 be channels that do not occur in Q_0 .

Let Q_1 be a process that contains the parallel composition of processes

$$!^{i_1 \leq n_1} \dots !^{i_{k'} \leq n_{k'}} c_j[i_1, \dots, i_{k'}](x : T_{\text{all}}).\overline{d_0}(\langle (j, i_1, \dots, i_{k'}), x \rangle)$$

for each output $\overline{c_j}[i'_1, \dots, i'_{k'}]\langle N \rangle$ that occurs under $!^{i'_1 \leq n_1} \dots !^{i'_{k'} \leq n_{k'}}$ in Q_0 . Since, in the initial game Q_0 , the channels of all outputs use the current replication indices as channel indices, as in $c_j[i'_1, \dots, i'_{k'}]$, a single output is executed for each value of the indices and for each syntactic occurrence of the output, so the inputs in Q_1 can receive all outputs made by Q_0 . The process Q_1 forwards all these outputs to the same channel d_0 , with a message that specifies both the channel $c_j[i_1, \dots, i_{k'}]$ on which Q_0 emitted (encoded as a bitstring) and the message x sent by Q_0 .

In addition, Q_1 also contains the parallel composition of processes

$$!^{i_1 \leq n_1} \dots !^{i_{k'} \leq n_{k'}} \text{yield}().\overline{d_0}(\langle \text{yield}, () \rangle)$$

for each occurrence of `yield` that occurs under $!^{i'_1 \leq n_1} \dots !^{i'_{k'} \leq n_{k'}}$ in Q_0 , to receive all outputs that come from the `yield` construct.

Let $C = \text{newChannel } d_0; \text{newChannel } d_1; \text{newChannel } d_2; (\text{start}().\overline{d_1}[1]\langle s_0 \rangle \mid Q_1 \mid Q_2 \mid [])$, where the process Q_2 is defined in Figure 11. Let us explain how the context C can simulate any Turing machine interacting with the process Q_0 .

The current state of the Turing machine is sent on channel $d_1[i]$ where i is a loop index that starts at 1 and increases during execution. As shown in the semantics of `CryptoVerif`,

```

1   $Q_2 = !^{i \leq n} d_1[i](s : \text{bitstring});$ 
2  let  $(s', o, v) = f(s)$  in
    Lines 3–6 are repeated for each  $j \leq k$  and each  $k'$ 
    such that there is an input on channel  $c_j[i'_1, \dots, i'_k]$  in  $Q_0$ .
3  if  $o = (j, k')$  then
4      let  $(a_1, \dots, a_{k'}, b) = v$  in  $\overline{c_j[a_1, \dots, a_{k'}]}(b);$ 
5       $d_0(s'' : T_{\text{all}}; \overline{d_1[i+1]}(f'(s', s''))$ 
6  else
7
7  if  $o = \text{random}$  then
8      new  $x : \text{bool}; \overline{d_1[i+1]}(f''(s', x))$ 
9  else
10 if  $o = \text{abort}$  then
11     event_abort  $e$ 
12 else
13      $\overline{d_2}()$ 

```

Figure 11: Looping process

upon startup, a message is sent on channel *start*. When *C* receives that message, it sends the initial state of the Turing machine s_0 on channel $d_1[1]$. This message is received by process Q_2 (line 1). Then Q_2 calls the function f on the current state s of the Turing machine (line 2). This function executes the Turing machine, until one of the following situations happens:

