



HAL
open science

CryptoVerif: a Computationally-Sound Security Protocol Verifier

Bruno Blanchet, Charlie Jacomme

► **To cite this version:**

Bruno Blanchet, Charlie Jacomme. CryptoVerif: a Computationally-Sound Security Protocol Verifier. RR-9526, Inria. 2023, pp.194. hal-04253820

HAL Id: hal-04253820

<https://inria.hal.science/hal-04253820v1>

Submitted on 23 Oct 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Inria

CryptoVerif: a Computationally-Sound Security Protocol Verifier

Bruno Blanchet, Charlie Jacomme

**RESEARCH
REPORT**

N° 9526

October 2023

Project-Team Prosecco

ISRN INRIA/RR--9526--FR+ENG

ISSN 0249-6399



CryptoVerif: a Computationally-Sound Security Protocol Verifier

Bruno Blanchet*, Charlie Jacomme*

Project-Team Prosecco

Research Report n° 9526 — October 2023 — 194 pages

Abstract: This document presents the security protocol verifier CryptoVerif. CryptoVerif does not rely on the symbolic, Dolev-Yao model, but on the computational model. It can verify secrecy, correspondence properties (which include authentication), and indistinguishability properties. It produces proofs presented as sequences of games, like those manually written by cryptographers; these games are formalized in a probabilistic process calculus. CryptoVerif provides a generic method for specifying security properties of the cryptographic primitives. It produces proofs valid for 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. CryptoVerif is post-quantum sound: when the used cryptographic assumptions are valid for quantum adversaries, the proofs hold for quantum adversaries. It can work automatically, or the user can guide it with manual proof indications.

Key-words: security protocols, verification, computational model

* Inria

**RESEARCH CENTRE
PARIS**

2 rue Simone Iff - CS 42112
75589 Paris Cedex 12

CryptoVerif: un vérificateur de protocoles cryptographiques sûr dans le modèle calculatoire

Résumé : Ce document présente le vérificateur de protocoles cryptographiques CryptoVerif. CryptoVerif ne s'appuie pas sur le modèle symbolique de Dolev-Yao, mais sur le modèle calculatoire. Il peut vérifier le secret, les correspondances (qui comprennent l'authentification) et les propriétés d'indistinguabilité. Il produit des preuves par suites de jeux, comme celles écrites manuellement par les cryptographes ; ces jeux sont formalisés dans un calcul de processus probabiliste. CryptoVerif fournit une méthode générique pour spécifier les propriétés de sécurité des primitives cryptographiques. Il produit des preuves valables pour un nombre quelconque de sessions du protocole, et fournit une borne supérieure sur la probabilité de succès d'une attaque contre le protocole en fonction de la probabilité de casser chaque primitive et du nombre de sessions. CryptoVerif est sûr contre des attaquants quantiques : quand les hypothèses cryptographiques utilisées sont valides pour des attaquants quantiques, les preuves produites sont valides pour des attaquants quantiques. Il peut fonctionner automatiquement, ou l'utilisateur peut guider la preuve manuellement.

Mots-clés : protocoles cryptographiques, vérification, modèle calculatoire

Contents

| | | |
|----------|--|------------|
| 1 | Introduction | 5 |
| 2 | A Calculus for Cryptographic Games | 8 |
| 2.1 | Syntax and Informal Semantics | 8 |
| 2.2 | Example | 17 |
| 2.3 | Type System | 20 |
| 2.4 | Formal Semantics | 25 |
| 2.4.1 | Definition of the Semantics | 25 |
| 2.4.2 | Properties | 35 |
| 2.4.3 | Each Variable is Defined at Most Once | 44 |
| 2.4.4 | Each Oracle is defined at most once | 47 |
| 2.4.5 | Variables are Defined Before Being Used | 48 |
| 2.4.6 | Typing | 50 |
| 2.5 | Subset for the Initial Game | 51 |
| 2.6 | Input Language with Channels instead of Oracles | 53 |
| 2.6.1 | Input Language Description | 53 |
| 2.6.2 | Translating from the Channel Calculus to the Oracle Calculus | 54 |
| 2.7 | Subsets used inside the Sequence of Games | 56 |
| 2.8 | Security Properties, Indistinguishability | 58 |
| 2.8.1 | Secrecy | 66 |
| 2.8.2 | Secrecy for a Bit | 69 |
| 2.8.3 | Correspondences | 73 |
| 2.8.4 | Computation of Advantages | 79 |
| 2.8.5 | Proof of Indistinguishability | 88 |
| 2.8.6 | Proof of query_equiv | 89 |
| 3 | Collecting True Facts | 92 |
| 3.1 | User-defined Rewrite Rules | 96 |
| 3.2 | Collecting True Facts from a Game | 97 |
| 3.3 | Local Dependency Analysis | 102 |
| 3.4 | Equational Prover | 105 |
| 4 | success: Criteria for Proving Security Properties | 108 |
| 4.1 | Secrecy | 109 |
| 4.2 | Correspondences | 120 |
| 4.2.1 | Example | 120 |
| 4.2.2 | Non-unique Events | 123 |
| 4.2.3 | Non-injective Correspondences | 123 |
| 4.2.4 | Injective Correspondences | 128 |
| 5 | Game Transformations | 135 |
| 5.1 | Syntactic Game Transformations | 135 |
| 5.1.1 | auto_SArename | 135 |
| 5.1.2 | expand_tables | 136 |
| 5.1.3 | expand | 136 |
| 5.1.4 | prove_unique | 137 |
| 5.1.5 | remove_assign | 137 |
| 5.1.6 | use_variable | 138 |

| | | |
|--------|--|-----|
| 5.1.7 | SArename | 139 |
| 5.1.8 | move | 140 |
| 5.1.9 | move array | 141 |
| 5.1.10 | move up | 141 |
| 5.1.11 | move_if_fun | 143 |
| 5.1.12 | insert_event | 144 |
| 5.1.13 | insert | 144 |
| 5.1.14 | replace | 144 |
| 5.1.15 | merge_branches | 145 |
| 5.1.16 | merge_arrays | 146 |
| 5.1.17 | guess i | 147 |
| 5.1.18 | guess $x[c_1, \dots, c_m]$ | 156 |
| 5.1.19 | guess_branch | 159 |
| 5.1.20 | global_dep_anal | 161 |
| 5.1.21 | simplify | 163 |
| 5.1.22 | all_simplify | 167 |
| 5.1.23 | success simplify | 167 |
| 5.2 | crypto: Applying the Security Assumptions on Primitives | 175 |
| 5.2.1 | Basic Transformation | 175 |
| 5.2.2 | Optimizations for $\text{transf}_\phi(FB)$ | 184 |
| 5.2.3 | Guiding the Application of Equivalences | 184 |
| 5.2.4 | Relaxing Hypothesis H6 | 185 |
| 5.2.5 | Relaxing Hypothesis H'2 | 185 |
| 5.2.6 | Soundness and Example | 186 |
| 6 | Proof Strategy | 187 |
| 7 | Conclusion | 188 |

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 only 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 previous 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 [22, 66, 67]: 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 defining oracles interacting with an adversary. This calculus tries to mimic how classical cryptographic games are defined.

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 used for representing cryptographic games, 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 represent global variables that are shared between all processes, and each cell can be given a value by a process for instance through an assignment or a random sampling. After being given a value, cells cannot be updated anymore. 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. As described in Section 5.2, 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 receive the arguments of the oracle and return its result. 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, simple 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 often sufficient to specify a few well-chosen case distinctions and 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 instance 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.

Finally, as we demonstrate in the rest of this document, CryptoVerif is post-quantum sound: when the used cryptographic assumptions are valid for quantum adversaries, the proofs hold for quantum adversaries. The game transformations performed by CryptoVerif are valid not only against a classical probabilistic Turing-machine adversary, but also against a quantum adversary. More generally, they are valid against any kind of interactive black-box adversary that can simulate bounded time probabilistic Turing machines, limited by the cryptographic assumptions.

In addition to the oracle-based process calculus, CryptoVerif also allows to define games in a channel-based process calculus inspired by the applied pi calculus [1, 2] and by the calculi of [59] and of [55]. The channel-based calculus is presented in Section 2.6 with a translation to the oracle-based calculus. The channel-based calculus was the initial version of the calculus used by CryptoVerif, as detailed in a previous report [29]. Future developments will use the oracle-based calculus with the semantics presented in this document. This language allowed us to prove the post-quantum soundness of CryptoVerif. It is closer to classical computational proofs and paves the way for important extensions, such as sharing parts of the code in calls to internal oracles, composition results, and loops implemented by recursive oracle calls, as it is easier to inline oracle calls than channel communications.

CryptoVerif has been implemented in OCaml (more than 60000 lines of code) and is available at <http://cryptoverif.inria.fr/>.

Related Work Various methods have been proposed for verifying security protocols in the computational model. Following the seminal paper by Abadi and Rogaway [3], many results show the soundness of the Dolev-Yao model with respect to the computational model, which makes it possible to use Dolev-Yao provers in order to prove protocols in the computational model (see, e.g., [7, 37, 41, 42, 50] and the survey [40]). However, these results have limitations, in particular in terms of allowed cryptographic primitives (they must satisfy strong security properties so that they correspond to Dolev-Yao style primitives), and they require some restrictions on

protocols (such as the absence of key cycles). A tool [39] was developed based on [41] to obtain computational proofs using the formal verifier AVISPA, for protocols that rely on public-key encryption and signatures.

Several frameworks exist for formalizing proofs of protocols in the computational model. Backes, Pfitzmann, and Waidner [9, 10] designed an abstract cryptographic library and showed its soundness with respect to computational primitives, under arbitrary active attacks. This framework has been used for a computationally-sound machine-checked proof of the Needham-Schroeder-Lowe protocol [70, 71]. Canetti [34] introduced the notion of universal composability. With Herzog [36], they show how a Dolev-Yao-style symbolic analysis can be used to prove security properties of protocols within the framework of universal composability, for a restricted class of protocols using public-key encryption as only cryptographic primitive. Then, they use the automatic Dolev-Yao verification tool ProVerif [24] for verifying protocols in this framework. Process calculi have been designed for representing cryptographic games, such as the probabilistic polynomial-time calculus of [59] and the cryptographic lambda-calculus of [62]. Logics have also been designed for proving security protocols in the computational model, such as the computational variant of PCL (Protocol Composition Logic) [46, 47] and CIL (Computational Indistinguishability Logic) [13]. Canetti *et al.* [35] use the framework of time-bounded task-PIOAs (Probabilistic Input/Output Automata) to prove security protocols in the computational model. This framework makes it possible to combine probabilistic and non-deterministic behaviors. These frameworks can be used to prove security properties of protocols in the computational sense, but except for [36] which relies on a Dolev-Yao prover, they have not been automated up to now, as far as we know.

Several techniques have been used for directly mechanizing proofs in the computational model. Type systems [45, 55, 57, 69] provide computational security guarantees. For instance, [55] handles shared-key and public-key encryption, with an unbounded number of sessions, by relying on the Backes-Pfitzmann-Waidner library. A type inference algorithm is given in [8]. The recent tool OWL [48] also relies on a type system that provides computational security guarantees. It supports MACs, public-key signatures, authenticated symmetric and public key encryption, random oracles, and the gap Diffie-Hellman assumption [63]. It can prove secrecy and integrity properties. In another line of research, a specialized Hoare logic was designed for proving asymmetric encryption schemes in the random oracle model [43, 44].

The tool CertiCrypt [14, 15, 17, 19, 20] enables the machine-checked construction and verification of cryptographic proofs by sequences of games [22, 68]. It relies on the general-purpose proof assistant Coq, which is widely believed to be correct. Nowak *et al.* [5, 60, 61] follow a similar idea by providing Coq proofs for several cryptographic primitives. More recently, frameworks for cryptographic proofs in Coq, FCF [64], and in Isabelle, CryptHOL [18], have been designed and used for proving cryptographic schemes. EasyCrypt [16], the successor of CertiCrypt, no longer generates Coq proofs, but provides a higher automation level by relying on SMT solvers, which makes the tool easier to use. Even if it focuses more on cryptographic primitives and schemes than on protocols, it has been used for proving some protocols such as one-round key exchange [12], e-voting [38], AWS key management [6], and distance bounding [31]. These frameworks and tools can perform more subtle reasoning than CryptoVerif, at the cost of more user effort: the user has to give all games and guide the proof that the games are indistinguishable. That becomes tedious for large protocols, which require many large games.

The tool Squirrel [11] relies on a computationally sound logic that allows to write interactive proofs of, e.g., stateful protocols. Still, it currently proves a security notion weaker than the standard one: the number of sessions of the protocol must be bounded independently of the security parameter (instead of being polynomial in the security parameter).

Independently, we have built the tool CryptoVerif [26] to help cryptographers, not only for

the verification, but also by generating the proofs by sequences of games [22, 68], automatically or with little user interaction. In particular, CryptoVerif generates the games, possibly using the indications of which transformations to perform. This tool extends considerably early work by Laud [53, 54] which was limited either to passive adversaries or to a single session of the protocol. More recently, Tšahhrov and Laud [56, 72] developed a tool similar to CryptoVerif but that represents games by dependency graphs. It handles public-key and shared-key encryption and proves secrecy properties; it does not provide bounds on the probability of success of an attack.

Outline The next section presents our process calculus for representing games, with its syntax, type system, formal semantics, as well as the definition of security properties. Section 3 collects information about games and reasons using it. Section 4 gives criteria for proving security properties of protocols. Section 5 describes the game transformations that we use for proving protocols. Section 6 explains how the prover chooses which transformation to apply at each point.

Notations We recall the following notations. We denote by $\{M_1/x_1, \dots, M_m/x_m\}$ the substitution that replaces x_j with M_j for each $j \in \{1, \dots, m\}$. The cardinal of a set or multiset S is denoted by $|S|$. Multisets S are represented by functions that map each element x of S to the number of occurrences of x in S , that is, when S is a multiset, $S(x)$ is the number of elements of S equal to x . We use \uplus for multiset union, defined by $(S_1 \uplus S_2)(x) = S_1(x) + S_2(x)$. When S and S' are multisets, $\max(S, S')$ is the multiset such that $\max(S, S')(x) = \max(S(x), S'(x))$. The notation $\{x_1 \mapsto a_1, \dots, x_m \mapsto a_m\}$ designates the function that maps x_j for a_j for each $j \in \{1, \dots, m\}$ and is undefined for other inputs. When f is a function, $f[x \mapsto a]$ is the function that maps x to a and all other elements as f . 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 , and we denote by $a \stackrel{R}{\leftarrow} D$ the random sampling of a value following the distribution D . We denote by $\text{Distr}(S)$ the set of discrete probability distributions over a set S . We denote by \emptyset the empty stack, and a non empty stack St can be denoted by $a :: St_l$, where a is then the head of St and St_l its tail stack.

2 A Calculus for Cryptographic Games

2.1 Syntax and Informal Semantics

CryptoVerif represents games in the syntax of Figure 1. This syntax assumes a countable set of oracles names, denoted by \mathcal{O} . It uses parameters, denoted by n , which are integers that bound the number of executions of oracles.

It also uses types, denoted by T , which are sets of values. So that types can be meaningfully interpreted by an attacker, we assume that there exists an efficient bijection from the disjoint union of all types to bitstrings. That implies in particular that all types are countable.

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: $\text{bool} = \{\text{true}, \text{false}\}$, where false is 0 and true is 1; bitstring is the set of all bitstrings; $\text{bitstring}_\perp = \text{bitstring} \cup \{\perp\}$ where \perp is a special symbol; $[1, n]$ is the set of integers $\{1, \dots, n\}$, where n is a parameter. (We consider integers as bitstrings without leading zeroes.) We denote by *any* the specific type corresponding to the empty set of

| | |
|---|---|
| $M, N ::=$ i $x[M_1, \dots, M_m]$ $f(M_1, \dots, M_m)$ $x[\tilde{i}] \stackrel{R}{\leftarrow} T; N$ $\text{let } p = M \text{ in } N \text{ else } N'$ $\text{let } x[\tilde{i}] : T = M \text{ in } N$ $\text{if } M \text{ then } N \text{ else } N'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M'_j \text{ then } N_j) \text{ else } N'$ $\text{insert } Tbl(M_1, \dots, M_l); N$ $\text{get}[\text{unique?}] Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } N \text{ else } N'$ $\text{event } e(M_1, \dots, M_l); N$ $\text{event_abort } e$ | 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 |
| $p ::=$ $x[\tilde{i}] : T$ $f(p_1, \dots, p_m)$ $= M$ | pattern variable function application comparison with a term |
| $Q ::=$ 0 $Q \mid Q'$ $\text{foreach } i \leq n \text{ do } Q$ $\text{newOracle } O; Q$ $O[\tilde{i}](p) := P$ | oracle definitions nil parallel composition replication n times oracle restriction oracle declaration |
| $P ::=$ $\text{return } (N); Q$ $\text{let } p = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N) \text{ in } P \text{ else } P'$ $x[\tilde{i}] \stackrel{R}{\leftarrow} T; P$ $\text{let } p = M \text{ in } P \text{ else } P'$ $\text{if } M \text{ then } P \text{ else } P'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P$ $\text{insert } Tbl(M_1, \dots, M_l); P$ $\text{get}[\text{unique?}] Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } P \text{ else } P'$ $\text{event } e(M_1, \dots, M_l); P$ $\text{event_abort } e$ yield | oracle body return oracle call random number assignment conditional array lookup insert in table get from table event event e and abort end |

Figure 1: Syntax for games

values, which is thus contained in any other type. No other type can be empty, and *any* cannot be used within the syntax of the games.

The syntax 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); test if *_fun* (M_1, M_2, M_3) (with a first argument of type *bool* and the last two arguments of the same type T ; it returns a value of type T : M_2 when M_1 is true and M_3 when M_1 is false).

In this syntax, terms represent computations on bitstrings. The replication index i is an integer which serves in distinguishing different copies of a replicated oracle *foreach* $i \leq n$ *do* . (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 oracle bodies and explained below. These constructs conclude by evaluating a term, instead of executing an oracle body. The construct *event_abort* e executes event e (without argument) and aborts the game.

The calculus distinguishes two kinds of processes, oracle definitions and oracle bodies.

An oracle definition Q defines a set of parallel oracles ready to be called, each oracle being defined with a body P that may end with a return value, followed by a potential new oracle definition Q' .

The oracle definition 0 defines no oracle; $Q \mid Q'$ is the parallel composition of Q and Q' , it defines both oracles in Q and in Q' ; *foreach* $i \leq n$ *do* Q represents n copies of Q in parallel, each with a different value of $i \in [1, n]$; *newOracle* $O; Q$ creates a new private oracle O and executes Q ; this construct is useful in proofs, but does not occur in games manipulated by *CryptoVerif*. The oracle declaration $O[\tilde{i}](p) := P$ defines the oracles $O[\tilde{i}]$, where the indices \tilde{i} are bound by the replications above the oracle declaration $O[\tilde{i}](p) := P$. The oracle $O[\tilde{i}]$ executes the oracle body P when it is called with an argument that matches the pattern p . It will be explained in more detail below together with the semantics of the oracle call *let* $p = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N)$ *in* P *else* P' .

The oracle body construct $x[\tilde{i}] \stackrel{R}{\leftarrow} T; P$ 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 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

$$d(D_1, D_2) = \frac{1}{2} \sum_{a \in T} |D_1(a) - D_2(a)|$$

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 euclidean 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

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

so we can take $\epsilon_T = \frac{m}{2^{k+1}}$. 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 $x[\tilde{i}] \stackrel{R}{\leftarrow} T; P$. Function symbols represent deterministic functions, so all random numbers must be chosen by $x[\tilde{i}] \stackrel{R}{\leftarrow} T$. Deterministic functions make automatic syntactic manipulations easier: we can duplicate a term without changing its value.

The construct $\text{let } x[\tilde{i}] : T = M \text{ in } P$ 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 $\text{let } f(x_1, \dots, x_m) = M \text{ in } P \text{ else } P'$ as an abbreviation for $\text{let } y : T = M \text{ in let } x'_1 : T_1 = f_1^{-1}(y) \text{ in } \dots \text{let } x'_m : T_m = f_m^{-1}(y) \text{ in if } f(x'_1, \dots, x'_m) = y \text{ then (let } x_1 : T_1 = x'_1 \text{ in } \dots \text{let } x_m : T_m = x'_m \text{ in } P) \text{ else } P'$ 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 $\text{let } p = M \text{ in } P \text{ else } P'$ 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. In this case, the alternative syntax $x \leftarrow M; P$ is a synonym for $\text{let } x = M \text{ in } P$.

The construct $\text{event } e(M_1, \dots, M_l); P$ 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 oracle. Events are used in particular for specifying security properties.

The construct `event_abort` e executes event e (without argument) and aborts the game.

Next, we explain the construct `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` `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` `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` `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` `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 `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 $\epsilon_{\text{find}}/2$, that is, $d(D_{\text{find}}(S), \text{uniform}) \leq \epsilon_{\text{find}}/2$, that is, $\frac{1}{2} \sum_{v \in S} \left| D_{\text{find}}(S)(v) - \frac{1}{|S|} \right| \leq \frac{\epsilon_{\text{find}}}{2}$, so that, when $|S| = |S'|$, for any bijection $\phi : S \rightarrow S'$, $\frac{1}{2} \sum_{v \in S} |D_{\text{find}}(S)(v) - D_{\text{find}}(S')(\phi(v))| \leq \epsilon_{\text{find}}$. Moreover $D_{\text{find}}(S)(v_i) = D_{|S|}(i)$ where $S = \{v_1, \dots, v_{|S|}\}$ with the values v_i ordered in increasing order lexicographically, for some distribution $D_{|S|}$ that depends only on the cardinal of S . In other words, the probability of a value v_i in the distribution $D_{\text{find}}(S)$ does not depend on the values in the set S but only on the number $|S|$ of elements of S and on the position i of the value v_i in S ordered in increasing order lexicographically. Therefore, transformations that do not modify the number of successful values nor their order, that is, transformations that map elements v of S to elements $\phi(v)$ of S' at the same position i , preserve the probabilities exactly and we do not need to add ϵ_{find} to the probability when we apply such a transformation. This is true for instance when we remove a branch of `find` that is never taken. 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 make the condition of the `find` evaluate to true, we execute the event e and abort the game. When there is zero or one solution, 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` M `then` P `else` P' executes P if M evaluates to true. Otherwise, it executes

P' . CryptoVerif also supports the conditional `if defined(M_1, \dots, M_l) \wedge M then P else P'` , which executes P if M_1, \dots, M_l are defined and M evaluates to true. Otherwise, it executes P' . This conditional is internally encoded as `find suchthat defined(M_1, \dots, M_l) \wedge M then P else P'` . The conjunct M can be omitted when it is true, writing `if defined(M_1, \dots, M_l) then P else P'` .

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 $D_{\text{get}}(S)$ where S is the set of indices of suitable elements, with $d(D_{\text{get}}(S), \text{uniform}) \leq \epsilon_{\text{find}}/2$.) When no such element is found, P' is executed. We can generalize this construct to patterns instead of variables similarly to the `let` case. As in the case of `find`, the indication `[unique?]` stands for either `[uniquee]` or empty. The empty case has just been explained. When the `get` is marked `[uniquee]` and there are several solutions, we execute the event e and abort the game. When there is zero or one solution, the `get` is executed as when `[unique?]` is empty. CryptoVerif internally translates the `insert` and `get` constructs into `find`.