- The Turing machine sends a message b on a channel $c_j[a_1, \dots, a_{k'}]$; in this case, f returns $(s', (j, k'), (a_1, \dots, a_{k'}, b))$, where s' is the new state of the Turing machine. The test at line 3 is then going to succeed for the appropriate value of j, k' , and the desired message is going to be sent at line 4. After receiving a message, the process Q_0 always replies by sending a message (except if it aborts). This message is going to be received by Q_1 , which is going to forward on d_0 the channel and the received message. These channel and message are then received as s'' at line 5. Then $f'(s', s'')$ is the new state of the Turing machine after receiving that message. This state is sent on channel $d_1[i+1]$, which restarts a new iteration of Q_2 .
- The Turing machine generates a fresh random bit; in this case, f returns $(s', \text{random}, ())$ where s' is the new state of the Turing machine. The test at line 7 is then going to succeed. At line 8, a random bit x is chosen. Then $f''(s', x)$ is the new state of the Turing machine with that random bit. This state is sent on channel $d_1[i+1]$, which restarts a new iteration of Q_2 as in the previous case.
- The Turing machine aborts; in this case, f returns $(s', \text{abort}, ())$. The test at line 10 is then going to succeed, and the process aborts at line 11. (The event e is any event not used elsewhere; the event is not really useful, it is present because the CryptoVerif language always executes an event before aborting.)
- The Turing machine stops; in this case, f returns $(s', \text{stop}, ())$. No test succeeds, so line 13 is executed. The process tries to send a message on channel d_2 , but there is no input on this channel, so the process blocks.

The constants *random*, *abort*, and *stop* are assumed to be pairwise distinct, and distinct from all pairs.

The function f is a CryptoVerif primitive, because it can be implemented by a deterministic bounded-time Turing machine. (Recall that f stops when the initial probabilistic Turing machine makes a random choice, and the random choice is performed by CryptoVerif at lines 7–8.) Similarly, the function f' that computes the new state of the Turing machine from the old state and the received message, and the function f'' that computes the new state of the Turing machine from the old state and a random bit are CryptoVerif primitives.

The replication bound n (used in Q_2 , line 1) is chosen large enough so that the loop never stops due to that bound: the Turing machine aborts or stops before the bound is reached. This is possible since the Turing machine runs in bounded time, so sends a bounded number of messages and chooses a bounded number of random bits.

Notice that, if Q_0 sends and receives messages on the same channels, it may happen that a message sent by Q_0 is immediately received by Q_0 without being intercepted by the adversary. In this case, since both Q_0 and Q_1 are going to listen on the same channels, the destination of the message (either the honest process Q_0 or the adversary Q_1) is chosen randomly with uniform probability, depending on the number of available receivers. Therefore, adding more copies of the receiving processes in Q_1 increases the probability that the adversary receives the message. Moreover, when the same channel is used for both inputs and outputs, the messages sent by Q_2 at line 4 may be received back by the adversary via Q_1 , instead of being received by Q_0 . We recommend avoiding this strange situation, by using distinct channels for inputs on the one hand and outputs on the other hand. More generally, we recommend using distinct channels for each input and output, so that the adversary gets full control of the network, as already mentioned page 7.

As a slight extension, it would still be possible to allow Q_0 to output on $c_j[i_1, \dots, i_{k'}]$ after receiving a message on the same channel $c_j[i_1, \dots, i_{k'}]$. In this case, a message sent by Q_0 on $c_j[i_1, \dots, i_{k'}]$ cannot be received by Q_0 , because the input on $c_j[i_1, \dots, i_{k'}]$ is no longer available when the output on $c_j[i_1, \dots, i_{k'}]$ is performed by Q_0 . Moreover, the problem that messages sent by Q_2 at line 4 may be received back by the adversary via Q_1 , instead of being received by Q_0 , can be avoided by putting the receiver process

$$c_j[a_1, \dots, a_{k'}](x : T_{\text{all}}).\overline{d_0}(\langle(j, a_1, \dots, a_{k'}), x\rangle)$$

after $\overline{c_j[a_1, \dots, a_{k'}]}(b)$ in parallel with $d_0(s'' : T_{\text{all}}); \overline{d_1[i+1]}(f'(s', s''))$ in Q_2 , instead of including

$$!^{i_1 \leq n_1} \dots !^{i_{k'} \leq n_{k'}} c_j[i_1, \dots, i_{k'}](x : T_{\text{all}}).\overline{d_0}(\langle(j, i_1, \dots, i_{k'}), x\rangle)$$

in Q_1 .

The context C does not allow the Turing machine to execute events of its choice, while a CryptoVerif context can execute events. We could obviously extend the model to allow the Turing machine to execute events, but this is not needed for the cases we consider. Indeed, if the adversary represented as CryptoVerif context executes events, these events can be deleted without changing the final result returned by the distinguisher: for correspondences, by Definition 9, the context is not allowed to contain events used by $\psi \Rightarrow \phi$, and all other events are ignored by the distinguisher $\neg(\psi \Rightarrow \phi)$; for one-session secrecy and secrecy, by Definitions 6 and 7, the context is not allowed to contain S nor \overline{S} , and all other events are ignored by the distinguishers S and \overline{S} .

To sum up, the context given in this section allows us to run any probabilistic bounded-time Turing machine as a CryptoVerif context, so CryptoVerif contexts are powerful enough to represent the adversaries usually considered by cryptographers.

Acknowledgments I warmly thank David Pointcheval for his advice and explanations of the computational proofs of protocols. This project would not have been possible without him. I

also thank Jacques Stern for initiating this work. The design and implementation of CryptoVerif was partly done while I was at CNRS and at Ecole Normale Supérieure. It was partly supported by the ANR through the projects FormaCrypt (ARA SSIA 2005) and ProSe (VERSO 2010, decision number 2010-VERS-004).

References

- [1] M. Abdalla, P.-A. Fouque, and D. Pointcheval. Password-based authenticated key exchange in the three-party setting. *IEE Proceedings Information Security*, 153(1):27–39, Mar. 2006.
- [2] M. Bellare and C. Namprempe. Authenticated encryption: Relations among notions and analysis of the generic composition paradigm. In T. Okamoto, editor, *Advances in Cryptology – ASIACRYPT’00*, volume 1976 of *Lecture Notes in Computer Science*, pages 531–545, Berlin, Heidelberg, Dec. 2000. Springer.
- [3] M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In S. Vaudenay, editor, *Advances in Cryptology – Eurocrypt 2006 Proceedings*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426, Berlin, Heidelberg, May 2006. Springer. Extended version available at <http://eprint.iacr.org/2004/331>.
- [4] B. Blanchet. Computationally sound mechanized proofs of correspondence assertions. In *20th IEEE Computer Security Foundations Symposium (CSF’07)*, pages 97–111, Los Alamitos, CA, July 2007. IEEE Computer Society Press. Extended version available as ePrint Report 2007/128, <http://eprint.iacr.org/2007/128>.
- [5] B. Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Transactions on Dependable and Secure Computing*, 5(4):193–207, Oct.–Dec. 2008.
- [6] D. Cadé and B. Blanchet. From computationally-proved protocol specifications to implementations. In *7th International Conference on Availability, Reliability and Security (ARES 2012)*, pages 65–74, Los Alamitos, CA, Aug. 2012. IEEE Computer Society Press.
- [7] P. Laud. Secrecy types for a simulatable cryptographic library. In *12th ACM Conference on Computer and Communications Security (CCS’05)*, pages 26–35, New York, NY, Nov. 2005. ACM Press.
- [8] J. C. Mitchell, A. Ramanathan, A. Scedrov, and V. Teague. A probabilistic polynomial-time calculus for the analysis of cryptographic protocols. *Theoretical Computer Science*, 353(1–3):118–164, Mar. 2006.
- [9] V. Shoup. A proposal for an ISO standard for public-key encryption, Dec. 2001. ISO/IEC JTC 1/SC27.
- [10] V. Shoup. OAEP reconsidered. *Journal of Cryptology*, 15(4):223–249, Sept. 2002.
- [11] V. Shoup. Sequences of games: a tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332, Nov. 2004. Available at <http://eprint.iacr.org/2004/332>.