Let us explain the `return (N)`, `yield`, and oracle call commands. A semantic configuration always contains an oracles execution stack, the head oracle being the one currently executed and the stack corresponding to the nested calls. It also contains a set of callable oracles.

Whenever we encounter an oracle declaration $O[\tilde{i}](p) := P$, this oracle is added to the set of callable oracles. Intuitively, it is then possible to call this oracle for the given indices and with an argument that matches the pattern p . In the following, we will consider that all oracle declarations are of the form $O[\tilde{i}](x[\tilde{i}] : T) := P$, where we encode the input pattern for $O[\tilde{i}](p) := P$ as $O[\tilde{i}](x[\tilde{i}] : T) := \text{let } p = x \text{ in } P \text{ else yield}$.

Whenever the head oracle executes an oracle call `let $x[\tilde{i}] : T = O[M_1, \dots, M_l](N)$ in P else P'` , one looks for an available oracle $O[a_1, \dots, a_l]$ where M_1, \dots, M_l evaluate to a_1, \dots, a_l . (Terms M_1, \dots, M_l are intuitively analogous to IP addresses and ports, which are numbers that the adversary may guess.) If no such callable oracle is found, the execution blocks. Otherwise, by the invariants given below, there is at most one such oracle $O[a_1, \dots, a_l](x'[\tilde{i}'] : T') := P''$. This oracle is then executed: the argument N is evaluated and stored in $x'[\tilde{i}']$ if it is in T' (otherwise the execution blocks). The well-typed condition is not a restriction for the attacker, as it can always use type conversion functions.

Finally, the oracle $O[a_1, \dots, a_l](x'[\tilde{i}'] : T') := P''$ is removed from the set of callable oracles, the body P'' of the called oracle is pushed on top of the stack, and executed. In this situation, we then have a new oracle body currently being executed, and the second one in the stack is the latest oracle call `let $x[\tilde{i}] : T = O[M_1, \dots, M_l](N)$ in P else P'` waiting for a return value.

When an oracle performs a return command `return (N'); Q` , it returns normally, and the in branch of the caller is executed. In more details, we pop the return command from the execution stack. The head of the execution stack is then the latest oracle call, of the form `let $x[\tilde{i}] : T = O[M_1, \dots, M_l](N)$ in P else P'` . We evaluate N' , store it in $x[\tilde{i}]$ if it is in T and continue the execution with P , in a similar fashion to an assignment. The oracles in Q , which follow the executed return command, are added to the set of callable oracles.

When an oracle performs a `yield` command, it returns abnormally, and the else branch of the caller is executed. In more detail, we also pop the `yield` command from the execution stack. The head of the execution stack is then the latest oracle call, of the form `let $x[\tilde{i}] : T = O[M_1, \dots, M_l](N)$ in P else P'` , and we execute P' .

In the case of an oracle call with pattern matching `let $p = O[M_1, \dots, M_l](N)$ in P else P'` , the pattern-matching is encoded as in assignments. The oracle body P' is executed in case the

pattern-matching fails. We note that in contrast to the assignment encoding, we still need an else branch for the specific failure case of `yield` that cannot be encoded with a conditional. Oracles that take several arguments or return several results are encoded using tuples and pattern-matching on tuples.

Oracle calls can be made with some unspecified indices, following the syntax

$$\text{let } x[\tilde{i}] : T = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N) \text{ in } P \text{ else } P' .$$

In this case, we evaluate M_1, \dots, M_l to a_1, \dots, a_l , and consider the set S of oracles $O[a_1, \dots, a_l, b_1, \dots, b_m]$ among the set of callable oracles. We choose one of these oracles randomly following the probability distribution $D_{\text{call}}(S)$, which is approximately uniform like $D_{\text{find}}(S)$. Then we perform the call as explained above. In the remainder of the execution, $?u_1[\tilde{i}], \dots, ?u_m[\tilde{i}]$ are bound to b_1, \dots, b_m . In a call to oracle O , only indices corresponding to replications that are directly above the oracle declaration of O can be unspecified. Indices that are used inside a previous oracle definition must then be explicitly specified. For instance, if oracle O_2 is defined by

$$\text{foreach } i_1 \leq n_1 \text{ do } O_1[i_1](p_1) := \dots \text{return } (N); \text{foreach } i_2 \leq n_2 \text{ do } O_2[i_1, i_2](p_2) := \dots$$

in call to O_2 , the first index i_1 must always be explicitly specified by a term M_1 ; the second index may either be specified by a term M_2 or left unspecified by $?u_2[\tilde{i}] \leq n_2$.

The execution depends on a parametric attacker \mathcal{A} . Whenever the execution stack is empty, the attacker can perform an oracle call, and the attacker receives the return of the last oracle in the stack. Oracle O can be declared private with `newOracle` O , in which case only oracles under the scope of `newOracle` O can call O , and in particular the attacker cannot. (This is useful in the proofs, although all oracles of protocols are often public.) In effect, the attacker can then be seen as controlling the network between communicating oracles.

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`, $\stackrel{R}{\leftarrow}$, or `insert` and `in yield` may be omitted after `let`. A trailing 0 after a return may be omitted. We also omit oracle parameters and replication indices if there are none.

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 `foreach` $i \leq n$ `do` Q binds the replication index i in Q . The `find` construct

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

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 at the definition of x , but it should be kept in mind that this is only an abbreviation. Variables defined under a replication must be arrays: for example `foreach` $i_1 \leq n_1$ `do` \dots `foreach` $i_m \leq n_m$ `do` `let` $x[i_1, \dots, i_m] : T = M$ `in` \dots . More formally, we require the following invariant:

Invariant 1 (Single definition) The oracle definition 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`, if (or `let`), oracle call, or `get`.

In a `let` with pattern-matching, `let $p = M$ in P else P'` , the variables bound by p are considered to be defined in the `in` branch; however, the variables defined in M and in terms included in the pattern p are defined before the branching.

Similarly, in an oracle call, `let $p = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N)$ in P else P'` , the variables bound by p are considered to be defined in the `in` branch; however, the variables defined in N , M_1, \dots, M_l and in terms included in the pattern p as well as the variables u_1, \dots, u_m are defined before the branching.

In `get $[\text{unique?}]$ $Tbl(p_1, \dots, p_l)$ suchthat M in P else P'` , the variables bound by p_j for $j \leq l$ and the variables defined in M and in terms included in the patterns p_j are (temporarily) defined before the branching.

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 $x[\tilde{i}] \stackrel{R}{\leftarrow} T$, a `let`, `get`, or oracle declaration that contains the pattern $x[\tilde{i}] : T$, the construct `find... $x[\tilde{i}] = i \leq n \dots$` , or the oracle call `let $p = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N)$ in P else P'` where p contains the pattern $x[\tilde{i}] : T$ or when $x[\tilde{i}]$ is $u_j[\tilde{i}]$ for some $j \leq m$.

Invariant 2 (Defined variables) The oracle definition 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` command or term;
- or in M'_j in a process or term of the form `find ($\bigoplus_{j=1}^{m''} \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j$ suchthat $\text{defined}(M'_{j_1}, \dots, M'_{j_{l_j}}) \wedge M'_j$ then P_j) 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 `find ($\bigoplus_{j=1}^{m''} \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j$ suchthat $\text{defined}(M'_{j_1}, \dots, M'_{j_{l_j}}) \wedge M'_j$ then P_j) 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). The scope of variable definitions is defined as follows: $x[\tilde{i}]$ is syntactically under its definition when it is

- inside P in $x[\tilde{i}] \stackrel{R}{\leftarrow} T; P$;
- inside N in $x[\tilde{i}] \stackrel{R}{\leftarrow} T; N$;
- inside P in `let $p = M$ in P else P'` when $x[\tilde{i}] : T$ is bound in the pattern p ;
- inside N in `let $p = M$ in N else N'` when $x[\tilde{i}] : T$ is bound in the pattern p ;
- inside N in `let $x[\tilde{i}] : T = M$ in N` ;

- inside P_j in $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P$ when x is u_{jk} for some $k \leq m_j$;
- inside N_j in $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } N_j) \text{ else } N$ when x is u_{jk} for some $k \leq m_j$;
- inside M or P in $\text{get}[\text{unique?}] \text{Tbl}(p_1, \dots, p_l) \text{ suchthat } M \text{ in } P \text{ else } P'$ when $x[\tilde{i}] : T$ is bound in one of the patterns p_1, \dots, p_l ;
- inside M or N in $\text{get}[\text{unique?}] \text{Tbl}(p_1, \dots, p_l) \text{ suchthat } M \text{ in } N \text{ else } N'$ when $x[\tilde{i}] : T$ is bound in one of the patterns p_1, \dots, p_l ;
- inside P in $O[\tilde{i}](p) := P$ when $x[\tilde{i}] : T$ is bound in the pattern p .
- inside P in $\text{let } p = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N) \text{ in } P \text{ else } P'$ when $x[\tilde{i}] : T$ is bound in the pattern p , or when x is u_m for some $j \leq m$.

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. In this invariant, we consider a finite set of public variables V : these variables are defined in the considered oracle definition Q_0 (though it may sometimes happen that they are never defined) and can be accessed by the adversary via the `find` construct.

Invariant 3 (Variables defined in find and get conditions) The oracle definition 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 oracle definition Q_0 satisfies Invariant 4 if and only if `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.

Definition 1 A term is *simple* when it contains only replication indices, variables, and function applications.

Invariant 5 (Terms in defined conditions) The oracle definition Q_0 satisfies Invariant 5 if and only if the terms M_1, \dots, M_l in conditions `defined`(M_1, \dots, M_l) in `find` are simple.

Terms that are not simple are handled by expanding them into their corresponding oracle body commands. Invariant 5 is needed because terms in `defined` conditions cannot be expanded (see the transformation `expand` in Section 5.1.3). By combining this invariant with Invariant 2, we see that the terms of all variable accesses $x[M_1, \dots, M_m]$ are simple.

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

- Shoup events occur only in processes or terms of the form `event.abort e` in Q_0 ,

- non-unique events occur only in `find[uniquee]` or `get[uniquee]` in Q_0 , and
- other events occur in `event` $e(M_1, \dots, M_l)$ or in `event_abort` e in Q_0 .

The name “Shoup events” is used because these events are introduced when applying Shoup’s lemma [68] (see Section 5.1.12). The non-unique events are those triggered when a `find[uniquee]` or `get[uniquee]` actually has several solutions.

To provide an invariant expressing that oracles for a given name are defined only once, we define the following multiset by induction on oracle definitions and bodies:

$$\begin{aligned}
\text{Defined}^O(0) &= \emptyset \\
\text{Defined}^O(Q_1 \mid Q_2) &= \text{Defined}^O(Q_1) \uplus \text{Defined}^O(Q_2) \\
\text{Defined}^O(\text{foreach } i \leq n \text{ do } Q) &= \text{Defined}^O(Q) \\
\text{Defined}^O(O[\tilde{i}](x[\tilde{i}] : T) := P) &= \{O\} \uplus \text{Defined}^O(P) \\
\text{Defined}^O(\text{newOracle } O; Q) &= \text{Defined}^O(Q) \setminus \{O\} \\
\text{Defined}^O(C[Q_1, \dots, Q_k]) &= \max_{j=1}^k \text{Defined}^O(Q_j)
\end{aligned}$$

Here, we denote by C an oracle body context, that is, a context with only oracle body constructs and several holes for oracle definitions as possible continuations, and $C[Q_1, \dots, Q_k]$ corresponds to instantiating the continuations with the given processes. We use $C[Q_1, \dots, Q_k]$ to define by induction the function at once for all oracle body constructs.

Invariant 7 (Single Oracle definition) The oracle definitions Q_0 satisfies Invariant 7 if and only if

1. in every oracle declaration of $O[\tilde{i}](p)$ in Q_0 , the indices \tilde{i} of O are the current replication indices at that declaration, and
2. $\text{Defined}^O(Q_0)$ does not contain duplicate elements.

Similar to the invariant 1 for variables, we guarantee that each oracle is correctly defined at most once. Thanks to this invariant, we can omit the indices of oracles in oracle declarations, writing $O(p) := P$ for $O[\tilde{i}](p) := P$ where \tilde{i} are the current replication indices at that declaration.

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 $\text{fo}(P)$ the set of free oracles of P . (We use similar notations for oracle definitions.)

2.2 Example

Let us introduce two cryptographic primitives that we use below.

Definition 2 Let T_{mk} and T_{ms} be types that correspond intuitively to keys and message authentication codes, respectively; T_{mk} is a fixed-length type. A message authentication code scheme MAC [21] consists of two function symbols:

- $\text{mac} : \text{bitstring} \times T_{mk} \rightarrow T_{ms}$ 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 random coins as an additional argument.)
- $\text{verify} : \text{bitstring} \times T_{mk} \times T_{ms} \rightarrow \text{bool}$ 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 $\forall m \in \text{bitstring}, \forall k \in T_{mk}, \text{verify}(m, k, \text{mac}(m, k)) = \text{true}$.

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

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

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 [25, 26], 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 3 Let T_k , T_r , and T_e be types for random coins, keys, and ciphertexts respectively. T_k and T_r are fixed-length types. A symmetric encryption scheme SE [21] consists of two function symbols:

- $\text{enc} : \text{bitstring} \times T_k \times T_r \rightarrow T_e$ 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 $\forall k \in T_k, \forall m \in \text{bitstring}, \forall r \in T_r, \text{dec}(\text{enc}(m, k, r), k) = m$.

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

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

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 on 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:

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

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 [21]. 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 oracle definition Q_0 :

$$\begin{aligned} Q_0 &= O_0() := x_k \xleftarrow{R} T_k; x_{mk} \xleftarrow{R} T_{mk}; \text{return } (); (Q_A \mid Q_B) \\ Q_A &= \text{foreach } i \leq n \text{ do } O_A() := x'_k \xleftarrow{R} T_k; x_r \xleftarrow{R} T_r; \\ &\quad \text{let } x_m : \text{bitstring} = \text{enc}(\text{k2b}(x'_k), x_k, x_r) \text{ in return } (x_m, \text{mac}(x_m, x_{mk})) \\ Q_B &= \text{foreach } i' \leq n' \text{ do } O_B(x'_m, x_{ma}) := \text{if verify}(x'_m, x_{mk}, x_{ma}) \text{ then} \\ &\quad \text{let } i_{\perp}(\text{k2b}(x''_k)) = \text{dec}(x'_m, x_k) \text{ in return } () \end{aligned}$$

When the oracle O_0 is called, it begins execution: it generates the keys x_k and x_{mk} randomly. Then it returns control to the adversary, and provides n copies of oracles O_A and n' copies of O_B for A and B respectively. The adversary can call them by calling $O_A[i]$ and $O_B[i']$ respectively. In a session that runs as expected, the adversary first calls $O_A[i]$. Then O_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 returns the ciphertext and the MAC. 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 returned message to $O_B[i]$. When O_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 call the oracles 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 oracle, 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) and oracle calls, the type-checking 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 $i_j \leq n_j$;
- oracle names O to types $[1, n_1] \times \dots \times [1, n_m] \rightarrow T \rightarrow T'$, where the definition of $O[i_1, \dots, i_m](p) := P$ of type $T \rightarrow T'$ 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 `foreach $i_1 \leq n_1$ do , . . . , foreach $i_m \leq n_m$ do` 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 $x[\tilde{i}] \stackrel{R}{\leftarrow} T$ or in a pattern $x[\tilde{i}] : T$ in an assignment, an oracle call, an oracle declaration, or a `get`. In the `find` construct, `find . . . $x[\tilde{i}] = i \leq n$` , the type T of x is $T = [1, n]$. In the oracle call `let $p = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N)$ in P else P'` , the type T_j of u_j is $[1, n_j]$ for $j \leq m$. 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. Similarly, all definitions of the same oracle O must yield the same value of $\mathcal{E}(O)$. While the type of variables can be determined by a preliminary pass on their definitions, provided their type is explicit at their definition, determining the type of oracles requires typing the oracle bodies. The type environment for oracles would then not be fully known at the beginning, but determined during typing by collecting constraints on the type of each oracle, provided the pattern in which the result of an oracle call is stored is explicitly typed. However, as we explain below, CryptoVerif currently does not accept oracle calls in its input language, which simplifies the typing algorithm. We sketch this algorithm below.

Oracles are typechecked in the type environment \mathcal{E} using the rules of Figures 2, 3, and 4. 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} .
- $\mathcal{E} \vdash P : T$ means that the oracle body P will return a value of type T in environment \mathcal{E} .
- $\mathcal{E} \vdash Q$ means that the oracle definition Q is well-typed in environment \mathcal{E} .

When a term M never evaluates to a value (it always ends with `event_abort e`), it has the special type *any*, which can be converted to any other type by subtyping. Similarly, when an oracle body P never returns (it always ends with `yield` or `event_abort e`), it has type *any*. The type *any* is not allowed in declarations of function symbols, tables, and events. Variables can have type *any* only when they are actually never defined (their definition is a term that always aborts). Patterns can have type *any* only when the pattern-matching will always fail (because it is an equality with a term that always aborts). The result type of oracles can be *any*, when they never return.

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 x[\tilde{i}] \stackrel{R}{\leftarrow} 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{\mathcal{E} \vdash M : \text{bool} \quad \mathcal{E} \vdash N : T \quad \mathcal{E} \vdash N' : T}{\mathcal{E} \vdash \text{if } M \text{ then } N \text{ else } N' : T} \quad \text{(TIIfT)} \\
\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{\text{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 } \text{Tbl}(M_1, \dots, M_l); N : T} \quad \text{(TInsertT)} \\
\frac{\text{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}[\text{unique?}] \text{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{(TEventT)} \\
\frac{e : ()}{\mathcal{E} \vdash \text{event.abort } e : \text{any}} \quad \text{(TEventAbortT)} \\
\frac{\mathcal{E} \vdash M : \text{any}}{\mathcal{E} \vdash M : T} \quad \text{(SubtypingT)}
\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 oracle definitions:

$$\begin{array}{c}
\mathcal{E} \vdash 0 \quad \text{(TNil)} \\
\frac{\mathcal{E} \vdash Q \quad \mathcal{E} \vdash Q'}{\mathcal{E} \vdash Q \mid Q'} \quad \text{(TPar)} \\
\frac{\mathcal{E}[i \mapsto [1, n]] \vdash Q}{\mathcal{E} \vdash \text{foreach } i \leq n \text{ do } Q} \quad \text{(TRepl)} \\
\frac{\mathcal{E}[O \mapsto T_1 \times \dots \times T_m \rightarrow T \rightarrow T'] \vdash Q}{\mathcal{E} \vdash \text{newOracle } O; Q} \quad \text{(TNewOracle)} \\
\frac{\mathcal{E} \vdash p : T \quad \mathcal{E} \vdash P : T' \quad \forall j \leq m, \mathcal{E} \vdash i_j : T_j \quad \mathcal{E}(O) = T_1 \times \dots \times T_m \rightarrow T \rightarrow T'}{\mathcal{E} \vdash O[i_1, \dots, i_m](p) := P} \quad \text{(TODecl)}
\end{array}$$

Figure 3: Typing rules (2)

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 $x[\tilde{i}] \stackrel{R}{\leftarrow} T; N$ 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 $x[\tilde{i}] \stackrel{R}{\leftarrow} T; N$.

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 if M then N else N' , M must be of type *bool* and N and N' must have the same type, which is also the type of if M then N else N' .

In

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

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[*unique?*] $Tbl(p_1, \dots, p_l)$ 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 is of type *any*, and can thus be of any type (because it aborts the game); the event e must be declared without argument, which we denote by $e : ()$. Terms of type

Typing rules for oracle bodies:

$$\begin{array}{c}
\frac{\mathcal{E} \vdash N : T \quad \mathcal{E} \vdash Q}{\mathcal{E} \vdash \text{return } (N); Q : T} \quad (\text{TRet}) \\
\\
\frac{\begin{array}{c} \forall j \leq l, \mathcal{E} \vdash M_j : T'_j \quad \forall j \leq m, \mathcal{E} \vdash u_j[\tilde{i}] : [1, n_j] \\ \mathcal{E} \vdash N : T \quad \mathcal{E} \vdash p : T' \quad \mathcal{E} \vdash P : T'' \quad \mathcal{E} \vdash P' : T'' \\ \mathcal{E} \vdash O : T'_1 \times \dots \times T'_j \times [1, n_1] \times \dots \times [1, n_m] \rightarrow T \rightarrow T' \end{array}}{\mathcal{E} \vdash \text{let } p = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N) \text{ in } P \text{ else } P' : T''} \quad (\text{TCall}) \\
\\
\frac{\begin{array}{c} T \text{ fixed, bounded, or nonuniform} \quad \mathcal{E} \vdash x[\tilde{i}] : T \quad \mathcal{E} \vdash P : T' \\ \mathcal{E} \vdash x[\tilde{i}] \stackrel{R}{\leftarrow} T; P : T' \end{array}}{\mathcal{E} \vdash x[\tilde{i}] \stackrel{R}{\leftarrow} T; P : T'} \quad (\text{TNew}) \\
\\
\frac{\mathcal{E} \vdash M : T \quad \mathcal{E} \vdash p : T \quad \mathcal{E} \vdash P : T' \quad \mathcal{E} \vdash P' : T'}{\mathcal{E} \vdash \text{let } p = M \text{ in } P \text{ else } P' : T'} \quad (\text{TLet}) \\
\\
\frac{\mathcal{E} \vdash M : \text{bool} \quad \mathcal{E} \vdash P : T \quad \mathcal{E} \vdash P' : T}{\mathcal{E} \vdash \text{if } M \text{ then } P \text{ else } P' : T} \quad (\text{TIf}) \\
\\
\frac{\begin{array}{c} \forall j \leq m, \forall k \leq m_j, \mathcal{E} \vdash u_{jk}[\tilde{i}] : [1, n_{jk}] \\ \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} \\ \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} \\ \forall j \leq m, \mathcal{E} \vdash P_j : T \quad \mathcal{E} \vdash P : T \end{array}}{\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} \\ \text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P : T} \quad (\text{TFind}) \\
\\
\frac{\begin{array}{c} 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 : T \\ \mathcal{E} \vdash \text{insert } Tbl(M_1, \dots, M_l); P : T \end{array}}{\mathcal{E} \vdash \text{insert } Tbl(M_1, \dots, M_l); P : T} \quad (\text{TInsert}) \\
\\
\frac{\begin{array}{c} 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 : T \quad \mathcal{E} \vdash P' : T \\ \text{get}[\text{unique?}] Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } P \text{ else } P' : T \end{array}}{\text{get}[\text{unique?}] Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } P \text{ else } P' : T} \quad (\text{TGet}) \\
\\
\frac{\begin{array}{c} 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 : T \\ \mathcal{E} \vdash \text{event } e(M_1, \dots, M_l); P : T \end{array}}{\mathcal{E} \vdash \text{event } e(M_1, \dots, M_l); P : T} \quad (\text{TEvent}) \\
\\
\frac{e : ()}{\mathcal{E} \vdash \text{event_abort } e : \text{any}} \quad (\text{TEventAbort}) \\
\\
\mathcal{E} \vdash \text{yield} : \text{any} \quad (\text{TYield}) \\
\\
\frac{\mathcal{E} \vdash P : \text{any}}{\mathcal{E} \vdash P : T} \quad (\text{Subtyping}) \\
\\
\frac{\mathcal{E}(O) = T_1 \times \dots \times T_m \rightarrow T \rightarrow T'}{\mathcal{E} \vdash O : T_1 \times \dots \times T_m \rightarrow T \rightarrow T'} \quad (\text{TO}) \\
\\
\frac{\mathcal{E} \vdash O : T_1 \times \dots \times T_m \rightarrow T \rightarrow \text{any}}{\mathcal{E} \vdash O : T_1 \times \dots \times T_m \rightarrow T \rightarrow T'} \quad (\text{SubtypingO})
\end{array}$$

Figure 4: Typing rules (3)

any can then be given any type through the subtyping rule.

The type system for oracle definitions and bodies requires each subterm to be well-typed. In `foreach` $i \leq n$ `do` Q , i is of type $[1, n]$ in Q . The oracle commands $\overset{R}{\text{let}}$, `if`, `find`, `insert`, `get`, `event`, and `event.abort` are typed similarly to the corresponding terms. The command `yield`, which aborts the execution of an oracle without returning any value, is typed with *any*. A `return` (N) is typed with the type of the term N . In an oracle call `let` $p = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N) in P `else` P' , the indices $M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m$ must have the same type as the indices at the definition of oracle O , the type of N must correspond to the type expected by the oracle O , the type of the result of oracle O must correspond to the type of the pattern p , and both branches P and P' must have the same type, which is also the type of the whole oracle call construct. Subtyping is lifted from terms to oracles with dedicated rules, and can be used either on an oracle body or an oracle definition.$

We say that an occurrence of a term M in an oracle definition 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 `foreach` $i \leq n$ `do` above M in Q and with $i_{j1} \mapsto [1, n_{j1}], \dots, i_{jm_j} \mapsto [1, n_{jm_j}]$ for each `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` `defined`(M_{j1}, \dots, M_{jl_j}) $\wedge M_j$ `then` P_j) `else` P such that the considered occurrence of M is in the condition `defined`(M_{j1}, \dots, M_{jl_j}) $\wedge M_j$.

Invariant 8 (Typing) The oracle definition Q_0 satisfies Invariant 8 if and only if there exists a type environment \mathcal{E} for Q_0 such that $\mathcal{E} \vdash Q_0$.

On the implementation side, we only have to check that input processes following the subset for the initial game (see Figure 13) are well-typed, so we do not need to deal with oracle calls and `newOracle`. In practice, the type of oracles in the type environment is determined by maintaining a list of oracles when going through the process, storing for each oracle its name, the type of indices, the type of arguments, and the potential type of the returns. For distinct branches of execution, we merge the oracle lists allowing duplicated names but ensuring that the same oracle name always corresponds to the same number of indices, the same type of arguments, and that the return type either corresponds to none or the same. For parallel execution branches and for oracles defined syntactically one under the other, we forbid oracles with the same names, to guarantee Invariant 7.

The type of variables in the type environment is determined by inspecting all variables definitions in a preliminary pass before the actual typing. Variables with several incompatible definitions (different indices, different types, or variables defined in the same branch of a test, which would contradict Invariant 1) are automatically renamed to distinct variable names. Such variables cannot be used queries and nor in `defined` conditions of `find`. For simplicity, variables defined without giving them an explicit type also cannot be used in queries and nor in `defined` conditions of `find`, even if their type is later inferred by the typing algorithm, and are also renamed to distinct variable names.

Looking forward, an adversary against a protocol will be made of two components:

- an attacker \mathcal{A} that can be instantiated by any black-box interactive and stateful function;
- an adversary context C , inside our game syntax, that allows the adversary to, e.g., raise events.

We require the adversary context 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.

The attacker calls to oracles will also be well-typed by the semantics, which once again does not restrict its computing power.

2.4 Formal Semantics

2.4.1 Definition of the Semantics

The formal semantics of our programming language is presented in Figures 5, 6, 8, 9, 10, and 11.

In this semantics, each term M or process P or Q is labeled by a program point μ , replacing M with ${}^\mu M$ and similarly for P and Q . We still use the notations M , P , Q for terms and processes tagged with program points. The program points are used in order to track from where each term or process comes from in the initial oracle definition. These program points are simply constant tags, and the initial oracle definition is tagged with a distinct program point at each subterm and subprocess.

The semantics depends on an attacker that may interact with the oracles. We consider attackers \mathcal{A} from a recursively defined attacker space \mathcal{BB} , that are probabilistic functions of type $\mathcal{A} : \text{Bitstring} \mapsto \text{Distr}(\mathcal{BB} \times \text{Bitstring})$. To each input, the attacker \mathcal{A} gives out a distribution yielding the corresponding output and next attacker step with some probability. It models an interactive black-box attacker, where given an input, it gives back with some probability both an output and the next attacker step. Here, $\text{Bitstring} = \{0, 1\}^*$ is the set of concrete bitstrings, and we implicitly rely on the efficient bijections from types and oracle names to concrete bitstrings so that the attacker can interact with the oracles.

A semantic configuration is a sextuple $E, St, Q, Or, \mathcal{T}, \mu\mathcal{E}v$, where

- E is an environment mapping array cells to values.
- St is an execution stack of the form $(\sigma_0, P_0) :: \dots :: (\sigma_m, P_m) :: \mathcal{A}$ where P_j are oracle bodies and σ_j are the associated mapping sequences which give values of replication indices. The oracle body P_0 is the one currently being executed, and the oracle bodies P_1, \dots, P_m correspond to oracle calls that have not returned yet. The stack always contains at least one pair (σ_j, P_j) and always contains as a last element an attacker \mathcal{A} .

The mapping sequence $\sigma = [i_1 \mapsto a_1, \dots, i_m \mapsto a_m]$ is a sequence of mappings $i_j \mapsto a_j$, which can also be interpreted as a function: $\sigma(i_j) = a_j$ for all $j \leq m$, and $\sigma(\tilde{i})$ is defined by the natural extension to sequences. However, using a sequence allows us to define $\text{Dom}(\sigma) = [i_1, \dots, i_m]$ to be the sequence of current replication indices and $\text{Im}(\sigma) = [a_1, \dots, a_m]$ to be the sequence of their values. When $\sigma = [i_1 \mapsto a_1, \dots, i_m \mapsto a_m]$, $\sigma[i_{m+1} \mapsto a_{m+1}, \dots, i_l \mapsto a_l] = [i_1 \mapsto a_1, \dots, i_l \mapsto a_l]$.

- Q is the multiset of oracle declarations running in parallel with P , with their associated mapping sequences giving values of replication indices. It is the set of callable oracles.
- $Or = (Or_{\text{pub}}, Or_{\text{priv}})$ represents the oracles already created. Or_{pub} is the set of public oracles, which can be called by the adversary; Or_{priv} is the set of private oracles.
- \mathcal{T} defines the contents of tables. It is a list of $Tbl(a_1, \dots, a_m)$ indicating that table Tbl contains the element (a_1, \dots, a_m) .
- $\mu\mathcal{E}v$ is a sequence representing the events executed so far. Each element of the sequence is of the form $(\mu, \tilde{a}) : e(a_1, \dots, a_m)$, meaning that the event $e(a_1, \dots, a_m)$ has been executed at program point μ with replication indices evaluating to \tilde{a} .

We define $\mathcal{E}v = \text{removepp}(\mu\mathcal{E}v) = [e(a_1, \dots, a_m) \mid (\mu, \tilde{a}) : e(a_1, \dots, a_m) \in \mu\mathcal{E}v]$ to be the sequence of events $\mu\mathcal{E}v$ without their associated program points and replication indices.

In addition to the grammar given in Figure 1, the terms M of the semantics can be values a and abort event values `event_abort` $(\mu, \tilde{a}) : e$, and the oracle bodies P can be `abort`, corresponding to the situation in which the game has been aborted. These additional terms and oracle bodies are not tagged with program points. (They do not occur in the initial oracle definition.)

The semantics is defined by reduction rules of the form $E, St, Q, Or, T, \mu\mathcal{E}v \xrightarrow{p}_t E', St', Q', Or', T', \mu\mathcal{E}v'$ meaning that $E, St, Q, Or, T, \mu\mathcal{E}v$ reduces to $E', St', Q', Or', T', \mu\mathcal{E}v'$ 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:

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

The probability of a trace $Tr = E_1, St_1, Q_1, Or_1, T_1, \mu\mathcal{E}v_1 \xrightarrow{p_1}_{t_1} \dots \xrightarrow{p_{m-1}}_{t_{m-1}} E_m, St_m, Q_m, Or_m, T_m, \mu\mathcal{E}v_m$ is $\Pr[Tr] = p_1 \times \dots \times p_{m-1}$. 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 5 and 6, we define an auxiliary relation for evaluating terms: $E, \sigma, M, T, \mu\mathcal{E}v \xrightarrow{p}_t E', \sigma, M', T', \mu\mathcal{E}v'$ means that the term M reduces to M' in environment E with the replication indices defined by σ , the table contents T , and the sequence of events $\mu\mathcal{E}v$, with probability p . 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) 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 $x[\sigma(\tilde{i})]$. Rule (IfT1) evaluates the `then` branch of `if` when the condition is true, and Rule (IfT2) evaluates the `else` branch otherwise.

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 an `event_abort` e or a `find[unique_e]`), the whole `find` executes the same event; in case the evaluations of the conditions execute several different events, one of them is chosen randomly, according to distribution $D_{\text{find}}(S)$, that is, almost uniformly over the choices of branches and indices (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 `[unique_e]`, 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 `[unique_e]` and S has a single element, we do the same. If the `find` is marked `[unique_e]` 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). If an element of the `defined` condition D is not defined, then 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

$$\begin{array}{c}
\frac{}{E, \sigma, \mu i, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, \sigma(i), \mathcal{T}, \mu \mathcal{E}v} \quad (\text{ReplIndex}) \\
\frac{x[a_1, \dots, a_m] \in \text{Dom}(E)}{E, \sigma, \mu x[a_1, \dots, a_m], \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, E(x[a_1, \dots, a_m]), \mathcal{T}, \mu \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, \mu f(a_1, \dots, a_m), \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, a, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{Fun}) \\
\frac{a \in T \quad E' = E[x[\sigma(\tilde{i})] \mapsto a]}{E, \sigma, \mu x[\tilde{i}] \stackrel{R}{\leftarrow} T; N, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{D_T(a)}_{N(a)} E', \sigma, N, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{NewT}) \\
\frac{a \in T \quad E' = E[x[\sigma(\tilde{i})] \mapsto a]}{E, \sigma, \mu \text{let } x[\tilde{i}] : T = a \text{ in } N, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E', \sigma, N, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{LetT}) \\
\frac{}{E, \sigma, \mu \text{if true then } N \text{ else } N', \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, N, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{IfT1}) \\
\frac{a \neq \text{true}}{E, \sigma, \mu \text{if } a \text{ then } N \text{ else } N', \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, N', \mathcal{T}, \mu \mathcal{E}v} \quad (\text{IfT2}) \\
\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}, \mu \mathcal{E}v \xrightarrow{p_k^*}_{t_k} E_k, \sigma_k, r_k, \mathcal{T}, \mu \mathcal{E}v \\
\text{where } v_k = (j, a_1, \dots, a_{m_j}) \text{ and } r_k \text{ is a value or } \text{event_abort}(\mu', \tilde{a}) : e \\
S = \{k \mid \exists(\mu', \tilde{a}, e), r_k = \text{event_abort}(\mu', \tilde{a}) : e\} \quad k_0 \in S
\end{array} \\
\frac{E, \sigma, \mu \text{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} \text{ suchthat} \\
D_j \wedge M_j \text{ then } N_j) \text{ else } N, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{p_1 \dots p_l D_{\text{find}}(S)(k_0)}_{t_1 \dots t_l FE(k_0)} E_{k_0}, \sigma_{k_0}, r_{k_0}, \mathcal{T}, \mu \mathcal{E}v}{\quad} \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}, \mu \mathcal{E}v \xrightarrow{p_k^*}_{t_k} E'', \sigma', r_k, \mathcal{T}, \mu \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 } [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} \\
\frac{E, \sigma, \mu \text{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} \text{ suchthat} \\
D_j \wedge M_j \text{ then } N_j) \text{ else } N, \mathcal{T}, \mu \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}, \mu \mathcal{E}v}{\quad} \quad (\text{FindT1}) \\
\frac{\text{First four lines as in (FindT1)} \quad S = \{v_k \mid r_k = \text{true}\} = \emptyset}{E, \sigma, \mu \text{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} \text{ suchthat} \\
D_j \wedge M_j \text{ then } N_j) \text{ else } N, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{p_1 \dots p_l}_{t_1 \dots t_l F2} E, \sigma, N, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{FindT2}) \\
\frac{\text{First four lines as in (FindT1)} \quad S = \{v_k \mid r_k = \text{true}\} \quad |S| > 1}{E, \sigma, \mu \text{find}[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}, \mu \mathcal{E}v \xrightarrow{p_1 \dots p_l}_{t_1 \dots t_l F3} E, \sigma, \text{event_abort}(\mu, \text{Im}(\sigma)) : e, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{FindT3})
\end{array}$$

Figure 5: Semantics (1): terms, first part

$$\begin{array}{c}
E, \sigma, \mu \text{insert } Tbl(a_1, \dots, a_l); N, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, N, (\mathcal{T}, Tbl(a_1, \dots, a_l)), \mu \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)] \\
\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}, \mu \mathcal{E}v \xrightarrow{pk}_{tk}^* E_k, \sigma_k, r_k, \mathcal{T}, \mu \mathcal{E}v \\
\text{where } v_k = Tbl(a_1, \dots, a_l) \text{ and } r_k \text{ is a value or event_abort } (\mu', \tilde{a}) : e \\
S = \{k \in [1, m] \mid \exists (\mu', \tilde{a}, e), r_k = \text{event_abort } (\mu', \tilde{a}) : e\} \quad k_0 \in S
\end{array}
}{
E, \sigma, \mu \text{get}[\text{unique?}] 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}, \mu \mathcal{E}v \quad (\text{GetTE}) \\
\frac{
p_1 \dots p_m D_{\text{get}(S)}(k_0) \xrightarrow{t_1 \dots t_m} GE(k_0) E_{k_0}, \sigma_{k_0}, r_{k_0}, \mathcal{T}, \mu \mathcal{E}v
}{
}
}
\\
\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}, \mu \mathcal{E}v \xrightarrow{pk}_{tk}^* E'', \sigma, r_k, \mathcal{T}, \mu \mathcal{E}v \\
\text{where } v_k = Tbl(a_1, \dots, a_l) \text{ and } r_k \text{ is a value} \\
S = \{k \in [1, m] \mid r_k = \text{true}\} \\
|S| = 1 \text{ or } [\text{unique?}] \text{ is empty} \\
k_0 \in S \quad Tbl(a_1, \dots, a_l) = v_{k_0} \quad E' = E[x_1[\sigma(\tilde{i})] \mapsto a_1, \dots, x_l[\sigma(\tilde{i})] \mapsto a_l]
\end{array}
}{
E, \sigma, \mu \text{get}[\text{unique?}] 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}, \mu \mathcal{E}v \quad (\text{GetT1}) \\
\frac{
p_1 \dots p_m D_{\text{get}(S)}(k_0) \xrightarrow{t_1 \dots t_m} G1(k_0) E', \sigma, N, \mathcal{T}, \mu \mathcal{E}v
}{
}
}
\\
\frac{
\text{First four lines as in (GetT1)} \quad S = \emptyset
}{
E, \sigma, \mu \text{get}[\text{unique?}] 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}, \mu \mathcal{E}v \quad (\text{GetT2}) \\
\frac{
p_1 \dots p_m \xrightarrow{t_1 \dots t_m} G2 E, \sigma, N', \mathcal{T}, \mu \mathcal{E}v
}{
}
}
\\
\frac{
\text{First four lines as in (GetT1)} \quad |S| > 1
}{
E, \sigma, \mu \text{get}[\text{unique}_e] 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}, \mu \mathcal{E}v \quad (\text{GetT3}) \\
\frac{
p_1 \dots p_m \xrightarrow{t_1 \dots t_m} G3 E, \sigma, \text{event_abort } (\mu, \text{Im}(\sigma)) : e, \mathcal{T}, \mu \mathcal{E}v
}{
}
}
\\
E, \sigma, \mu \text{event } e(a_1, \dots, a_l); N, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, N, \mathcal{T}, (\mu \mathcal{E}v, (\mu, \text{Im}(\sigma)) : e(a_1, \dots, a_l)) \quad (\text{EventT}) \\
E, \sigma, \mu \text{event_abort } e, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, \text{event_abort } (\mu, \text{Im}(\sigma)) : e, \mathcal{T}, \mu \mathcal{E}v \quad (\text{EventAbortT}) \\
\frac{
E, \sigma, N, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{p}_t E', \sigma', N', \mathcal{T}', \mu \mathcal{E}v'
}{
E, \sigma, C[N], \mathcal{T}, \mu \mathcal{E}v \xrightarrow{p}_t E', \sigma', C[N'], \mathcal{T}', \mu \mathcal{E}v' \quad (\text{CtxT})
} \\
E, \sigma, C[\text{event_abort } (\mu, \tilde{a}) : e], \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, \text{event_abort } (\mu, \tilde{a}) : e, \mathcal{T}, \mu \mathcal{E}v \quad (\text{CtxEventT}) \\
\frac{
\neg \forall j \leq l, \exists a_j, E, \sigma, M_j, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1}^* E, \sigma, a_j, \mathcal{T}, \mu \mathcal{E}v
}{
E, \sigma, \text{defined}(M_1, \dots, M_l) \wedge M, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, \text{false}, \mathcal{T}, \mu \mathcal{E}v \quad (\text{DefinedNo})
} \\
\frac{
\forall j \leq l, \exists a_j, E, \sigma, M_j, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1}^* E, \sigma, a_j, \mathcal{T}, \mu \mathcal{E}v
}{
E, \sigma, \text{defined}(M_1, \dots, M_l) \wedge M, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, \sigma, M, \mathcal{T}, \mu \mathcal{E}v \quad (\text{DefinedYes})
}
\end{array}$$

Figure 6: Semantics (2): terms, second part, and defined conditions

$$\begin{aligned}
 C ::= & \mu x[a_1, \dots, a_{k-1}, [], M_{k+1}, \dots, M_m] \\
 & \mu f(a_1, \dots, a_{k-1}, [], M_{k+1}, \dots, M_m) \\
 & \mu \text{let } x[\tilde{i}] : T = [] \text{ in } N \\
 & \mu \text{if } [] \text{ then } N \text{ else } N' \\
 & \mu \text{event } e(a_1, \dots, a_{k-1}, [], M_{k+1}, \dots, M_l); N \\
 & \mu \text{insert } Tbl(a_1, \dots, a_{k-1}, [], M_{k+1}, \dots, M_l); N
 \end{aligned}$$

Figure 7: Term contexts

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) to (GetT3) 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 ; when several events may be executed, one of them is chosen randomly according to distribution $D_{\text{get}}(S)$, that is, almost uniformly in the elements of the table Tbl . Rules (GetT1), (GetT2), and (GetT3) compute the set S of elements of indices of elements of table Tbl in \mathcal{T} that satisfy condition M . If S is empty, we execute N' (Rule (GetT2)). When S is not empty, two cases can happen. If the `get` is not marked `[uniquee]` or S has a single element, then one of its elements is chosen randomly according to distribution $D_{\text{get}}(S)$, 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 (GetT1)). If the `get` is marked `[uniquee]` and S contains several elements, then we execute event e and abort (Rule (GetT3)). 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 (used for the initial game), and one without tables (used for the other games). The semantics of the calculus without tables can be defined without the component \mathcal{T} . We then need to relate the two semantics.

Rule (EventT) adds the executed event to $\mu\mathcal{E}v$. Rule (EventAbortT) executes `event_abort` e . The event e is not immediately added to $\mu\mathcal{E}v$, because for terms that occur in conditions of `find`, in case several branches execute `event_abort`, we may need to choose randomly which event will be added to $\mu\mathcal{E}v$. Hence, we use the result `event_abort` $(\mu, \tilde{a}) : e$ instead.

Rules (CtxT) and (CtxEventT) allow evaluating terms under a context. In these rules, C is an elementary context, of one of the forms defined in Figure 7. 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

$$\begin{array}{l}
\{(\sigma, {}^\mu 0)\} \uplus \mathcal{Q}, \mathcal{O}r \rightsquigarrow \mathcal{Q}, \mathcal{O}r \quad (\text{Nil}) \\
\{(\sigma, {}^\mu(Q_1 \mid Q_2))\} \uplus \mathcal{Q}, \mathcal{O}r \rightsquigarrow \{(\sigma, Q_1), (\sigma, Q_2)\} \uplus \mathcal{Q}, \mathcal{O}r \quad (\text{Par}) \\
\{(\sigma, {}^\mu \text{foreach } i \leq n \text{ do } Q)\} \uplus \mathcal{Q}, \mathcal{O}r \rightsquigarrow \{(\sigma[i \mapsto a], Q) \mid a \in [1, n]\} \uplus \mathcal{Q}, \mathcal{O}r \quad (\text{Repl}) \\
\frac{\mathcal{O}r = (\mathcal{O}r_{\text{pub}}, \mathcal{O}r_{\text{priv}}) \quad \mathcal{O}' \notin \mathcal{O}r_{\text{pub}} \cup \mathcal{O}r_{\text{priv}}}{\{(\sigma, {}^\mu \text{newOracle } \mathcal{O}; Q)\} \uplus \mathcal{Q}, \mathcal{O}r \rightsquigarrow \{(\sigma, Q\{\mathcal{O}'/\mathcal{O}\})\} \uplus \mathcal{Q}, (\mathcal{O}r_{\text{pub}}, \mathcal{O}r_{\text{priv}} \cup \{\mathcal{O}'\})} \quad (\text{NewOracle}) \\
\{(\sigma, {}^\mu \mathcal{O}[\tilde{i}](x[\tilde{i}] : T) := P)\} \uplus \mathcal{Q}, \mathcal{O}r \rightsquigarrow \{(\sigma, {}^\mu \mathcal{O}[\sigma(\tilde{i})](x[\tilde{i}] : T) := P)\} \uplus \mathcal{Q}, \mathcal{O}r \quad (\text{DefOracle}) \\
\text{reduce}(\mathcal{Q}, \mathcal{O}r) \text{ is the normal form of } \mathcal{Q}, \mathcal{O}r \text{ by } \rightsquigarrow
\end{array}$$

Figure 8: Semantics (3): oracle definitions

value a or an abort event value $\text{event_abort}(\mu, \tilde{a}) : e$. We have $E, \sigma, M, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1}^* E, \sigma, M, \mathcal{T}, \mu \mathcal{E}v$ and, if $E, \sigma, M, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mu \mathcal{E}v'$ and $E', \sigma', M', \mathcal{T}', \mu \mathcal{E}v' \xrightarrow{p'}_{t'} E'', \sigma'', M'', \mathcal{T}'', \mu \mathcal{E}v''$, then $E, \sigma, M, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{p \times p'}_{t, t'}^* E'', \sigma'', M'', \mathcal{T}'', \mu \mathcal{E}v''$: 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 8 defines the semantics of oracle definitions. We use an auxiliary reduction relation \rightsquigarrow , for reducing oracle definitions. This relation transforms configurations of the form $\mathcal{Q}, \mathcal{O}r$. Rule (Nil) removes nil definitions. Rules (Par) and (Repl) expand parallel compositions and replications, respectively. Rule (NewOracle) creates a new oracle and adds it to $\mathcal{O}r_{\text{priv}}$. Semantic configurations are considered equivalent modulo renaming of oracles in $\mathcal{O}r_{\text{priv}}$, so that a single semantic configuration is obtained after applying (NewOracle). Rule (DefOracle) evaluates the replication indices in an oracle declaration. The oracle itself is not executed: the call is done by the (OracleCall) rule. The relation \rightsquigarrow is convergent (confluent and terminating), so it has normal forms. Oracle definitions in \mathcal{Q} in configurations $E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v$ are always in normal form by \rightsquigarrow , so they always start with an oracle declaration.

Finally, Figures 9, 10, and 11 define the semantics of oracles bodies. Most of these rules are very similar to those for terms: they just use oracles instead of terms as continuations, and include a whole semantic configuration. Rule (EventAbort) executes event e and aborts the game, by reducing to the configuration with oracle body `abort`. Similarly to the case of terms, Rules (Ctx) and (CtxEvent) allow evaluating terms under a context inside an oracle body. In these rules, C is an elementary context, of one of the forms defined in Figure 12.

Rule (OracleCall) performs an oracle call: it selects an available oracle, and immediately executes it. (The oracle call blocks if no suitable oracle is available.) The head of the execution stack after this rule is the body of the called oracle. In this rule, S is a multiset. When we take probabilities over multisets, we consider that $D_{\text{call}}(S)(\sigma', Q_0)$ is the probability of choosing *one* of the elements equal to σ', Q_0 in S according to the distribution $D_{\text{call}}(S)$, so that the probability of choosing any element equal to (σ', Q_0) is in fact $S(\sigma', Q_0) \times D_{\text{call}}(S)(\sigma', Q_0)$.

The rule (Return) corresponds to an oracle body finishing its execution and returning the result to the caller, which is the next oracle P_0 in the execution stack. The oracle definitions that follow the return are added to the set of available oracles, after reducing them by rules of Figure 8. The in branch of the caller is then executed. The rule (Yield) corresponds to an oracle body finishing its execution by yield. The else branch of the caller is then executed.

Rule (AttIO) corresponds to the attacker receiving a return and then triggering a new oracle call. The probability is over all possible attacker choices, as well as the corresponding oracles.

$$\begin{array}{c}
\text{Same assumption as in (FindTE)} \quad r_{k_0} = \text{event_abort}(\mu', \tilde{a}) : e \\
\hline
E, (\sigma, \mu \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) :: St, Q, Or, T, \mu \mathcal{E}v \\
\frac{p_1 \dots p_{l_0} D_{\text{find}}(S)(k_0)}{\rightarrow_{t_1 \dots t_{l_0} FE(k_0)}} E_{k_0}, (\sigma_{k_0}, \text{abort}) :: St, Q, Or, T, (\mu \mathcal{E}v, (\mu', \tilde{a}) : e) \\
\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'}}] \\
\hline
E, (\sigma, \mu \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) :: St, Q, Or, T, \mu \mathcal{E}v \\
\frac{p_1 \dots p_l D_{\text{find}}(S)(v_0)}{\rightarrow_{t_1 \dots t_l F1(v_0)}} E', (\sigma, P_{j'}) :: St, Q, Or, T, \mu \mathcal{E}v \\
\text{(Find1)} \\
\hline
\text{First four lines as in (FindT1)} \quad S = \{v_k \mid r_k = \text{true}\} = \emptyset \\
\hline
E, (\sigma, \mu \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) :: St, Q, Or, T, \mu \mathcal{E}v \xrightarrow{p_1 \dots p_l}_{t_1 \dots t_l F2} E, (\sigma, P) :: St, Q, Or, T, \mu \mathcal{E}v \\
\text{(Find2)} \\
\hline
\text{First four lines as in (FindT1)} \quad S = \{v_k \mid r_k = \text{true}\} \quad |S| > 1 \\
\hline
E, (\sigma, \mu \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) :: St, Q, Or, T, \mu \mathcal{E}v \\
\frac{p_1 \dots p_l}{\rightarrow_{t_1 \dots t_l F3}} E, (\sigma, \text{abort}) :: St, Q, Or, T, (\mu \mathcal{E}v, (\mu, \text{Im}(\sigma)) : e) \\
\text{(Find3)} \\
\hline
E, (\sigma, \mu \text{insert } Tbl(a_1, \dots, a_l); P) :: St, Q, Or, T, \mu \mathcal{E}v \\
\stackrel{1}{\mapsto} E, (\sigma, P) :: St, Q, Or, (T, Tbl(a_1, \dots, a_l)), \mu \mathcal{E}v \\
\text{(Insert)} \\
\hline
\text{Same assumption as in (GetTE)} \quad r_{k_0} = \text{event_abort}(\mu', \tilde{a}) : e \\
\hline
E, (\sigma, \mu \text{get}[\text{unique?}] Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } P \text{ else } P') :: St, Q, Or, T, \mu \mathcal{E}v \\
\frac{p_1 \dots p_m D_{\text{get}}(S)(k_0)}{\rightarrow_{t_1 \dots t_m GE(k_0)}} E_{k_0}, (\sigma_{k_0}, \text{abort}) :: St, Q, Or, T, (\mu \mathcal{E}v, (\mu', \tilde{a}) : e) \\
\text{(GetE)} \\
\hline
\text{First four lines as in (GetT1)} \quad |S| = 1 \text{ or } [\text{unique?}] \text{ is empty} \\
k_0 \in S \quad Tbl(a_1, \dots, a_l) = v_{k_0} \quad E' = E[x_1[\sigma(\tilde{i})] \mapsto a_1, \dots, x_l[\sigma(\tilde{i})] \mapsto a_l] \\
\hline
E, (\sigma, \mu \text{get}[\text{unique?}] Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } P \text{ else } P') :: St, Q, Or, T, \mu \mathcal{E}v \\
\frac{p_1 \dots p_m D_{\text{get}}(S)(k_0)}{\rightarrow_{t_1 \dots t_m G1(k_0)}} E', (\sigma, P) :: St, Q, Or, T, \mu \mathcal{E}v \\
\text{(Get1)} \\
\hline
\text{First four lines as in (GetT1)} \quad S = \emptyset \\
\hline
E, (\sigma, \mu \text{get}[\text{unique?}] Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } P \text{ else } P') :: St, Q, Or, T, \mu \mathcal{E}v \\
\frac{p_1 \dots p_m}{\rightarrow_{t_1 \dots t_m G2}} E, (\sigma, P') :: St, Q, Or, T, \mu \mathcal{E}v \\
\text{(Get2)} \\
\hline
\text{First four lines as in (GetT1)} \quad |S| > 1 \\
\hline
E, (\sigma, \mu \text{get}[\text{unique}_e] Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } M \text{ in } P \text{ else } P') :: St, Q, Or, T, \mu \mathcal{E}v \\
\frac{p_1 \dots p_m}{\rightarrow_{t_1 \dots t_m G3}} E, (\sigma, \text{abort}) :: St, Q, Or, T, (\mu \mathcal{E}v, (\mu, \text{Im}(\sigma)) : e) \\
\text{(Get3)}
\end{array}$$

Figure 9: Semantics (4): oracle bodies, first part

$$\begin{array}{c}
\frac{a \in T \quad E' = E[x[\sigma(\tilde{i})] \mapsto a]}{E, (\sigma, \mu x[\tilde{i}] \stackrel{R}{\leftarrow} T; P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{D_T(a)}_{N(a)} E', (\sigma, P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{New}) \\
\frac{a \in T \quad E' = E[x[\sigma(\tilde{i})] \mapsto a]}{E, (\sigma, \mu \text{let } x[\tilde{i}] : T = a \text{ in } P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E', (\sigma, P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{Let}) \\
E, (\sigma, \mu \text{if true then } P \text{ else } P') :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, (\sigma, P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \quad (\text{If1}) \\
\frac{a \neq \text{true}}{E, (\sigma, \mu \text{if } a \text{ then } P \text{ else } P') :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, (\sigma, P') :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{If2}) \\
\frac{E, (\sigma, \mu \text{event } e(a_1, \dots, a_l); P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, (\sigma, P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, (\mu \mathcal{E}v, (\mu, \text{Im}(\sigma)) : e(a_1, \dots, a_l))}{E, (\sigma, \mu \text{event_abort } e) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, (\sigma, \text{abort}) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, (\mu \mathcal{E}v, (\mu, \text{Im}(\sigma)) : e)} \quad (\text{Event}) \\
\text{(EventAbort)} \\
\frac{E, \sigma, N, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{P} E', \sigma', N', \mathcal{T}', \mu \mathcal{E}v'}{E, (\sigma, C[N]) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{P} E', (\sigma', C[N']) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}', \mu \mathcal{E}v'} \quad (\text{Ctx}) \\
\frac{E, (\sigma, C[\text{event_abort } (\mu, \tilde{a}) : e]) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, (\sigma, \text{abort}) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, (\mu \mathcal{E}v, (\mu, \tilde{a}) : e)}{\text{(CtxEvent)}} \\
\frac{P_0 = \text{let } x[\tilde{i}] : T = O[a_1, \dots, a_l, ?u_{l+1}[\tilde{i}] \leq n_{l+1}, \dots, ?u_k[\tilde{i}] \leq n_k](b) \text{ in } P_1 \text{ else } P_2 \\ S = \{(\sigma', Q) \in \mathcal{Q} \mid Q = \mu'' O[a_1, \dots, a_l, c_{l+1}, \dots, c_k](x'[\tilde{i}] : T') := P' \text{ and } b \in T' \\ \text{for some } \mu'', \sigma', x', T', P', c_{l+1} \in [1, n_{l+1}], \dots, c_k \in [1, n_k]\} \\ (\sigma', Q_0) \in S \quad Q_0 = \mu' O[a_1, \dots, a_k](x'[\tilde{i}] : T') := P' \\ P'_0 = \text{let } x[\tilde{i}] : T = O[a_1, \dots, a_k](b) \text{ in } P_1 \text{ else } P_2}{E, (\sigma, \mu P_0) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{S(\sigma', Q_0) \times D_{\text{call}}(S)(\sigma', Q_0)}_{O(\sigma', Q_0)} E[x'[\sigma'(\tilde{i})] \mapsto b, u_{l+1}[\sigma(\tilde{i})] \mapsto c_{l+1}, \dots, u_k[\sigma(\tilde{i})] \mapsto c_k], (\sigma', P') :: (\sigma, \mu P'_0) :: St, \mathcal{Q} \uplus \{(\sigma', Q_0)\}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{OracleCall}) \\
\frac{d \in T \quad Q', \mathcal{O}r' = \text{reduce}(\{(\sigma, Q)\}, \mathcal{O}r) \\ P_0 = \text{let } x[\tilde{i}] : T = O[a_1, \dots, a_k](b) \text{ in } P_1 \text{ else } P_2}{E, (\sigma, \mu' \text{return } (d); Q) :: (\sigma', \mu P_0) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E[x[\sigma'(\tilde{i})] \mapsto d], (\sigma', P_1) :: St, \mathcal{Q} \uplus Q', \mathcal{O}r', \mathcal{T}, \mu \mathcal{E}v} \quad (\text{Return}) \\
\frac{P_0 = \text{let } x[\tilde{i}] : T = O[a_1, \dots, a_k](b) \text{ in } P_1 \text{ else } P_2}{E, (\sigma, \mu' \text{yield}) :: (\sigma', \mu P_0) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{1} E, (\sigma', P_2) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v} \quad (\text{Yield})
\end{array}$$

Figure 10: Semantics (5): oracle bodies, second part

$$\begin{array}{c}
 \mathcal{Q}, \mathcal{O}r' = \text{reduce}(\{(\sigma, Q)\}, \mathcal{O}r) \\
 S = \{(\sigma', Q) \in \mathcal{Q} \uplus \mathcal{Q}' \mid Q = \mu'' O[a_1, \dots, a_k](x'[\tilde{i}] : T') := P' \text{ and } b \in T' \\
 \quad \text{for some } \mu'', \sigma', x', \tilde{i}, T', P'\} \\
 (\sigma', Q_0) \in S \quad Q_0 = \mu' O[a_1, \dots, a_k](x[\tilde{i}] : T) := P' \\
 \mathcal{A}(d) = D_{\mathcal{A}}^d \quad \mathcal{A}_0, m_0 \stackrel{R}{\leftarrow} D_{\mathcal{A}}^d \quad m_0 = (O, a_1, \dots, a_k, b) \\
 \mathcal{O}r' = (\mathcal{O}r'_{\text{pub}}, \mathcal{O}r'_{\text{priv}}) \quad O \notin \mathcal{O}r'_{\text{priv}} \\
 \hline
 E, (\sigma, \mu \text{return } (d); Q) :: \mathcal{A}, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{S(\sigma', Q_0) \times D_{\text{call}}(S)(\sigma', Q_0) \times D_{\mathcal{A}}^d(\mathcal{A}_0, m_0)} \mathcal{O}(\sigma', Q_0, \mathcal{A}_0, a_1, \dots, a_k, b) \\
 E[x[\sigma'(\tilde{i})] \mapsto b], (\sigma', P') :: \mathcal{A}, \mathcal{Q} \uplus \mathcal{Q}' \uplus \{(\sigma', Q_0)\}, \mathcal{O}r', \mathcal{T}, \mu \mathcal{E}v \\
 \text{(AttIO)} \\
 \\
 S = \{(\sigma', Q) \in \mathcal{Q} \mid Q = \mu'' O[a_1, \dots, a_k](x'[\tilde{i}] : T') := P' \text{ and } b \in T' \text{ for some } \mu'', \sigma', x', \tilde{i}, T', P'\} \\
 (\sigma', Q_0) \in S \quad Q_0 = \mu' O[a_1, \dots, a_k](x[\tilde{i}] : T) := P' \\
 \mathcal{A}() = D_{\mathcal{A}}^\emptyset \quad \mathcal{A}_0, m_0 \stackrel{R}{\leftarrow} D_{\mathcal{A}}^\emptyset \quad m_0 = (O, a_1, \dots, a_k, b) \\
 \mathcal{O}r = (\mathcal{O}r_{\text{pub}}, \mathcal{O}r_{\text{priv}}) \quad O \notin \mathcal{O}r_{\text{priv}} \\
 \hline
 E, (\sigma, \mu \text{yield}) :: \mathcal{A}, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{S(\sigma', Q_0) \times D_{\text{call}}(S)(\sigma', Q_0) \times D_{\mathcal{A}}^\emptyset(\mathcal{A}_0, m_0)} \mathcal{O}(\sigma', Q_0, \mathcal{A}_0, a_1, \dots, a_k, b) \\
 E[x[\sigma'(\tilde{i})] \mapsto b], (\sigma', P') :: \mathcal{A}, \mathcal{Q} \uplus \{(\sigma', Q_0)\}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v \\
 \text{(AttYield)}
 \end{array}$$

Figure 11: Semantics (6): oracle bodies, third part, attackers interactions

We later show in Section 2.4.4 that there is in fact always at most one such possible oracle, and that the probability only depends on attacker choices. Rule (AttYield) handles the case of a yield returned to the attacker, where we then call the attacker with an empty input. This call is disjoint from any other call to the attacker that could be performed through rule (AttIO): it is assumed that the efficient bijection from types to concrete bitstrings does not return the empty string.

After finishing execution of an oracle definition, the system produces as a result the sequence of executed events $\mu \mathcal{E}v$. Looking forward, recall that an adversary will always be made of an attacker \mathcal{A} and an adversary context, this context thus allowing the adversary to raise arbitrary values in the sequence of events.

These events can be used to distinguish games, so we introduce an additional algorithm, a *distinguisher* D that takes as input a sequence of events $\mathcal{E}v$ (without program points and replication indices) and returns true or false.

An example of distinguisher is D_e defined by $D_e(\mathcal{E}v) = \text{true}$ 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 $(D \vee D')(\mathcal{E}v) = D(\mathcal{E}v) \vee D'(\mathcal{E}v)$, $(D \wedge D')(\mathcal{E}v) = D(\mathcal{E}v) \wedge D'(\mathcal{E}v)$, and $(\neg D)(\mathcal{E}v) = \neg D(\mathcal{E}v)$. We denote by $\text{Pr}[Q : D]$ the probability that Q executes a sequence of events $\mathcal{E}v$ such that $D(\mathcal{E}v) = \text{true}$. This is formally defined as follows.

Definition 4 The initial configuration for running oracle definition Q against attacker \mathcal{A} is $\text{initConfig}(Q, \mathcal{A}) = \emptyset, (\sigma_0, \text{return } (start)) :: \mathcal{A}, \mathcal{Q}, \mathcal{O}r, \emptyset, \emptyset$ where $\mathcal{Q}, \mathcal{O}r = \text{reduce}(\{(\sigma_0, Q)\}, (\text{fo}(Q), \emptyset))$ and σ_0 is the empty mapping sequence. A *trace* of Q for attacker \mathcal{A} is a trace that starts from $\text{initConfig}(Q, \mathcal{A})$: $Tr = \text{initConfig}(Q, \mathcal{A}) \xrightarrow{p_1} t_1 \dots \xrightarrow{p_{m-1}} t_{m-1} \text{ Conf}_m$. Let $\mathcal{T}r^{\mathcal{A}, Q}$ be the set of all traces of Q for an attacker $\mathcal{A} \in \mathcal{B}\mathcal{B}$.

$$\begin{aligned}
C ::= & \text{ }^\mu\text{let } x[\tilde{i}] : T = [] \text{ in } P \\
& \text{ }^\mu\text{if } [] \text{ then } P \text{ else } P' \\
& \text{ }^\mu\text{return } ([]) ; Q \\
& \text{ }^\mu\text{let } x[\tilde{i}] : T = O[a_1, \dots, a_{k-1}, [], M_{k+1}, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N) \text{ in } P \\
& \text{ else } P' \\
& \text{ }^\mu\text{let } x[\tilde{i}] : T = O[a_1, \dots, a_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m]([]) \text{ in } P \text{ else } P' \\
& \text{ }^\mu\text{insert } Tbl(a_1, \dots, a_{k-1}, [], M_{k+1}, \dots, M_l) ; P \\
& \text{ }^\mu\text{event } e(a_1, \dots, a_{k-1}, [], M_{k+1}, \dots, M_l) ; P
\end{aligned}$$

Figure 12: Oracle contexts

Let $\mathcal{T}_{\text{full}}^{\mathcal{A}, Q}$ be the set of *full traces* of Q , that is, the set of traces of $\mathcal{T}^{\mathcal{A}, Q}$ whose last configuration Conf_m cannot be reduced.

A trace Tr' is an *extension* of Tr when Tr' is obtained by continuing execution from the last configuration of Tr . Equivalently, Tr is a *prefix* of Tr' .

Let φ be a property of traces, that is, a function from traces to $\{\text{true}, \text{false}\}$. We say that Tr satisfies φ , and we write $Tr \vdash \varphi$, when $\varphi(Tr) = \text{true}$.

A property φ is *preserved by extension* when for all traces Tr such that $Tr \vdash \varphi$, for all extensions Tr' of Tr , $Tr' \vdash \varphi$.

Given a trace $Tr = \text{initConfig}(Q, \mathcal{A}) \xrightarrow{p_1}_{t_1} \dots \xrightarrow{p_{m-1}}_{t_{m-1}} \text{Conf}_m$, recall that $\Pr[Tr] = p_1 \times \dots \times p_{m-1}$. We define

$$\begin{aligned}
\Pr^{\mathcal{A}}[Q : \varphi] &= \sum_{Tr \in \mathcal{T}_{\text{full}}^{\mathcal{A}, Q}, Tr \vdash \varphi} \Pr[Tr], \\
\Pr^{\mathcal{A}}[Q \preceq \varphi] &= \sum_{Tr \in \mathcal{T}_{\text{full}}^{\mathcal{A}, Q}, \exists Tr' \text{ prefix of } Tr, Tr' \vdash \varphi} \Pr[Tr] \\
&= \sum_{Tr \in \mathcal{T}^{\mathcal{A}, Q}, Tr \vdash \varphi, \text{for any strict prefix } Tr' \text{ of } Tr, Tr' \not\vdash \varphi} \Pr[Tr].
\end{aligned}$$

Given a distinguisher D , we can consider it as a property of traces by defining $Tr \vdash D$ if and only if $D(\text{removepp}(\mu\mathcal{E}v_m)) = \text{true}$ where the last configuration of Tr is $\text{Conf}_m = E_m, St_m, Q_m, Or_m, T_m, \mu\mathcal{E}v_m$. The function removepp guarantees that the pair (program point, replication indices) is not used in the evaluation of the distinguisher. Actually, this pair could be removed from the semantics. It is useful for the proof of injective correspondences (Section 4.2.4).

Lemma 1 1. $\Pr^{\mathcal{A}}[Q : \varphi] \leq \Pr^{\mathcal{A}}[Q \preceq \varphi]$.

2. If φ is preserved by extension, then $\Pr^{\mathcal{A}}[Q : \varphi] = \Pr^{\mathcal{A}}[Q \preceq \varphi]$.

Proof The first property holds because, when $Tr \vdash \varphi$, there exists a prefix Tr' of Tr (take $Tr' = Tr$) such that $Tr' \vdash \varphi$.

The second property holds because, if φ is preserved by extension and there exists a prefix Tr' of Tr such that $Tr' \vdash \varphi$, then $Tr \vdash \varphi$. \square

For simple terms M , the evaluation can be defined without tables and events. We define $E, \rho, M \Downarrow a$ if and only if $E, \rho, M, \emptyset, \emptyset \xrightarrow{1^*} E, \rho, a, \emptyset, \emptyset$, where the environment E gives the values of process variables (in particular arrays), and the environment ρ gives the values of replication indices and other variables (e.g., those used in correspondences, see Section 2.8.3). The evaluation relation $E, \rho, M \Downarrow a$ can also be defined by induction as follows:

$$\begin{array}{c} E, \rho, i \Downarrow \rho(i) \\ \frac{\forall j \leq m, E, \rho, M_j \Downarrow a_j}{E, \rho, x[M_1, \dots, M_m] \Downarrow E(x[a_1, \dots, a_m])} \\ \frac{\forall j \leq m, E, \rho, M_j \Downarrow a_j}{E, \rho, f(M_1, \dots, M_m) \Downarrow f(a_1, \dots, a_m)} \end{array}$$

For terms that do not contain process variables, the environment E can be omitted, and we write $\rho, M \Downarrow a$. It can also be defined by induction as follows:

$$\begin{array}{c} \rho, i \Downarrow \rho(i) \\ \frac{\forall j \leq m, \rho, M_j \Downarrow a_j}{\rho, f(M_1, \dots, M_m) \Downarrow f(a_1, \dots, a_m)} \end{array}$$

2.4.2 Properties

Given an oracle definition Q_0 , we write I_μ for the current replication indices at μ in Q_0 . For ${}^\mu\text{return}$ (*start*) in the initial configuration, we let $I_\mu = \emptyset$ be the empty sequence.

Lemma 2 *Let Tr be a trace of Q_0 for any $A \in \mathcal{BB}$. In the derivation of Tr , for all configurations $E, \sigma, {}^\mu M, \mathcal{T}, \mu\mathcal{E}v$, $\text{Dom}(\sigma) = I_\mu$; for all configurations $E, St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v$, for all $(\sigma, {}^\mu P) \in St$, $\text{Dom}(\sigma) = I_\mu$; and for all configurations \mathcal{Q}, Or or $E, St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v$, for all $(\sigma, {}^\mu Q) \in \mathcal{Q}$, $\text{Dom}(\sigma) = I_\mu$.*

Proof sketch We say that

- a configuration $E, \sigma, {}^\mu M, \mathcal{T}, \mu\mathcal{E}v$ is ok when $\text{Dom}(\sigma) = I_\mu$;
- a configuration \mathcal{Q}, Or is ok when for all $(\sigma, {}^\mu Q) \in \mathcal{Q}$, $\text{Dom}(\sigma) = I_\mu$;
- a configuration $E, St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v$ is ok when for all $(\sigma, {}^\mu P) \in St$, $\text{Dom}(\sigma) = I_\mu$ and for all $(\sigma, {}^\mu Q) \in \mathcal{Q}$, $\text{Dom}(\sigma) = I_\mu$.

We show by induction on the derivations that

1. if $E_1, \sigma_1, M_1, \mathcal{T}_1, \mu\mathcal{E}v_1$ is ok, then all configurations $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ in the derivation of $E_1, \sigma_1, M_1, \mathcal{T}_1, \mu\mathcal{E}v_1 \xrightarrow{p}_t E_2, \sigma_2, M_2, \mathcal{T}_2, \mu\mathcal{E}v_2$ are ok;
2. if \mathcal{Q}_1, Or_1 is ok and $\mathcal{Q}_1, Or_1 \rightsquigarrow \mathcal{Q}_2, Or_2$, then \mathcal{Q}_2, Or_2 is ok;
3. if \mathcal{Q}_1, Or_1 is ok, then $\text{reduce}(\mathcal{Q}_1, Or_1)$ is ok and all configurations \mathcal{Q}, Or and in the derivation of $\mathcal{Q}_1, Or_1 \rightsquigarrow^* \text{reduce}(\mathcal{Q}_1, Or_1)$ are ok;
4. if $E_1, St_1, \mathcal{Q}_1, Or_1, \mathcal{T}_1, \mu\mathcal{E}v_1$ is ok, then all configurations $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ or \mathcal{Q}, Or or $E, St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v$ in the derivation of $E_1, St_1, \mathcal{Q}_1, Or_1, \mathcal{T}_1, \mu\mathcal{E}v_1 \xrightarrow{p}_t E_2, St_2, \mathcal{Q}_2, Or_2, \mathcal{T}_2, \mu\mathcal{E}v_2$ are ok.

Indeed, the changes in σ match the definition of replication indices.

For Property 1, in rules for `find`, the indices of the `find` condition are added to the domain of σ when evaluating the condition and in all other cases, σ is unchanged.

In the proof of Property 2, in (`Repl`), the replication index i is added to σ and in all other cases, σ is unchanged.

Property 3 follows immediately from Property 2 by induction.

For Property 4, in rules for `find`, the indices of the `find` condition are added to the domain of σ when evaluating the condition, as in Property 1; in (`Return`), we use Property 3. In all other cases, σ is unchanged and remains associated with the same program point or a program point with the same replication indices. We use Property 1 when evaluating terms, in conditions of `find` and `get` and in (`Ctx`).

Moreover, in the computation of `initConfig`(Q_0, \mathcal{A}), the configuration $\{(\sigma_0, Q_0)\}, (\text{fo}(Q), \emptyset)$ is ok (the current replication indices at the root of Q_0 are empty), so by Property 3, `initConfig`(Q_0, \mathcal{A}) is ok. \square

Lemma 3 *If $E, \sigma, N, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', \sigma', N', \mathcal{T}', \mu\mathcal{E}v'$, then E' is an extension of E , \mathcal{T} is a prefix of \mathcal{T}' , $\mu\mathcal{E}v$ is a prefix of $\mu\mathcal{E}v'$, σ' is an extension of σ , and if the term N' is not of the form $C_1[\dots C_k[\text{event_abort}(\mu, \tilde{a}) : e] \dots]$ for some $k \in \mathbb{N}$ and C_1, \dots, C_k contexts defined in Figure 7, then $\sigma' = \sigma$. For all configurations $E'', \sigma'', N'', \mathcal{T}'', \mu\mathcal{E}v''$ in the derivation of $E, \sigma, N, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', \sigma', N', \mathcal{T}', \mu\mathcal{E}v'$, we have that E'' is an extension of E and, if $E'', \sigma'', N'', \mathcal{T}'', \mu\mathcal{E}v''$ is not in the derivation of an hypothesis of a rule for `find` or `get`, then E' is an extension of E'' ; σ'' is an extension of σ ; \mathcal{T} is a prefix of \mathcal{T}'' , which is a prefix of \mathcal{T}' ; and $\mu\mathcal{E}v$ is a prefix of $\mu\mathcal{E}v''$, which is a prefix of $\mu\mathcal{E}v'$.*

If $\mathcal{Q}, (\mathcal{O}r_{\text{pub}}, \mathcal{O}r_{\text{priv}}) \rightsquigarrow \mathcal{Q}', (\mathcal{O}r'_{\text{pub}}, \mathcal{O}r'_{\text{priv}})$, then $\mathcal{O}r'_{\text{pub}} = \mathcal{O}r_{\text{pub}}$ and $\mathcal{O}r_{\text{priv}} \subseteq \mathcal{O}r'_{\text{priv}}$.

If $E, (\sigma, P) :: St, \mathcal{Q}, (\mathcal{O}r_{\text{pub}}, \mathcal{O}r_{\text{priv}}), \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', (\sigma', P') :: St', \mathcal{Q}', (\mathcal{O}r'_{\text{pub}}, \mathcal{O}r'_{\text{priv}}), \mathcal{T}', \mu\mathcal{E}v'$, then E' is an extension of E , \mathcal{T} is a prefix of \mathcal{T}' , $\mu\mathcal{E}v$ is a prefix of $\mu\mathcal{E}v'$, $\mathcal{O}r_{\text{pub}} = \mathcal{O}r'_{\text{pub}}$ and $\mathcal{O}r_{\text{priv}} \subseteq \mathcal{O}r'_{\text{priv}}$.

- *If $E, (\sigma, P) :: St, \mathcal{Q}, (\mathcal{O}r_{\text{pub}}, \mathcal{O}r_{\text{priv}}), \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', (\sigma', P') :: St', \mathcal{Q}', (\mathcal{O}r'_{\text{pub}}, \mathcal{O}r'_{\text{priv}}), \mathcal{T}', \mu\mathcal{E}v'$ is derived by (`Return`) or (`AttIO`), then $\mathcal{T}' = \mathcal{T}$, $\mu\mathcal{E}v' = \mu\mathcal{E}v$, and for all configurations $\mathcal{Q}'', \mathcal{O}r''$ in that derivation, $\mathcal{O}r_{\text{priv}} \subseteq \mathcal{O}r''_{\text{priv}} \subseteq \mathcal{O}r'_{\text{priv}}$ and $\mathcal{O}r_{\text{pub}} = \mathcal{O}r''_{\text{pub}} = \mathcal{O}r'_{\text{pub}}$.*
- *If $E, (\sigma, P) :: St, \mathcal{Q}, (\mathcal{O}r_{\text{pub}}, \mathcal{O}r_{\text{priv}}), \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', (\sigma', P') :: St', \mathcal{Q}', (\mathcal{O}r'_{\text{pub}}, \mathcal{O}r'_{\text{priv}}), \mathcal{T}', \mu\mathcal{E}v'$ is derived by (`Yield`), (`OracleCall`), or (`AttYield`), then $\mathcal{T}' = \mathcal{T}$, $\mu\mathcal{E}v' = \mu\mathcal{E}v$, $\mathcal{O}r_{\text{priv}} = \mathcal{O}r'_{\text{priv}}$, and $\mathcal{O}r_{\text{pub}} = \mathcal{O}r'_{\text{pub}}$.*
- *In all other cases, $\mathcal{O}r'_{\text{pub}} = \mathcal{O}r_{\text{pub}}$, $\mathcal{O}r'_{\text{priv}} = \mathcal{O}r_{\text{priv}}$, $St = St'$, σ' is an extension of σ , and if the oracle body P' is not `abort` or of the form $C_0[C_1[\dots C_k[\text{event_abort}(\mu, \tilde{a}) : e] \dots]]$ for some C_0 context defined in Figure 12, $k \in \mathbb{N}$, and C_1, \dots, C_k contexts defined in Figure 7, then $\sigma' = \sigma$. For all configurations $E'', \sigma'', N'', \mathcal{T}'', \mu\mathcal{E}v''$ in the derivation of $E, (\sigma, P) :: St, \mathcal{Q}, (\mathcal{O}r_{\text{pub}}, \mathcal{O}r_{\text{priv}}), \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', (\sigma', P') :: St', \mathcal{Q}', (\mathcal{O}r'_{\text{pub}}, \mathcal{O}r'_{\text{priv}}), \mathcal{T}', \mu\mathcal{E}v'$, we have that E'' is an extension of E and, if $E'', \sigma'', N'', \mathcal{T}'', \mu\mathcal{E}v''$ is not in the derivation of an hypothesis of a rule for `find` or `get`, then E' is an extension of E'' ; σ'' is an extension of σ ; \mathcal{T} is a prefix of \mathcal{T}'' , which is a prefix of \mathcal{T}' ; and $\mu\mathcal{E}v$ is a prefix of $\mu\mathcal{E}v''$, which is a prefix of $\mu\mathcal{E}v'$.*

Proof sketch By induction on the derivation of $E, \sigma, N, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', \sigma', N', \mathcal{T}', \mu\mathcal{E}v'$ and by cases on the reductions $\mathcal{Q}, \mathcal{O}r \rightsquigarrow \mathcal{Q}', \mathcal{O}r'$ and $E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', St', \mathcal{Q}', \mathcal{O}r', \mathcal{T}', \mu\mathcal{E}v'$. \square

We say that a term N is *in evaluation position* in a configuration $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ when $M = C_1[\dots C_k[N]\dots]$ for some $k \in \mathbb{N}$ and C_1, \dots, C_k contexts defined in Figure 7.

We say that a term N is *in evaluation position* in a configuration $E, (\sigma, P) :: \mathcal{S}t, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ when $P = C_0[C_1[\dots C_k[N]\dots]]$ for some C_0 context defined in Figure 12, $k \in \mathbb{N}$, and C_1, \dots, C_k contexts defined in Figure 7, or $(\sigma', \text{let } x[\tilde{i}] : T = O[a_1, \dots, a_k](b) \text{ in } P_1 \text{ else } P_2) \in \mathcal{S}t$ and N is a_1, \dots, a_k , or b .

The oracle bodies P such that $(\sigma, P) \in \mathcal{S}t$ for some σ are *in evaluation position* in configuration $E, \mathcal{S}t, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$. The oracle definitions Q such that $(\sigma, Q) \in \mathcal{Q}$ for some σ are *in evaluation position* in configurations $\mathcal{Q}, \mathcal{O}r$ and $E, \mathcal{S}t, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$.

Lemma 4 Consider a trace Tr of Q_0 for any $\mathcal{A} \in \mathcal{B}\mathcal{B}$.

All subterms of M that occur in non-evaluation position in a configuration $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ in the derivation of Tr are subterms of Q_0 .

All subterms and subprocesses of oracle definitions in \mathcal{Q} that occur in non-evaluation position in a configuration $\mathcal{Q}, \mathcal{O}r$ in the derivation of Tr are subterms, resp. subprocesses, of Q_0 up to renaming of oracles. All subterms and subprocesses of $\mathcal{S}t$ and of \mathcal{Q} that occur in non-evaluation position in a configuration $E, \mathcal{S}t, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ in the derivation of Tr are subterms, resp. subprocesses, of Q_0 up to renaming of oracles (except for the process 0 that follows return (start) in $\text{initConfig}(Q_0, \mathcal{A})$).

Proof We say that

- a configuration $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ is ok when all subterms of M that are not in evaluation position are subterms of Q_0 ;
- a configuration $\mathcal{Q}, \mathcal{O}r$ is ok when all subterms and subprocesses of \mathcal{Q} that are not in evaluation position are subterms, resp. subprocesses, of Q_0 up to renaming of oracles;
- a configuration $E, \mathcal{S}t, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ is ok when all subterms and subprocesses of $\mathcal{S}t$ and of \mathcal{Q} that are not in evaluation position are subterms, resp. subprocesses, of Q_0 up to renaming of oracles.

We show by induction on the derivations that

1. if $E_1, \sigma_1, M_1, \mathcal{T}_1, \mu\mathcal{E}v_1$ is ok, then all configurations $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ in the derivation of $E_1, \sigma_1, M_1, \mathcal{T}_1, \mu\mathcal{E}v_1 \xrightarrow{p}_t E_2, \sigma_2, M_2, \mathcal{T}_2, \mu\mathcal{E}v_2$ are ok;
2. if $\mathcal{Q}_1, \mathcal{O}r_1$ is ok and $\mathcal{Q}_1, \mathcal{O}r_1 \rightsquigarrow \mathcal{Q}_2, \mathcal{O}r_2$, then $\mathcal{Q}_2, \mathcal{O}r_2$ is ok;
3. if $\mathcal{Q}_1, \mathcal{O}r_1$ is ok, then $\text{reduce}(\mathcal{Q}_1, \mathcal{O}r_1)$ is ok and all configurations $\mathcal{Q}, \mathcal{O}r$ in the derivation of $\mathcal{Q}_1, \mathcal{O}r_1 \rightsquigarrow^* \text{reduce}(\mathcal{Q}_1, \mathcal{O}r_1)$ are ok;
4. if $E_1, \mathcal{S}t_1, \mathcal{Q}_1, \mathcal{O}r_1, \mathcal{T}_1, \mu\mathcal{E}v_1$ is ok, then all configurations $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ or $\mathcal{Q}, \mathcal{O}r$ or $E, \mathcal{S}t, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ in the derivation of $E_1, \mathcal{S}t_1, \mathcal{Q}_1, \mathcal{O}r_1, \mathcal{T}_1, \mu\mathcal{E}v_1 \xrightarrow{p}_t E_2, \mathcal{S}t_2, \mathcal{Q}_2, \mathcal{O}r_2, \mathcal{T}_2, \mu\mathcal{E}v_2$ are ok.

Property 1: In (ReplIndex), (Var), (Fun), and (EventAbortT), all terms are in evaluation position, so all configurations are ok. In (NewT), (LetT), (IfT1), (IfT2), (InsertT), and (EventT), N (resp. N') is not in evaluation position, so by hypothesis it is a subterm of Q_0 . Therefore, the target configuration is ok. In the rules for find, the recursive calls are on $D_j \wedge M_j$ which is not in evaluation position, so it is a subterm of Q_0 . Therefore, the initial configurations of the recursive calls are ok, and we conclude for the configurations inside the recursive calls by induction hypothesis. The target configuration is ok because in (FindTE) and (FindT3), the resulting

term is in evaluation position (it has no subterm), and in (FindT1) and (FindT2), the resulting term is a term that is not evaluation position in the initial configuration, so it is a subterm of Q_0 . In (DefinedNo) and (DefinedYes), the terms M_1, \dots, M_l are not in evaluation position in the initial configuration, so they are subterms of Q_0 . Hence the initial configurations of the recursive calls are ok, and we conclude for the configurations inside the recursive calls by induction hypothesis. The target configuration is ok because in (DefinedNo), false is in evaluation position (it has no subterm) and in (DefinedYes), the resulting term M is a term that is not evaluation position in the initial configuration, so it is a subterm of Q_0 . The case of `get` is similar to the one of `find`. In (CtxT), the initial configuration of the recursive call is ok because terms that are not in evaluation position in N are also not in evaluation position in $C[N]$, so they are subterms of Q_0 . We conclude for the configurations inside the recursive call by induction hypothesis. The target configuration is ok because terms that are not in evaluation position in $C[N']$ are either not in evaluation position inside C , in which case they are subterms of Q_0 because the initial configuration is ok, or they are not in evaluation position in N' , in which case they are also subterms of Q_0 because the target configuration of the recursive call is ok. In (CtxEventT), the target configuration is ok because the term is in evaluation position (it has no subterm).

Property 2: The desired property is preserved for unmodified oracle definitions in \mathcal{Q} . This is enough to conclude for (Nil). For (Par) and (Repl), the resulting oracle definitions Q_1, Q_2, Q are not in evaluation position in the initial configuration, so they are subprocesses of Q_0 up to renaming of oracles. For (NewOracle), Q is not in evaluation position in the initial configuration, so it is a subprocess of Q_0 up to renaming of oracles, and so is $Q\{O'/O\}$. For (DefOracle), the target configuration is ok because the subterms or subprocesses not in evaluation position are subterms or subprocesses of P , which are not in evaluation position in the initial configuration, so that are subterms or subprocesses of Q_0 up to renaming of oracles.

Property 3 follows immediately from Property 2 by induction.

Property 4: In (Return), (AttYield) or (AttIO), Q is not in evaluation position in the initial configuration, so it is a subprocess of Q_0 up to renaming of oracles. Therefore, the configuration $\{(\sigma, Q)\}, \mathcal{O}r$ is ok. By Property 3, all configurations in the computation of $Q', \mathcal{O}r'$ and $Q', \mathcal{O}r'$ itself are ok. Moreover, in (Return), P_1 is not in evaluation position in the initial configuration, so it is a subprocess of Q_0 up to renaming of oracles. We can then conclude that the target configuration is ok. In (AttIO) or (AttYield), P' is not in evaluation position in the initial configuration, so it is a subprocess of Q_0 up to renaming of oracles. We can then conclude that the target configuration is ok. Rule (OracleCall) is similar with P', P_1, P_2 also not in evaluation position in the initial configuration. In (Yield), P_2 is not in evaluation position in the initial configuration. All other cases can be treated similarly to terms in Property 1. Moreover, in the computation of $\text{initConfig}(Q_0, \mathcal{A})$, the configuration $\{(\sigma_0, Q_0)\}, (\text{fo}(Q), \emptyset)$ is ok, so by Property 3, $\text{initConfig}(Q_0, \mathcal{A})$ is ok except for the process 0 that follows `return` (*start*) in $\text{initConfig}(Q_0, \mathcal{A})$. That process disappears in the first reduction, which is by (AttIO). \square

Corollary 1 *Consider a trace Tr of Q_0 for any $\mathcal{A} \in \mathcal{BB}$.*

In Tr , the target oracle body of rules (New), (Let), (If1), (If2), (Find1), (Find2), (Insert), (Get1), (Get2), (OracleCall), (Return), (Yield), (AttIO), (AttYield), (Event) is a subprocess of Q_0 up to renaming of oracles.

In Tr , the target term of rules (NewT), (LetT), (IfT1), (IfT2), (FindT1), (FindT2), (InsertT), (GetT1), (GetT2), (EventT), (DefinedYes) is a subterm Q_0 .

Proof The target term or oracle body of these rules appears in non-evaluation position in the initial configuration of these rules, so by Lemma 4, it is a subterm of Q_0 or a subprocess of Q_0 up to renaming of oracles. \square

We say that a configuration $Conf$ is at program point μ in a trace Tr of Q_0 when $Conf$ occurs in the derivation of Tr , and either $Conf = E, \sigma, {}^\mu N, \mathcal{T}, \mu \mathcal{E}v$ for some subterm ${}^\mu N$ of Q_0 or $Conf = E, (\sigma, {}^\mu P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v$ for some subprocess ${}^\mu P$ of Q_0 up to renaming of oracles.

Lemma 5 *Let Tr be a trace of Q_0 for any $\mathcal{A} \in \mathcal{BB}$. Let $Conf$ be a configuration at program point μ in Tr . If μ is not inside a condition of find or get, then $Conf$ is not in the derivation of an hypothesis of a rule for find or get inside the derivation of Tr .*

Proof The only rules that can conclude with a term find or get are (NewT), (LetT), (IfT1), (IfT2), (FindT1), (FindT2), (InsertT), (GetT1), (GetT2), (EventT), (DefinedYes) and, by Corollary 1, their target term is a subterm of Q_0 . The situation is similar find and get processes. So the executed term or oracle in the initial configuration of rules for find and get is always a subterm or subprocess of Q_0 up to renaming of oracles. When a configuration $Conf$ is in the derivation of an hypothesis of a rule for find or get, it therefore always deals with program points syntactically in conditions of find or get in Q_0 . (The semantic rules do not create program points.) Since μ is not inside a condition of find or get, we conclude that $Conf$ is not in the derivation of an hypothesis of a rule for find or get. \square

Given a trace Tr , we define a partial ordering relation \preceq_{Tr} (reflexive, transitive, antisymmetric) on the occurrences of configurations in the derivation of Tr : if $Conf_1 \xrightarrow{P}_t Conf_2$ occurs in the derivation of Tr , then $Conf_1 \preceq_{Tr} Conf_2$ and for all $Conf$ that occur in the derivation of the assumptions of $Conf_1 \xrightarrow{P}_t Conf_2$, $Conf_1 \preceq_{Tr} Conf \preceq_{Tr} Conf_2$, and similarly for \rightsquigarrow instead of \xrightarrow{P}_t . If Tr is a trace of Q_0 for \mathcal{A} , then for all $Conf$ that occur in the derivation of $\{(\sigma_0, Q_0)\}, (fo(Q_0), \emptyset) \rightsquigarrow^* \mathcal{Q}, \mathcal{O}r$, we have $Conf \preceq_{Tr} \text{initConfig}(Q_0, \mathcal{A})$. When $Conf_1 \preceq_{Tr} Conf_2$, we say that $Conf_1$ occurs before $Conf_2$ in Tr , or equivalently, that $Conf_2$ occurs after $Conf_1$ in Tr .

We say that a configuration $Conf = E, \sigma, {}^\mu N, \mathcal{T}, \mu \mathcal{E}v$, $Conf = E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v$ with $(\sigma', {}^\mu Q) \in \mathcal{Q}$, or $(\sigma', {}^\mu P) \in St$, or $Conf = \mathcal{Q}, \mathcal{O}r$ with $(\sigma', {}^\mu Q) \in \mathcal{Q}$ is inside program point μ . A configuration may be inside several program points: $Conf = \mathcal{Q}, \mathcal{O}r$ is inside the program points of the oracle definitions in \mathcal{Q} , $Conf = E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v$ is inside the program points of the oracle bodies in St as well as the oracle definitions in \mathcal{Q} .

We say that the program point μ is immediately above the program points μ_j in a process Q_0 when Q_0 contains one of the following constructs:

$$\begin{aligned}
 & {}^\mu x[{}^{\mu_1} M_1, \dots, {}^{\mu_m} M_m] \\
 & {}^\mu f({}^{\mu_1} M_1, \dots, {}^{\mu_m} M_m) \\
 & {}^\mu x[\tilde{i}] \stackrel{R}{\leftarrow} T; {}^{\mu_1} N \\
 & {}^\mu \text{let } x[\tilde{i}] : T = {}^{\mu_1} M \text{ in } {}^{\mu_2} N \\
 & {}^\mu \text{if } {}^{\mu_1} M \text{ then } {}^{\mu_2} N \text{ else } {}^{\mu_3} N' \\
 & {}^\mu \text{find}[\text{unique?}] \left(\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat} \right. \\
 & \quad \left. \text{defined}({}^{\mu_{j,1}} M_{j1}, \dots, {}^{\mu_{j,l_j}} M_{jl_j}) \wedge {}^{\mu_{j,l_j+1}} M'_j \text{ then } {}^{\mu_{j,l_j+2}} N_j \right) \text{ else } {}^{\mu_0} N' \\
 & {}^\mu \text{insert } Tbl({}^{\mu_1} M_1, \dots, {}^{\mu_l} M_l); {}^{\mu_0} N \\
 & {}^\mu \text{get}[\text{unique?}] Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } {}^{\mu_1} M \text{ in } {}^{\mu_2} N \text{ else } {}^{\mu_3} N' \\
 & {}^\mu \text{event } e({}^{\mu_1} M_1, \dots, {}^{\mu_l} M_l); {}^{\mu_0} N \\
 & {}^\mu ({}^{\mu_1} Q \mid {}^{\mu_2} Q')
 \end{aligned}$$

$$\begin{aligned}
&{}^\mu \text{foreach } i \leq n \text{ do } {}^{\mu_1} Q \\
&{}^\mu \text{newOracle } O; {}^{\mu_1} Q \\
&{}^\mu O[\tilde{i}](x[\tilde{i}] : T) := {}^{\mu_0} P \\
&{}^\mu \text{return } ({}^{\mu_1} M); {}^{\mu_2} Q \\
&{}^\mu \text{let } x[\tilde{i}] : T = O[{}^{\mu_1} M_1, \dots, {}^{\mu_l} M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m]({}^{\mu_0} M) \text{ in } {}^{\mu'_1} P_1 \text{ else } {}^{\mu'_2} P_2 \\
&{}^\mu x[\tilde{i}] \stackrel{R}{\leftarrow} T; {}^{\mu_1} P \\
&{}^\mu \text{let } x[\tilde{i}] = {}^{\mu_1} M \text{ in } {}^{\mu_2} P \\
&{}^\mu \text{if } {}^{\mu_1} M \text{ then } {}^{\mu_2} P \text{ else } {}^{\mu_3} P' \\
&{}^\mu \text{find}[unique?] \left(\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat} \right. \\
&\quad \left. \text{defined}({}^{\mu_{j,1}} M_{j1}, \dots, {}^{\mu_{j,l_j}} M_{jl_j}) \wedge {}^{\mu_{j,l_j+1}} M_j \text{ then } {}^{\mu_{j,l_j+2}} P_j \right) \text{ else } {}^{\mu_0} P \\
&{}^\mu \text{insert } Tbl({}^{\mu_1} M_1, \dots, {}^{\mu_l} M_l); {}^{\mu_0} P \\
&{}^\mu \text{get}[unique?] Tbl(x_1[\tilde{i}] : T_1, \dots, x_l[\tilde{i}] : T_l) \text{ suchthat } {}^{\mu_1} M \text{ in } {}^{\mu_2} P \text{ else } {}^{\mu_3} P' \\
&{}^\mu \text{event } e({}^{\mu_1} M_1, \dots, {}^{\mu_l} M_l); {}^{\mu_0} P
\end{aligned}$$

The relation “ μ is above μ' ” is the reflexive and transitive closure of “ μ is immediately above μ' ”.

Lemma 6 *Let Tr be a trace of Q_0 for any $\mathcal{A} \in \mathcal{BB}$. If $Conf$ is a configuration in Tr inside program point μ and μ is a program point in Q_0 , then either $Conf$ is at the program point μ at the top of Q_0 , or there exists a configuration $Conf' \neq Conf$ such that $Conf' \preceq_{Tr} Conf$ and $Conf'$ is inside program point μ or inside the program point μ' immediately above μ in Q_0 .*

As a consequence, if a configuration $Conf$ inside program point μ in Q_0 is in Tr , then there are configurations inside all program points above μ in Q_0 before $Conf$ in Tr .

Proof First property. Since $Conf$ is a configuration in Tr , we are in one of the following cases:

- $Conf = \{(\sigma_0, Q_0)\}, (\text{fo}(Q_0), \emptyset)$, the very first configuration of Tr . The configuration $Conf$ is at the program point μ at the top of Q_0 .
- $Conf = \text{initConfig}(Q_0, \mathcal{A}) = \emptyset, (\sigma_0, {}^{\mu''} \text{return}(start)) :: \mathcal{A}, \mathcal{Q}, \mathcal{Or}, \emptyset, \emptyset$. Since μ'' is not in Q_0 , μ is the program point of an oracle declaration in \mathcal{Q} . Then $Conf' = \mathcal{Q}, \mathcal{Or}$ is also inside μ and $Conf' \preceq_{Tr} Conf$, by definition of \preceq_{Tr} .
- $Conf$ is the initial configuration of an assumption of a semantic rule. In rules for **find**, the initial configuration of the assumption is not inside a program point (because it evaluates the **defined** condition, not a term). In rules for **get**, (**CtxT**), and (**Ctx**), the initial configuration of the conclusion is inside the program point μ' immediately above μ . In rules (**DefinedNo**) and (**DefinedYes**), these rules are used to conclude assumptions of **find**, and the initial configuration of the **find** rule is inside the program point μ' immediately above μ . In rules (**Return**) and (**AttIO**), $Conf = \{(\sigma, Q)\}, \mathcal{Or}$, and the initial configuration of the conclusion of the rule is inside the program point μ' immediately above μ .
- $Conf$ is the target configuration of a semantic rule. In rules (**NewT**), (**LetT**), (**IfT1**), (**IfT2**), (**FindT1**), (**FindT2**), (**InsertT**), (**GetT1**), (**GetT2**), and (**EventT**), the initial configuration of the rule is inside the program point μ' immediately above μ . In rule (**CtxT**), the initial

configuration of the rule is inside the same program point μ . In rule (DefinedYes), this rule is used to conclude assumptions of `find`, and the initial configuration of the `find` rule is inside the program point μ' immediately above μ . In the rules for oracle definitions, if μ is the program point of an unchanged element of \mathcal{Q} , the initial configuration of the rule is also inside μ . This is sufficient for (Nil). If μ is the program point of a modified oracle, then for rules (Par), (Repl), and (NewOracle), the initial configuration of the rule is inside the program point μ' immediately above μ , and for rule (DefOracle), the initial configuration of the rule is inside the same program point μ . In the rules for oracle bodies other than (AttIO), (AttYield), (OracleCall), and (Return), \mathcal{Q} is unchanged, so if μ is the program point of an oracle declaration in \mathcal{Q} , then the initial configuration of the rule is inside the same program point μ . If μ is the program point of an oracle body not at the root of the stack, then the initial configuration of the rule is also inside the same program point μ , because that part of the stack is unchanged. In rules (New), (Let), (If1), (If2), (Find1), (Find2), (Insert), (Get1), (Get2), (Event), and (Yield), if μ is the program point of an oracle body at the root of the stack, then the initial configuration of the rule is inside the program point μ' immediately above μ . In rule (Ctx), the initial configuration of the rule is inside the same program point μ .

In rules (OracleCall), (AttYield), and (AttIO), if μ is the program point of an oracle definition in \mathcal{Q} , then the initial configuration of the rule is inside the same program point μ . In rules (OracleCall) and (AttYield), if μ is the program point of P' , then the initial configuration of the rule is inside the program point μ' immediately above μ , because $(\sigma', Q_0) \in S \subseteq \mathcal{Q}$ and μ' is the program point of Q_0 . Additionally, in rule (OracleCall), if μ is the program point of an oracle body in $(\sigma, P'_0) :: St$, then the initial configuration of the rule is inside the same program point μ . In rule (AttIO), we have $(\sigma', Q_0) \in S \subseteq \mathcal{Q} \uplus \mathcal{Q}'$. If $(\sigma', Q_0) \in \mathcal{Q}$ and μ is the program point of P' , then the initial configuration of the rule is inside the program point μ' immediately above μ , as above because μ' is the program point of Q_0 . If $(\sigma', Q_0) \in \mathcal{Q}'$ and μ is the program point of P' , then $\mathcal{Q}', \mathcal{O}r'$ is inside the program point μ' immediately above μ .

In rule (Return), if μ is the program point of an oracle definition in \mathcal{Q} , then the initial configuration of the rule is inside the same program point μ . If μ is the program point of an oracle definition in \mathcal{Q}' , then the configuration $\mathcal{Q}', \mathcal{O}r'$ in the assumption of the rule is inside the same program point μ . If μ is the program point of P_1 , then the initial configuration of the rule is inside the program point μ' immediately above μ as (σ', P_0) is in the stack and μ' is the program point of P_0 . If μ is the program point of an oracle body in St , then the initial configuration of the rule is inside the same program point μ .

Second property. Suppose that $Conf$ is inside μ and there is a program point μ' immediately above μ in Q_0 . Let us show that there exists $Conf' \preceq_{Tr} Conf$ such that $Conf'$ is inside μ' . The proof proceeds by well-founded induction on \preceq_{Tr} . The program point μ is not at the top of Q_0 , so by the first property, there exists $Conf' \neq Conf$ such that $Conf' \preceq_{Tr} Conf$ and $Conf'$ is inside μ or inside μ' . If $Conf'$ is inside μ , we conclude by applying the induction hypothesis on $Conf'$. If $Conf'$ is inside μ' , we have the result. By applying this property repeatedly, we obtain the second property. \square

Lemma 7 *Let Tr be a trace of Q_0 for any $\mathcal{A} \in \mathcal{BB}$.*

1. *If $Conf = E, \sigma, C_m[\dots C_1[{}^\mu M] \dots], \mathcal{T}, \mu \mathcal{E}v$ is the target configuration of a semantic rule in Tr , where μ is a program point in Q_0 , C_1, \dots, C_m are term contexts defined in Figure 7, and ${}^\mu M$ is not a subterm of Q_0 , then the reduction that yields $Conf$ is obtained by m*

applications of (CtxT) with contexts C_m, \dots, C_1 from a reduction with target configuration $E, \sigma, {}^\mu M, \mathcal{T}, \mu\mathcal{E}v$, itself proved by (CtxT).

2. If $Conf = E, (\sigma, C_0[C_m[\dots C_1[{}^\mu M]\dots]]) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ is the target configuration of a semantic rule in Tr , where μ is a program point in Q_0 , C_0 is an oracle context defined in Figure 12, C_1, \dots, C_m are term contexts defined in Figure 7, and ${}^\mu M$ is not a subterm of Q_0 , then the reduction that yields $Conf$ is obtained by one application of (Ctx) with context C_0 and m applications of (CtxT) with contexts C_m, \dots, C_1 from a reduction with target configuration $E, \sigma, {}^\mu M, \mathcal{T}, \mu\mathcal{E}v$, itself proved by (CtxT).

Proof Property 1. This property is proved by induction of m . The reduction that yields $Conf$ cannot be obtained by (NewT), (LetT), (IfT1), (IfT2), (FindT1), (FindT2), (InsertT), (GetT1), (GetT2), (EventT), (DefinedYes), because in this case, by Corollary 1, $C_m[\dots C_1[{}^\mu M]\dots]$ would be a subterm of Q_0 , so ${}^\mu M$ would be a subterm of Q_0 . So it is obtained by (CtxT). For $m = 0$, this is enough to conclude. For $m > 0$, the rule (CtxT) is applied with context C_m . Indeed, since ${}^\mu M$ is not a value, the hole of the context cannot be after $C_{m-1}[\dots C_1[{}^\mu M]\dots]$ and, since ${}^\mu M$ is not a subterm of Q_0 , the hole of the context cannot be before $C_{m-1}[\dots C_1[{}^\mu M]\dots]$. (If it were before, $C_{m-1}[\dots C_1[{}^\mu M]\dots]$ would not be in evaluation position in the initial configuration of rule (CtxT), so by Lemma 4, $C_{m-1}[\dots C_1[{}^\mu M]\dots]$ would be a subterm of Q_0 .) Hence the reduction that yields $Conf$ is obtained from a reduction that yields $E, \sigma, C_{m-1}[\dots C_1[{}^\mu M]\dots], \mathcal{T}, \mu\mathcal{E}v$ by applying (CtxT) with context C_m . We conclude by applying the induction hypothesis.

Property 2. The reduction that yields $Conf$ cannot be obtained by (New), (Let), (If1), (If2), (Find1), (Find2), (Insert), (Get1), (Get2), (Return), (Yield), (OracleCall), (AttYield), (AttIO), or (Event), because in this case, by Corollary 1, the oracle body of $Conf$ would be a subprocess of Q_0 up to renaming of oracles, so ${}^\mu M$ would be a subterm of Q_0 . Therefore, $Conf$ is the target configuration of (Ctx). Furthermore, by the same reasoning as in Property 1, this rule is applied with context C_0 . Hence the reduction that yields $Conf$ is obtained from a reduction that yields $E, \sigma, C_m[\dots C_1[{}^\mu M]\dots], \mathcal{T}, \mu\mathcal{E}v$ by applying (Ctx) with context C_0 . We conclude by Property 1. \square

Lemma 8 *Let Tr be a trace of Q_0 for any $\mathcal{A} \in \mathcal{BB}$.*

1. If $Conf = E, (\sigma, {}^\mu P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ is a configuration in Tr where μ is a program point in Q_0 , then this configuration is derived in Tr from a configuration $Conf' = E', (\sigma, {}^\mu P') :: St', \mathcal{Q}, \mathcal{O}r, \mathcal{T}', \mu\mathcal{E}v'$ where ${}^\mu P'$ is a subprocess of Q_0 up to renaming of oracles, by (Ctx) any number of times.
2. If $Conf = \{(\sigma, {}^\mu Q)\} \uplus \mathcal{Q}, \mathcal{O}r$ is a configuration in Tr where μ is a program point in Q_0 , then, possibly after swapping reductions in Tr , this configuration is derived from a configuration $Conf' = \{(\sigma, {}^\mu Q')\} \uplus \mathcal{Q}, \mathcal{O}r$ where ${}^\mu Q'$ is a subprocess of Q_0 up to renaming of oracles, by (DefOracle) any number of times.
3. If $Conf = E, (\sigma, C_0[C_m[\dots C_1[{}^\mu M]\dots]]) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ is a configuration in Tr where μ is a program point in Q_0 , C_0 is an oracle context defined in Figure 12, and C_1, \dots, C_m are term contexts defined in Figure 7, then we have

$$E', \sigma, {}^\mu M', \mathcal{T}', \mu\mathcal{E}v' \xrightarrow{p_t} \dots \xrightarrow{p'_t} E, \sigma, {}^\mu M, \mathcal{T}, \mu\mathcal{E}v$$

by (CtxT) any number of times, where ${}^\mu M'$ is a subterm of Q_0 and in Tr , these reductions are in fact performed starting from a configuration $Conf' = E', (\sigma, C_0[C_m[\dots C_1[{}^\mu M']\dots]])$, $\mathcal{Q}, \mathcal{O}r, \mathcal{T}', \mu\mathcal{E}v'$ under one application of rule (Ctx) with context C_0 and m applications of rule (CtxT) with contexts C_m, \dots, C_1 .

4. If $Conf = E, \sigma, C_l[\dots C_1[\mu M]\dots], \mathcal{T}, \mu\mathcal{E}v$ is a configuration in Tr and μ is a program point in Q_0 where C_1, \dots, C_l are term contexts defined in Figure 7, then we have

$$E', \sigma, \mu M', \mathcal{T}', \mu\mathcal{E}v' \xrightarrow{p'} \dots \xrightarrow{p'} E, \sigma, \mu M, \mathcal{T}, \mu\mathcal{E}v$$

by (CtxT) any number of times, where $\mu M'$ is a subterm of Q_0 and in Tr , these reductions are in fact performed starting

- from a configuration $Conf' = E', \sigma, C_m[\dots C_1[\mu M']\dots], \mathcal{T}', \mu\mathcal{E}v'$ where $m \geq l$, C_1, \dots, C_m are term contexts defined in Figure 7, under m applications of (CtxT) with contexts C_m, \dots, C_1 .
- or from a configuration $Conf' = E', (\sigma, C_0[C_m[\dots C_1[\mu M']\dots]]) :: St, \mathcal{Q}, Or, \mathcal{T}', \mu\mathcal{E}v'$ where $m \geq l$, C_0 is an oracle context defined in Figure 12, and C_1, \dots, C_l are term contexts defined in Figure 7, under one application of (Ctx) with context C_0 and m applications of (CtxT) with contexts C_m, \dots, C_1 .

Proof Property 1. The proof proceeds by well-founded induction on \preceq_{Tr} . The configuration $Conf$ cannot be $\text{initConfig}(Q_0, \mathcal{A})$ because μ is a program point in Q_0 . Then, $Conf$ is the target configuration of some semantic rule. If $Conf$ is the target configuration of (New), (Let), (If1), (If2), (Find1), (Find2), (Insert), (Get1), (Get2), (OracleCall), (Return), (AttYield), (AttIO) or (Event), then by Corollary 1, μP is a subprocess of Q_0 up to renaming of oracles, so the property holds with $Conf' = Conf$. If it is the target configuration of (Ctx), we conclude by applying the induction hypothesis to the initial configuration of this rule.

Property 2. The proof proceeds by well-founded induction on \preceq_{Tr} . If $Conf = \{(\sigma_0, Q_0)\}$, $(\text{fo}(Q_0), \emptyset)$, then the property holds with $Conf' = Conf$, since Q_0 is a subprocess of Q_0 . If $Conf$ is the initial configuration of the assumption of (Return) or (AttIO), then by Lemma 4, Q is a subprocess of Q_0 up to renaming of oracles, since Q occurs in non-evaluation position in the initial configuration of (Return) or (AttIO). So the property holds with $Conf' = Conf$. Otherwise, $Conf$ is the target configuration of a semantic rule.

- If that rule does not affect Q , then it is of the form $\{(\sigma, \mu Q)\} \uplus \mathcal{Q}', Or' \rightsquigarrow \{(\sigma, \mu Q)\} \uplus \mathcal{Q}, Or$. We apply the induction hypothesis to $\{(\sigma, \mu Q)\} \uplus \mathcal{Q}', Or'$, so possibly after swapping reductions in Tr , we have $\{(\sigma, \mu Q')\} \uplus \mathcal{Q}', Or' \rightsquigarrow^* \{(\sigma, \mu Q)\} \uplus \mathcal{Q}', Or'$ by (DefOracle) any number of times, where $\mu Q'$ is a subprocess of Q_0 up to renaming of oracles. By swapping reductions, we have $\{(\sigma, \mu Q')\} \uplus \mathcal{Q}', Or' \rightsquigarrow \{(\sigma, \mu Q')\} \uplus \mathcal{Q}, Or \rightsquigarrow^* \{(\sigma, \mu Q)\} \uplus \mathcal{Q}, Or$, so we obtain the desired property with $Conf' = \{(\sigma, \mu Q')\} \uplus \mathcal{Q}, Or$.
- Otherwise, that rule affects Q . If that rule is (Par), (Repl), or (NewOracle), then by Lemma 4, Q is a subprocess of Q_0 up to renaming of oracles, since Q occurs in non-evaluation position in the initial configuration of the rule. So the property holds with $Conf' = Conf$. If that rule is (DefOracle), then we obtain the result by induction hypothesis applied to the initial configuration of the rule.

Property 3. The proof proceeds by well-founded induction on \preceq_{Tr} . The configuration $Conf$ cannot be $\text{initConfig}(Q_0, \mathcal{A})$ because μ is a program point in Q_0 . Then, $Conf$ is the target configuration of some semantic rule. If μM is a subterm of Q_0 , then the property holds with $Conf' = Conf$. Otherwise, by Lemma 7, Property 2, the reduction that yields $Conf$ is obtained by one application of (Ctx) with context C_0 and m applications of (CtxT) with contexts C_m, \dots, C_1 from a reduction with target configuration $E, \sigma, \mu M, \mathcal{T}, \mu\mathcal{E}v$, itself proved by (CtxT). We obtain the desired property by applying the induction hypothesis to the configuration before this reduction step by (Ctx).

Property 4. The proof proceeds by well-founded induction on \preceq_{Tr} . First case: $Conf$ is the initial configuration of an assumption of a semantic rule. This rule cannot be a rule for `find` (in rules for `find`, the initial configuration of assumption is not inside a program point, because it evaluates the defined condition, not a term). In rules for `get`, (DefinedNo), and (DefinedYes), by Lemma 4, since the term in $Conf$ occurs in non-evaluation position in the initial configuration of the rule, it is a subterm of Q_0 , so ${}^\mu M$ is a subterm of Q_0 . The property holds with $Conf' = Conf$. In (CtxT), we let C_{l+1} be the context used in this rule, and we conclude by induction hypothesis applied to the initial configuration of this rule: $E, \sigma, C_{l+1}[C_l[\dots C_1[{}^\mu M]\dots]], \mathcal{T}, \mu\mathcal{E}v$. In (Ctx), we let $m = l$. The initial configuration of the rule is of the form $E, (\sigma, C_0[C_l[\dots C_1[{}^\mu M]\dots]])$, $\mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ for some process context C_0 defined in Figure 12. We conclude by applying Property 3 to that configuration.

Second case: $Conf$ is the target configuration of a semantic rule. If ${}^\mu M$ is a subterm of Q_0 , then the property holds with $Conf' = Conf$. Otherwise, by Lemma 7, Property 1, the reduction that yields $Conf$ is obtained by l applications of (CtxT) with contexts C_l, \dots, C_1 from a reduction with target configuration $E, \sigma, {}^\mu M, \mathcal{T}, \mu\mathcal{E}v$, itself proved by (CtxT). We apply the induction hypothesis to the initial configuration of the reduction that yields $Conf$, and we continue the reduction one more step by applying possibly (Ctx) with context C_0 , (CtxT) m times with contexts C_m, \dots, C_1 under the reduction with target configuration $E, \sigma, {}^\mu M, \mathcal{T}, \mu\mathcal{E}v$. This is also what happens in Tr . (By inspection of the rules, only the rule (Ctx) with context C_0 can reduce an oracle body $C_0[N]$, where C_0 is a process context defined in Figure 12; only the rule (CtxT) with context C_j can reduce a term $C_j[N]$, where C_j is a term context defined in Figure 7.) \square

2.4.3 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.

We define the multiset of variable accesses that may be defined by a term or a process (given the replication indices fixed by a mapping sequence σ) as follows:

$$Defined(\sigma, {}^\mu i) = \emptyset$$

$$Defined(\sigma, {}^\mu x[M_1, \dots, M_m]) = \bigsqcup_{j=1}^m Defined(\sigma, M_j)$$

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

$$Defined(\sigma, {}^\mu x[\tilde{i}] \stackrel{R}{\leftarrow} T; N) = \{x[\sigma(\tilde{i})]\} \uplus Defined(\sigma, N)$$

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

$$Defined(\sigma, {}^\mu \text{if } M \text{ then } N \text{ else } N') = Defined(\sigma, M) \uplus \max(Defined(\sigma, N), Defined(\sigma, N'))$$

$$Defined(\sigma, {}^\mu \text{find}[\text{unique?}] (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{v}_j \leq \tilde{n}_j \text{ suchthat defined}(\tilde{M}_j) \wedge M_j \text{ then } N_j) \text{ else } N) = \\ \max(\max_{j=1}^m \max_{\tilde{a} \leq \tilde{n}_j} Defined(\sigma[\tilde{i}_j \mapsto \tilde{a}], M_j), \max_{j=1}^m \{\tilde{u}_j[\sigma(\tilde{i})]\}) \uplus Defined(\sigma, N_j), Defined(\sigma, N))$$

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

$$\begin{aligned}
 \text{Defined}(\sigma, \text{get}[\text{unique?}] \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 \max(\text{Defined}(\sigma, M), \text{Defined}(\sigma, N)), \text{Defined}(\sigma, N')) & \\
 \text{Defined}(\sigma, \text{event } e(M_1, \dots, M_l); N) &= \biguplus_{j=1}^l \text{Defined}(\sigma, M_j) \uplus \text{Defined}(\sigma, N) \\
 \text{Defined}(\sigma, \text{event_abort } e) &= \emptyset \\
 \text{Defined}(\sigma, a) &= \text{Defined}(\sigma, \text{event_abort } (\mu, \tilde{a}) : e) = \emptyset \\
 \text{Defined}(\sigma, \text{0}) &= \emptyset \\
 \text{Defined}(\sigma, \text{ } (Q_1 \mid Q_2)) &= \text{Defined}(\sigma, Q_1) \uplus \text{Defined}(\sigma, Q_2) \\
 \text{Defined}(\sigma, \text{foreach } i \leq n \text{ do } Q) &= \biguplus_{a \in [1, n]} \text{Defined}(\sigma[i \mapsto a], Q) \\
 \text{Defined}(\sigma, \text{newOracle } O; Q) &= \text{Defined}(\sigma, Q) \\
 \text{Defined}(\sigma, \text{ } O[\tilde{i}](x[\tilde{i}] : T) := P) &= \{x[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, P) \\
 \text{Defined}(\sigma, \text{return } (M); Q) &= \text{Defined}(\sigma, M) \uplus \text{Defined}(\sigma, Q) \\
 \text{Defined}(\sigma, \text{let } x[\tilde{i}] : T = O[M_1, \dots, M_l, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](N) \text{ in } P \text{ else } P') &= \\
 \biguplus_{j=1}^l \text{Defined}(M_j) \uplus \{u_j[\sigma(\tilde{i})] \mid j \leq m\} \uplus \max(\{x[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, P), \text{Defined}(\sigma, P')) & \\
 \text{Defined}(\sigma, \text{ } x[\tilde{i}] \stackrel{R}{\leftarrow} 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{if } M \text{ then } P \text{ else } P') &= \text{Defined}(M) \uplus \max(\text{Defined}(P), \text{Defined}(P')) \\
 \text{Defined}(\sigma, \text{find}[\text{unique?}] (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(\tilde{M}_j) \wedge M_j \text{ then } P_j) \text{ else } P) &= \\
 \max(\max_{j=1}^m \max_{\tilde{a} \leq \tilde{n}_j} \text{Defined}(\sigma[\tilde{i}_j \mapsto \tilde{a}], M_j), \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) &= \biguplus_{j=1}^l \text{Defined}(\sigma, M_j) \uplus \text{Defined}(\sigma, P) \\
 \text{Defined}(\sigma, \text{get}[\text{unique?}] \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 \max(\text{Defined}(\sigma, M), \text{Defined}(\sigma, P)), \text{Defined}(\sigma, P')) & \\
 \text{Defined}(\sigma, \text{event } e(M_1, \dots, M_l); P) &= \biguplus_{j=1}^l \text{Defined}(\sigma, M_j) \uplus \text{Defined}(\sigma, P) \\
 \text{Defined}(\sigma, \text{event_abort } e) &= \emptyset \\
 \text{Defined}(\sigma, \text{abort}) &= \emptyset \\
 \text{Defined}(\sigma, \text{yield}) &= \emptyset
 \end{aligned}$$

Notice that, by Invariant 5, the terms \tilde{M}_j in defined conditions of find do not define any variable. By Invariant 3, the variables defined in conditions of find and get can be considered as defined temporarily only during the evaluation of the considered condition. Given a configuration $\text{Conf} = E, \sigma, N, \mathcal{T}, \mu \mathcal{E}v$ or $\text{Conf} = E, \mathcal{S}t, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v$, we denote by E_{Conf} the environment E in configuration Conf . For $\text{Conf} = \mathcal{Q}, \mathcal{O}r$, E_{Conf} is empty. We define

$$\text{Defined}^{\text{Fut}}(E, \sigma, M, \mathcal{T}, \mu \mathcal{E}v) = \text{Defined}(\sigma, M)$$

$$\begin{aligned}
\text{Defined}^{\text{Fut}}(\mathcal{Q}, \mathcal{O}r) &= \bigsqcup_{(\sigma, Q) \in \mathcal{Q}} \text{Defined}(\sigma, Q) \\
\text{Defined}^{\text{Fut}}(E, \text{St}, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v) &= \bigsqcup_{(\sigma, P) \in \text{St}} \text{Defined}(\sigma, P) \uplus \bigsqcup_{(\sigma, Q) \in \mathcal{Q}} \text{Defined}(\sigma, Q) \\
\text{Defined}(\text{Conf}) &= \text{Dom}(E_{\text{Conf}}) \uplus \text{Defined}^{\text{Fut}}(\text{Conf}).
\end{aligned}$$

Invariant 9 (Single definition, for executing games) The semantic configuration Conf (which can be $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ or $\mathcal{Q}, \mathcal{O}r$ or $E, \text{St}, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$) satisfies Invariant 9 if and only if $\text{Defined}(\text{Conf})$ does not contain duplicate elements.

Lemma 9 *Let Tr be trace of Q_0 for any $\mathcal{A} \in \mathcal{BB}$. If Q_0 satisfies Invariant 1, then all semantic configurations in the derivation of Tr satisfy Invariant 9.*

Proof sketch We first show that, for all program points μ in Q_0 , if $\text{Dom}(\sigma) = I_\mu$ are the current replication indices at μ and the process or term Q at μ satisfies Invariant 1, then all elements of $\text{Defined}(\sigma, Q)$ are of the form $x[\tilde{a}]$ where $x \in \text{vardef}(Q)$ and $\text{Im}(\sigma)$ is a prefix of \tilde{a} . The proof proceeds by induction on Q . At the definition of a variable $x[\tilde{i}]$, $x[\sigma(\tilde{i})]$ is added to $\text{Defined}(\sigma, Q)$ and we have $x \in \text{vardef}(Q)$; by Invariant 1, \tilde{i} are the current replication indices at that definition, so $\sigma(\tilde{i}) = \text{Im}(\sigma)$. All recursive calls $\text{Defined}(\sigma', \mu' Q')$ consider an extension σ' of σ and a subprocess or subterm Q' of Q (so $\text{vardef}(Q') \subseteq \text{vardef}(Q)$) such that $\text{Dom}(\sigma') = I_{\mu'}$ are the current replication indices at μ' .

Next, we show that, for all program points μ , if $\text{Dom}(\sigma) = I_\mu$ are the current replication indices at μ and the process or term Q at μ satisfies Invariant 1, then $\text{Defined}(\sigma, Q)$ does not contain duplicate elements. The proof proceeds by induction on Q . All multiset unions in the computation of $\text{Defined}(\sigma, Q)$ are disjoint unions by the property above, because either they use different extensions of σ (case of replication) or they use disjoint variable definitions or subprocesses or subterms in the same branch of find, if, get, or oracle call, which must define different variables by Invariant 1.

We show by induction on the derivations that, if $\text{Conf} \xrightarrow{p}_t \text{Conf}'$, then $\text{Defined}(\text{Conf}) \supseteq \text{Defined}(\text{Conf}')$ and for all semantic configurations Conf'' in the derivation of $\text{Conf} \xrightarrow{p}_t \text{Conf}'$, $\text{Defined}(\text{Conf}) \supseteq \text{Defined}(\text{Conf}'')$, and similarly with \rightsquigarrow instead of \xrightarrow{p}_t .

The result follows: since Q_0 satisfies Invariant 1, $\text{Defined}(\sigma_0, Q_0)$ does not contain duplicate elements, where σ_0 is the empty mapping sequence. Then $\{(\sigma_0, Q_0)\}$, $(\text{fo}(Q_0), \emptyset)$ satisfies Invariant 9 and so do reduce $(\{(\sigma_0, Q_0)\}, (\text{fo}(Q_0), \emptyset))$, $\text{initConfig}(Q_0, \mathcal{A})$, and the other configurations of Tr . \square

Corollary 2 *If Q_0 satisfies Invariant 1, then each variable that is not defined in a condition of find or get is defined at most once for each value of its array indices in a trace of Q_0 for any $\mathcal{A} \in \mathcal{BB}$.*

Proof Let x be the considered variable and Tr be the considered trace of Q_0 . The only semantic rules that can add $x[\tilde{a}]$ to the environment E are (NewT), (LetT), (FindT1), (GetT1), (New), (Let), (Find1), (Get1), (AttYield), (AttIO), (OracleCall) and (Return). By Corollary 1, the target term or process of these rules is a subterm or subprocess of Q_0 up to renaming of oracles. Hence, the target configuration Conf' of these rules is at some program point μ in Tr . By hypothesis, x is not defined in conditions of find or get, so μ is not inside the condition of find or get in Q_0 , so by Lemma 5, the configuration Conf' is not in the derivation of an assumption of a rule for find or get.

In order to derive a contradiction, assume that two transitions $Conf_1 \xrightarrow{p_1}_{t_1} Conf'_1$ and $Conf_2 \xrightarrow{p_2}_{t_2} Conf'_2$ inside Tr define the same variable $x[\tilde{a}]$.

- First case: one transition happens before the other, for instance $Conf'_1 \preceq_{Tr} Conf_2$. (The case $Conf'_2 \preceq_{Tr} Conf_1$ is symmetric.) Since $Conf_1 \xrightarrow{p_1}_{t_1} Conf'_1$ defines $x[\tilde{a}]$, we have $x[\tilde{a}] \in \text{Dom}(E_{Conf'_1})$. Since $Conf'_1$ is not in the derivation of an assumption of a rule for `find` or `get`, by Lemma 3, E_{Conf_2} extends $E_{Conf'_1}$, so $x[\tilde{a}] \in \text{Dom}(E_{Conf_2})$. Moreover, since $Conf_2 \xrightarrow{p_2}_{t_2} Conf'_2$ defines $x[\tilde{a}]$, we have $x[\tilde{a}] \in \text{Defined}^{\text{Fut}}(Conf_2)$, by inspecting all rules that add elements to the environment. Therefore $\text{Defined}(Conf_2) = \text{Dom}(E_{Conf_2}) \uplus \text{Defined}^{\text{Fut}}(Conf_2)$ contains twice $x[\tilde{a}]$. Contradiction with Invariant 9.
- Second case: the transitions cannot be ordered. By definition of \preceq_{Tr} , this can happen only when a semantic rule uses several derivations for its assumptions, which happens only in rules for `find` or `get`. Contradiction.

That concludes the proof. \square

Variables defined in conditions of `get` may be defined several times, once for each element of the table that is tested. Variables defined in conditions of `find` may be defined several times in case the same variable is used in several branches of the same `find`. (We use the indices of the `find` as indices of the variables defined in the condition, so when we evaluate several times the condition of a certain branch of a `find`, we use variables with different indices.) In Section 2.7, we define properties that exclude these situations (Properties 3 and 4), and we prove in Lemma 29 that every variable is defined at most once for each value of its indices when these properties are satisfied.

2.4.4 Each Oracle is defined at most once

Similarly to the variable definition case, we can show that Invariant 7 implies that each oracle is defined at most once during the execution of a process. This ensures that for every oracle call, a unique corresponding oracle is potentially available. We generalize the previous Defined^O definition of Invariant 7 to also take into account the replication indices fixed by a mapping sequence σ as follows:

$$\begin{aligned}
\text{Defined}_p^O(\sigma, \mu 0) &= \emptyset \\
\text{Defined}_p^O(\sigma, \mu(Q_1 \mid Q_2)) &= \text{Defined}_p^O(\sigma, Q_1) \uplus \text{Defined}_p^O(\sigma, Q_2) \\
\text{Defined}_p^O(\sigma, \mu \text{foreach } i \leq n \text{ do } Q) &= \bigsqcup_{a \in [1, n]} \text{Defined}_p^O(\sigma[i \mapsto a], Q) \\
\text{Defined}_p^O(\sigma, \mu \text{newOracle } O; Q) &= \text{Defined}_p^O(\sigma, Q) \setminus \{O[\tilde{i}] \mid k \in \mathbb{N}, \tilde{i} \in \mathbb{N}^k\} \\
\text{Defined}_p^O(\sigma, \mu O[\tilde{i}](x[\tilde{i}] : T) := P) &= \{O[\sigma(\tilde{i})]\} \uplus \text{Defined}_p^O(\sigma, P) \\
\text{Defined}_p^O(\sigma, \mu C[Q_1, \dots, Q_k]) &= \max_{j=1}^k \text{Defined}_p^O(\sigma, Q_j)
\end{aligned}$$

Here, similarly to the definition of Invariant 7, C is a context restricted to oracle body constructs only. We lift Defined_p^O to configurations in a similar fashion to Defined , as:

$$\begin{aligned}
\text{Defined}_p^O(\mathcal{Q}, \mathcal{O}r) &= \bigsqcup_{(\sigma, Q) \in \mathcal{Q}} \text{Defined}_p^O(\sigma, Q) \\
\text{Defined}_p^O(E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v) &= \bigsqcup_{(\sigma, P) \in St} \text{Defined}_p^O(\sigma, P) \uplus \bigsqcup_{(\sigma, Q) \in \mathcal{Q}} \text{Defined}_p^O(\sigma, Q)
\end{aligned}$$

Invariant 10 (Single definition, for executing games) The semantic configuration $Conf$ (which can be $\mathcal{Q}, \mathcal{O}r$ or $E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$) satisfies Invariant 10 if and only if $Defined_p^O(Conf)$ does not contain duplicate elements.

Lemma 10 *Let Tr be trace of Q_0 for any $\mathcal{A} \in \mathcal{BB}$. If Q_0 satisfies Invariant 7, then all semantic configurations in the derivation of Tr satisfy Invariant 10.*

We omit the proof, which is similar to the one of Lemma 9. The core difference is that $Defined_p^O$ does not directly diminish over the evolution of configurations, as $\text{newOracle } O$ introduces new definitions. We remark that however, $Defined_p^O \cap \mathcal{O}r_{\text{pub}}$ diminishes for all rules, and by construction, a $\text{newOracle } O$ command introduces new oracle definitions that are distinct from all existing ones.

2.4.5 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 11 (Defined variables, for executing games) The semantic configuration $E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ satisfies Invariant 11 if and only if every occurrence of a variable access $x[M_1, \dots, M_m]$ in St or \mathcal{Q} is either

1. present in $\text{Dom}(E)$: if $x[M_1, \dots, M_m]$ occurs in a process P' for $(\sigma', P') \in St \cup \mathcal{Q}$, then for all $j \leq m$, $E, \sigma', M_j \Downarrow a_j$ 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'_{jk} .
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'_{jk} such that $N\{\tilde{u}_j[\tilde{i}]/\tilde{i}_j\} = x[M_1, \dots, M_m]$.

Similarly, $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ satisfies Invariant 11 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 $j \leq m$, $E, \sigma, M_j \Downarrow a_j$ and $x[a_1, \dots, a_m] \in \text{Dom}(E)$) or satisfies one of the last four conditions above.

$E, \sigma, \text{defined}(M'_1, \dots, M'_l) \wedge M, \mathcal{T}, \mu\mathcal{E}v$ satisfies Invariant 11 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 $j \leq m$, $E, \sigma, M_j \Downarrow a_j$ 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 $x[M_1, \dots, M_m]$ are simple. That is why we can evaluate them by $E, \sigma', M_j \Downarrow a_j$.

Lemma 11 *If Q_0 satisfies Invariant 2, then $\text{initConfig}(Q_0, \mathcal{A})$ satisfies Invariant 11.*

Lemma 12 *Let M be a simple term. If $E, \sigma, M \Downarrow a$, then for all subterms $x[M_1, \dots, M_m]$ of M , for all $j' \leq m$, $E, \sigma, M_{j'} \Downarrow a_{j'}$ and $x[a_1, \dots, a_m]$ is in $\text{Dom}(E)$.*

Proof sketch By induction on M . □

Lemma 13 *Let N, M be simple terms. If $E, \sigma[i \mapsto a'], N \Downarrow a$ and $E, \sigma, M \Downarrow a'$, then we have $E, \sigma, N\{M/i\} \Downarrow a$.*

Proof sketch By induction on N . □

Lemma 14 *If $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mu\mathcal{E}v'$ and $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ satisfies Invariant 11, then so does $E', \sigma', M', \mathcal{T}', \mu\mathcal{E}v'$.*

If $E, \sigma, \text{defined}(M_1, \dots, M_m) \wedge M, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mu\mathcal{E}v'$ and $E, \sigma, \text{defined}(M_1, \dots, M_m) \wedge M, \mathcal{T}, \mu\mathcal{E}v$ satisfies Invariant 11, then so does $E', \sigma', M', \mathcal{T}', \mu\mathcal{E}v'$.

If $E, St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', St', \mathcal{Q}', Or', \mathcal{T}', \mu\mathcal{E}v'$ and $E, St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v$ satisfies Invariant 11, then so does $E', St', \mathcal{Q}', Or', \mathcal{T}', \mu\mathcal{E}v'$.

Moreover, if the rules that define $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mu\mathcal{E}v'$ (respectively $E, \sigma, \text{defined}(M_1, \dots, M_m) \wedge M, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mu\mathcal{E}v'$ or $E, St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', St', \mathcal{Q}', Or', \mathcal{T}', \mu\mathcal{E}v'$) require as assumption $E'', \sigma'', M'', \mathcal{T}'', \mu\mathcal{E}v'' \xrightarrow{p}_t \dots$ or $E'', \sigma'', \text{defined}(M''_1, \dots, M''_m) \wedge M'', \mathcal{T}'', \mu\mathcal{E}v'' \xrightarrow{p}_t \dots$, and the initial configuration $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v$ (resp. $E, \sigma, \text{defined}(M_1, \dots, M_m) \wedge M, \mathcal{T}, \mu\mathcal{E}v$ or $E, St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v$) satisfies Invariant 11, then so does the initial configuration of the assumption, $E'', \sigma'', M'', \mathcal{T}'', \mu\mathcal{E}v''$ or $E'', \sigma'', \text{defined}(M''_1, \dots, M''_m) \wedge M'', \mathcal{T}'', \mu\mathcal{E}v''$.

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 11, 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), (AttYield), (AttIO), (OracleCall), (Return) or (Get1) so it moves to the first case of Invariant 11.

If $x[M_1, \dots, M_m]$ is in the third case of Invariant 11, 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 11, 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 $E, \sigma[\tilde{i}_j \mapsto \tilde{a}], D_j \wedge M'_j, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{pk'}_{t_k}^* E'', \sigma', r_{k'}, \mathcal{T}, \mu\mathcal{E}v$ with $D_j \wedge M'_j = \text{defined}(M'_{j1}, \dots, M'_{jl_j}) \wedge M'_j$ and $\sigma' = \sigma[\tilde{i}_j \mapsto \tilde{a}]$ also satisfies Invariant 11. In case this assumption is reduced by (DefinedYes), we have $E, \sigma', M'_{jk}, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{1}^* E, \sigma', a_{jk}, \mathcal{T}, \mu\mathcal{E}v$, that is, $E, \sigma', M'_{jk} \Downarrow a_{jk}$. Therefore, by Lemma 12, for all $j' \leq m$, $E, \sigma', M'_{j'} \Downarrow a_{j'}$ 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 11 in $E, \sigma', M'_j, \mathcal{T}, \mu\mathcal{E}v$ after reduction by (DefinedYes).

If $x[M_1, \dots, M_m]$ is in the last case of Invariant 11, 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 $E, \sigma[\tilde{i}_j \mapsto \tilde{a}'], D_j \wedge M'_j, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{pk'}_{t_k}^* E, \sigma', r_{k'}, \mathcal{T}, \mu\mathcal{E}v$ where $r_{k'} = \text{true}$, $v_0 = (j, \tilde{a}') \in S$, $D_j \wedge M'_j = \text{defined}(M'_{j1}, \dots, M'_{jl_j}) \wedge M'_j$, 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 $E, \sigma', M'_{jk}, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{1}^* E, \sigma', a, \mathcal{T}, \mu\mathcal{E}v$ for some a , that is, $E, \sigma', M'_{jk} \Downarrow a$. The term $N = x[N_1, \dots, N_m]$ is a subterm of M'_{jk} . Therefore, by Lemma 12, for all $j' \leq m$, $E, \sigma', N_{j'} \Downarrow a_{j'}$ 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 $j' \leq m$, $E', \sigma', N_{j'} \Downarrow a_{j'}$ 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 13, for all $j' \leq m$, $E', \sigma, M_{j'} \Downarrow a_{j'}$ 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 11.

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 $j \leq m$, $E, \sigma, M_j \Downarrow a_j$ 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.

2.4.6 Typing

In this section, we show that our type system is compatible with the semantics of the oracles, 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 $\mathcal{E} \vdash P : T$, $\mathcal{E} \vdash Q$, and $\mathcal{E} \vdash M : T$ as in Section 2.3, with the additional rules $\mathcal{E} \vdash a : T$ if and only if $a \in T$, $\mathcal{E} \vdash \text{event_abort}(\mu, \tilde{a}) : e : T$ for all T , and $\mathcal{E} \vdash \text{abort} : T$ for all T . (These rules are useful to type evaluated terms and processes.)
- $\mathcal{E} \vdash (\sigma, P) : T$ if and only if $\mathcal{E}[i_1 \mapsto [1, n_1], \dots, i_m \mapsto [1, n_m]] \vdash P : T$ 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 $\mathcal{E} \vdash (\sigma, M) : T$ and $\mathcal{E} \vdash (\sigma, Q)$ are defined in the same way.
- $\mathcal{E} \vdash (\sigma, P) : T \mapsto T'$ if and only if $\mathcal{E} \vdash (\sigma, P) : T'$ and P is of the form $\text{let } x[\tilde{i}] : T = O[M_1, \dots, M_l](N) \text{ in } P_1 \text{ else } P_2$.
- $\mathcal{E} \vdash \mathcal{T}$ if and only if $\text{Tbl}(a_1, \dots, a_m) \in \mathcal{T}$ implies $\text{Tbl} : T_1 \times \dots \times T_m$ with for all $j \leq m$, $a_j \in T_j$.
- $\mathcal{E} \vdash \mu\mathcal{E}v$ if and only if $(\mu, \tilde{a}) : e(a_1, \dots, a_m) \in \mu\mathcal{E}v$ implies $e : T_1 \times \dots \times T_m$ with for all $j \leq m$, $a_j \in T_j$.
- $\mathcal{E} \vdash E, (\sigma_0, P_0) :: \dots :: (\sigma_m, P_m) :: \mathcal{A}, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ if and only if $\mathcal{E} \vdash E, (\sigma_0, P_0) : T_0$, for all $1 \leq i \leq m$, $\mathcal{E} \vdash (\sigma_i, P_i) : T_{i-1} \mapsto T_i$, $\mathcal{E} \vdash \mathcal{T}$, $\mathcal{E} \vdash \mu\mathcal{E}v$, and for all $(\sigma', Q) \in \mathcal{Q}$, $\mathcal{E} \vdash (\sigma', Q)$.
- $\mathcal{E} \vdash \mathcal{Q}, \mathcal{O}r$ if and only if for all $(\sigma', Q) \in \mathcal{Q}$, $\mathcal{E} \vdash (\sigma', Q)$.
- $\mathcal{E} \vdash E, \sigma, M : T, \mathcal{T}, \mu\mathcal{E}v$ if and only if $\mathcal{E} \vdash E$, $\mathcal{E} \vdash (\sigma, M) : T$, $\mathcal{E} \vdash \mathcal{T}$, and $\mathcal{E} \vdash \mu\mathcal{E}v$.

Lemma 15 *If $\mathcal{E} \vdash E, \sigma, M : T, \mathcal{T}, \mu\mathcal{E}v$ and $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', \sigma', M', \mathcal{T}', \mu\mathcal{E}v'$, then $\mathcal{E} \vdash E', \sigma', M' : T, \mathcal{T}', \mu\mathcal{E}v'$.*

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

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

Lemma 16 *If $\mathcal{E} \vdash \mathcal{Q}, \mathcal{O}r$ and $\mathcal{Q}, \mathcal{O}r \rightsquigarrow \mathcal{Q}', \mathcal{O}r'$, then $\mathcal{E} \vdash \mathcal{Q}', \mathcal{O}r'$.*

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

Proof sketch By cases on the derivation of $\mathcal{Q}, \mathcal{O}r \rightsquigarrow \mathcal{Q}', \mathcal{O}r'$. In the case of the replication, we have $\mathcal{E} \vdash (\sigma, \text{foreach } i \leq n \text{ do } Q)$, so $\mathcal{E}[i_1 \mapsto [1, n_1], \dots, i_m \mapsto [1, n_m]] \vdash \text{foreach } i \leq n \text{ do } Q$ 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 $a \in [1, n]$. \square

Lemma 17 *If $\mathcal{E} \vdash Q_0$, then $\mathcal{E} \vdash \text{initConfig}(Q_0, \mathcal{A})$.*

Proof sketch By Lemma 16 and the previous definitions. \square

Lemma 18 (Subject reduction) *If $\mathcal{E} \vdash E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ and $E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', St', \mathcal{Q}', \mathcal{O}r', \mathcal{T}', \mu\mathcal{E}v'$, then $\mathcal{E} \vdash E', St', \mathcal{Q}', \mathcal{O}r', \mathcal{T}', \mu\mathcal{E}v'$.*

Proof sketch By cases on the derivation of $E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p}_t E', St', \mathcal{Q}', \mathcal{O}r', \mathcal{T}', \mu\mathcal{E}v'$, using Lemmas 15 and 16. \square

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

As an immediate consequence of Lemmas 17, 18, and 15 and the observation above, we obtain: if Q_0 satisfies Invariant 8, 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 rules (IfT1), (IfT2), (If1), and (If2), the condition is in $\text{bool} = \{\text{false}, \text{true}\}$ (when it is a value, not an abort event value), so the condition $a \neq \text{true}$ is equivalent to $a = \text{false}$. In the rules for find, we have $r_k \in \{\text{false}, \text{true}\}$ when r_k is a value (not an abort event value). 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 $e : T_1 \times \dots \times T_l$. In the rule (Return), we have $d \in T$. Finally, in the rule (OracleCall), the set S of available oracles is equal to the one where we would not filter over the type of the oracle argument b .

2.5 Subset for the Initial Game

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 at its definition, we can also omit the indices.

Along similar lines, the oracles O in oracle declarations are used without indices, and the current replication indices are implicitly added by CryptoVerif. The construct `newOracle` and the oracle calls cannot occur in games manipulated by CryptoVerif. They are used only inside proofs. The grammar of the resulting syntax is summarized in Figure 13.

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

| | |
|--|--|
| $M, N ::=$ i $x[M_1, \dots, M_m]$ $f(M_1, \dots, M_m)$ $x \stackrel{R}{\leftarrow} T; N$ $\text{let } p = M \text{ in } N \text{ else } N'$ $\text{let } x : T = M \text{ in } N$ $\text{if } M \text{ then } N \text{ else } N'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } N_j) \text{ else } N'$ $\text{insert } Tbl(M_1, \dots, M_l); N$ $\text{get}[\text{unique?}] Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } N \text{ else } N'$ $\text{event } e(M_1, \dots, M_l); N$ $\text{event_abort } e$ | <p>terms</p> <ul style="list-style-type: none"> 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 |
| $p ::=$ $x : T$ $f(p_1, \dots, p_m)$ $= M$ | <p>pattern</p> <ul style="list-style-type: none"> variable function application comparison with a term |
| $Q ::=$ 0 $Q \mid Q'$ $\text{foreach } i \leq n \text{ do } Q$ $O(p) := P$ | <p>oracle definition</p> <ul style="list-style-type: none"> nil parallel composition replication n times oracle definition |
| $P ::=$ $\text{return } (N); Q$ $x \stackrel{R}{\leftarrow} T; P$ $\text{let } p = M \text{ in } P \text{ else } P'$ $\text{if } M \text{ then } P \text{ else } P'$ $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat}$ $\quad \text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P$ $\text{insert } Tbl(M_1, \dots, M_l); P$ $\text{get}[\text{unique?}] Tbl(p_1, \dots, p_l) \text{ suchthat } M \text{ in } P \text{ else } P'$ $\text{event } e(M_1, \dots, M_l); P$ $\text{event_abort } e$ yield | <p>oracle body</p> <ul style="list-style-type: none"> return random number assignment conditional array lookup insert in table get from table event event e and abort end |

Figure 13: Subset of the calculus for the initial game

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.
- Our compiler that translates CryptoVerif specifications into OCaml implementations [32] 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 with several incompatible definitions (different indices, different types, or variables defined in the same branch of a test) as well as variables declared without an explicit type are not allowed to occur in V nor in `defined` conditions of `find` or `if` and are renamed to distinct names, so that Invariant 1 is satisfied for these variables. CryptoVerif checks the rest of the invariants.

2.6 Input Language with Channels instead of Oracles

CryptoVerif also allows defining games in a channel-based process calculus inspired by the applied pi calculus [1, 2] and by the calculi of [59] and of [55]. This input language is mostly compatible with the symbolic protocol verifier ProVerif [28] (<http://proverif.inria.fr>), up to the following differences:

- ProVerif does not support `find`. Since ProVerif supports tables, one can encode most usages of `find` using tables instead for compatibility with ProVerif.
- ProVerif does not support channels with indices. As mentioned below, in CryptoVerif, the channels are indexed with the current replication indices. Since CryptoVerif automatically adds the current replication indices to channels without indices, one can easily use channels without indices.

So avoiding `find` and channels with indices allows us to have a language compatible with ProVerif for the definition of processes. (The assumptions on cryptographic primitives still need to be defined in a different way. In practice, these assumptions can be included in a library, using different libraries for ProVerif and for CryptoVerif.)

2.6.1 Input Language Description

Instead of defining oracles that the adversary can interact with, processes are now considered as communicating over a network with explicit input and output commands. The adversary has full control over the network: it receives any value output by a process and it is in charge of sending the input values to the processes. Inputs and outputs are made over channels, a set of constants often denoted by c .

Concretely, the following changes are made to the syntax previously defined in Figure 1:

- oracle definitions are called input processes and oracle bodies are called output processes;
- oracle declarations $O[\tilde{i}](p) := P$ do not exist anymore and are replaced with inputs $c[\tilde{i}](p); P$, where \tilde{i} are the current replication indices;
- return commands $\text{return } (N); Q$ are now output commands $\overline{c[\tilde{i}]}(N); Q$, where \tilde{i} are the current replication indices;
- the **yield** command is now considered as syntactic sugar for $\overline{\text{yield}}(\cdot)$, where *yield* is a specific channel not used elsewhere;
- oracle calls and oracle restrictions do no exist anymore. (Oracle calls can be encoded via channel inputs and outputs. We could introduce a channel restriction to match the oracle restriction, but it would be useful only in proofs, not in the input calculus of CryptoVerif.)

2.6.2 Translating from the Channel Calculus to the Oracle Calculus

Processes specified in the channel calculus are directly translated to the oracle calculus; this feature is only a form of syntactic sugar available for convenience. Notably, we do not define a formal semantics for the channel calculus, and all up to date theoretical developments of CryptoVerif are over the oracle calculus. For reference, the old semantics of the channel calculus variant can be found here [29]. We describe informally the translation here, which proceeds as follows:

- Any input of the form $c[\tilde{i}](p); P$ is translated as an oracle definition $O_c[\tilde{i}](p) := P$ where the oracle name O_c is built from the input channel c .
- Any output of the form $\overline{c[\tilde{i}]}(N); Q$ is translated as $\text{return } (N); Q$.

All input processes cannot be translated. We provide a list of invariants needed so that the translation preserves the semantics and produces valid instances of the oracle calculus. A first invariant is needed so that the translated process satisfies Invariant 7. We define the following multiset by induction over input and output processes, in a similar fashion to the previously defined $Defined^O$:

$$\begin{aligned}
Defined^c(0) &= \emptyset \\
Defined^c(Q_1 \mid Q_2) &= Defined^c(Q_1) \uplus Defined^c(Q_2) \\
Defined^c(\text{foreach } i \leq n \text{ do } Q) &= Defined^c(Q) \\
Defined^c(c[\tilde{i}](p); P) &= \{c\} \uplus Defined^c(P) \\
Defined^c(C[Q_1, \dots, Q_k]) &= \max_{j=1}^k Defined^c(Q_j)
\end{aligned}$$

where C is any oracle body context.

Channel Calculus Invariant 1 The input process Q_0 satisfies the Channel Invariant 1 if and only if

1. in every input and output over channel c in Q_0 , the indices \tilde{i} of c are the current replication indices at that declaration, and
2. $Defined^c(Q_0)$ does not contain duplicate elements.

In other words, Item 2 of Channel Invariant 1 means that inputs on the same channel occur in different branches of a `find`, `if` (or `let`), oracle call, or `get`.

A more complex issue is that in the channel calculus semantics, whenever an output can be received on the same channel by an input of the process, the output is probabilistically either caught by the attacker or directly transmitted to the input process. There is no equivalent to this probabilistic choice in the oracle setting. We thus add a channel invariant that makes it so that no output of the defined process can be received by the process itself. To this end, we define the following functions:

- $Ch_o(Q)$ is the set of pairs (c, k) of channels c that occur with k indices as output inside Q , that is, in a subprocess of the form $c[\tilde{i}](N); Q'$ inside Q , where the length of \tilde{i} is k ;
- $Ch_i(Q)$ is the set of pairs (c, k) of channels c that occur with k indices as input inside Q , that is, in a subprocess of the form $c[\tilde{i}](p); P$ inside Q , where the length of \tilde{i} is k ;
- $NoIntIO(Q)$ is true when no direct communication on a channel can occur inside Q , without going through the attacker:

$$\begin{aligned}
 NoIntIO(0) &= \text{true} \\
 NoIntIO(Q_1 \mid Q_2) &= NoIntIO(Q_1) \wedge NoIntIO(Q_2) \wedge \\
 &\quad (Ch_i(Q_1) \cap Ch_o(Q_2) = \emptyset) \wedge (Ch_i(Q_2) \cap Ch_o(Q_1) = \emptyset) \\
 NoIntIO(\text{foreach } i \leq n \text{ do } Q) &= NoIntIO(Q) \\
 NoIntIO(c[\tilde{i}](p); C[Q_1, \dots, Q_k]) &= \bigwedge_{j=1}^k NoIntIO(Q_j)
 \end{aligned}$$

where C is any oracle body context.

Channel Calculus Invariant 2 The input process Q_0 satisfies the Channel Invariant 2 if and only if $NoIntIO(Q_0) = \text{true}$.

This invariant guarantees that dynamically, there will be no input available in the process on a channel when the process performs an output on that channel. This is obtained by statically requiring that, when there is an output on channel c , the inputs on channel c with the same number of indices are syntactically above or under that output, or in a different branch, but not in a process in parallel.

When translating an output to a return, notice that the channel is dropped. The attacker might lose some distinguishing power if several outputs in distinct branches were made on distinct channels. We thus forbid this in the channel calculus.

Channel Calculus Invariant 3 The input process Q_0 satisfies the Channel Invariant 3 if and only if for all inputs on a given channel, all outputs directly below these inputs are over the same channel.

Most other previously defined invariants must also hold and can be seen as independent from oracles or channels.

Channel Calculus Invariant 4 The input process Q_0 satisfies the Channel Invariant 4 if only if it satisfies the Invariants 1 (with the minor modification that there are no oracle calls), 2, 3, 4, 5 and 6.

We do not explicitly state the typing Channel Invariant, which would correspond to Invariant 8. The old typing invariant for the channel calculus can be found in [29]. However, the translation to oracles strengthens this invariant: informally, inputs on the same channel must receive messages of same type (since they translate to oracle definitions of the same oracle) and outputs that follow inputs on the same channel must send messages of the same type (since they translate into returns of the same oracle).

2.7 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 returns values of interval types. The types of values chosen by $x[\tilde{i}] \stackrel{R}{\leftarrow} T$ are not interval types. The type T of the returned message in the (TRet) rule and of the receiving pattern in the (TODecl) rule are not interval types.

This property is satisfied by all games manipulated by CryptoVerif, but not by contexts that model parts of the adversary. Combined with Invariants 8, 2, and 5, it implies that the terms of variable accesses $x[M_1, \dots, M_m]$ contain only replication indices and variables. (Tuples, events, and tables can take interval types as arguments. The constraint on oracle declarations and returns could probably be relaxed.)

For oracles 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; see Section 5.2.)

Property 2 The construct `newOracle` O and the oracle calls do not appear in games.

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

Property 3 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, `expand_tables`, encodes `insert` and `get` using `find` (see Section 5.1.2). The constructs `insert` and `get` are never introduced by subsequent game transformations, so this property remains valid in the rest of the sequence.

Property 4 The variables defined in conditions of `find` have pairwise distinct names.

This property is enforced by the transformation `auto_SArename` (see Section 5.1.2) 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 4 is required as a precondition by many game transformations, and may be broken by game transformations that duplicate code. Therefore, we apply `auto_SArename` after these game transformations.

Property 5 The terms M are simple except for conditions of `find`.

| | |
|---|-------------------------------|
| $M, N ::=$ | terms |
| i | replication index |
| $x[M_1, \dots, M_m]$ | variable access |
| $f(M_1, \dots, M_m)$ | function application |
| $FC ::=$ | find condition |
| M | term |
| $x[\tilde{i}] \stackrel{R}{\leftarrow} T; FC$ | random number |
| let $p = M$ in FC else FC' | assignment (pattern-matching) |
| let $x[\tilde{i}] : T = M$ in N | assignment |
| if M then FC else FC' | conditional |
| find[<i>unique?</i>] ($\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j}$ <i>suchthat</i> defined(M_{j1}, \dots, M_{jl_j}) $\wedge FC'_j$ then FC_j) else FC'' | array lookup |
| event_abort e | event e and abort |
| $p ::=$ | pattern |
| $x[\tilde{i}] : T$ | variable |
| $f(p_1, \dots, p_m)$ | function application |
| $= M$ | comparison with a term |
| $Q ::=$ | oracle definition |
| 0 | nil |
| $Q \mid Q'$ | parallel composition |
| foreach $i \leq n$ do Q | replication n times |
| $O[\tilde{i}](p) := P$ | oracle declaration |
| $P ::=$ | oracle body |
| return (N); Q | return |
| $x[\tilde{i}] \stackrel{R}{\leftarrow} T; P$ | random number |
| let $p = M$ in P else P' | assignment |
| if M then P else P' | conditional |
| find[<i>unique?</i>] ($\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j}$ <i>suchthat</i> defined(M_{j1}, \dots, M_{jl_j}) $\wedge FC_j$ then P_j) else P | array lookup |
| event $e(M_1, \dots, M_l); P$ | event |
| event_abort e | event e and abort |
| yield | end |

Figure 14: Subset after game expansion

The grammar of the language taking into account this property as well as Properties 2 and 3 is shown in Figure 14. By Invariant 4, `event` does not occur in conditions of `find`, so `event` never occurs as term. Property 5 is enforced by the transformation **expand** (see Section 5.1.3) by converting other terms into processes. This transformation is applied on the initial game after **expand_tables**. Property 5 is broken by the cryptographic transformation of Section 5.2, so by default **expand** is called again after this transformation. Many game transformations require Property 5 as a precondition.

2.8 Security Properties, Indistinguishability

A context is a process containing a hole $[]$. An evaluation context C is a context built from $[]$, `newOracle` $O; C$, $Q \mid C$, and $C \mid Q$.

We use an evaluation context to represent an adversary in addition to the attacker \mathcal{A} . Adversarial contexts are useful, e.g., to push code inside the adversary in reduction proofs, but are also needed so that the attacker \mathcal{A} can raise events, access the public variables in V , and return its final result to a distinguisher by raising an event with that result as argument, by calling oracles of the context that execute these actions. We denote by $C[Q]$ the oracle definition obtained by replacing the hole $[]$ in the context C with the oracle definition Q .

We write $\text{event}(D)$ for the set of events that occur in the distinguisher D (i.e. are used by the distinguisher D). We write $\text{event}(Q)$ for the set of events that occur in the process Q . We use similar notations for oracle bodies, contexts, \dots We write $\text{event}(Q, Q')$ for $\text{event}(Q) \cup \text{event}(Q')$.

From now on, we consider in this section a fixed attacker space \mathcal{BB} , and all security definitions are given for this \mathcal{BB} . There must exist a function allowing to extract a run-time from the execution of an attacker. Classically, \mathcal{BB} would be the set of interactive probabilistic Turing machines, but the theory of CryptoVerif is sound for any kind of black-box interactive attacker, where the computational power is restricted by the cryptographic axioms. As cryptographic reductions sound for a black-box interactive attacker are also sound for quantum Turing Machines, this notably implies the post-quantum soundness of CryptoVerif, by considering as \mathcal{BB} the set of interactive quantum Turing machines.

Definition 5 (Indistinguishability) Let Q and Q' be two processes, V a set of variables, and \mathcal{E} a set of events. Assume that Q and Q' satisfy Invariants 1 to 8 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 8 with public variables V .

We write $Q \approx_p^{V, \mathcal{E}} Q'$ when, for all attackers $\mathcal{A} \in \mathcal{BB}$, 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 and such that $\text{event}(D) \cap \text{event}(Q, Q') \subseteq \mathcal{E}$, $|\Pr^{\mathcal{A}}[C[Q] : D] - \Pr^{\mathcal{A}}[C[Q'] : D]| \leq p(\mathcal{A}, C, t_D)$.

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

Property 6 All probabilities computed by CryptoVerif are built from the following components by mathematical operations:

- the runtime of the context plus the attacker;

- the maximum number of oracle calls made by the context or the attacker;
- the value of replication bounds, which is also determined from the number of calls performed by the context or attacker;
- the maximum length of the bitstring represented by a term, in particular a variable; this length may depend on messages computed by the context or the attacker; it is used only for unbounded types; for bounded types, we use the maximum length of the type instead;
- the maximum length of bitstring of a type T ;
- the length of the result of a function, expressed as a function of the length of its arguments;
- the time of some action, expressed as a function of other elements of the formula;
- probability functions, used in particular to express the probability of breaking each primitive from other elements of the formula;
- the cardinal $|T|$ of a type T ;
- the probability of collision between two random values of a type T , or between a random value and a value independent from that random value; these probabilities depend on the default distribution on the type D_T ;
- ϵ_T , the distance between the default distribution D_T of type T and the uniform distribution;
- ϵ_{find} , where the distance between $D_{\text{find}}(S)$ and the uniform distribution is $\epsilon_{\text{find}}/2$.

Among the elements above, the first four depend on the context and the attacker. In particular, probability formulas output by CryptoVerif do not depend on the variable, table, event names in the context. They also do not depend on the values of variables, but may depend on their length. For variables of bounded types, the probabilities do not depend at all on the values. We do not in practice distinguish between the calls of the context or the attacker.

The set of events \mathcal{E} corresponds to events that the adversary is allowed to observe. When $\mathcal{E} = \text{event}(Q, Q')$, we omit it and write $Q \approx_p^V Q'$.

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 the finite set of public variables 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 $Q \approx_p^\mathcal{E} Q'$.

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 5:

- Lemma 19**
1. *Reflexivity:* $Q \approx_0^{V, \mathcal{E}} Q$.
 2. *Symmetry:* If $Q \approx_p^{V, \mathcal{E}} Q'$, then $Q' \approx_p^{V, \mathcal{E}} Q$.
 3. *Transitivity:* If $Q \approx_p^{V, \mathcal{E}} Q'$ and $Q' \approx_{p'}^{V, \mathcal{E}} Q''$, then $Q \approx_{p+p'}^{V, \mathcal{E}} Q''$.

4. *Application of a context:* If $Q \approx_p^{V, \mathcal{E}} Q'$ and C is an evaluation context acceptable for Q and Q' with public variables V , then $C[Q] \approx_p^{V', \mathcal{E}'} C[Q']$, where $p'(\mathcal{A}, C', t_D) = p(\mathcal{A}, C'[C[\]])$, $V' \subseteq V \cup \text{var}(C)$, and $\mathcal{E}' = \mathcal{E} \cup (\text{event}(C) \setminus \text{event}(Q, Q'))$.

Next, we introduce a notion related to indistinguishability that treats Shoup and non-unique events specially.

Definition 6 (Property preservation 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 8 with public variables V , and the variables of V are defined in Q and Q' , with the same types.

Let $D_{\text{false}}(\mathcal{E}v) = \text{false}$ for all $\mathcal{E}v$. Let $\text{NonUnique}_Q = \bigvee \{e \mid [\text{unique}_e] \text{ occurs in } Q\}$ and $\text{NonUnique}_{Q,D} = \bigvee \{e \mid [\text{unique}_e] \text{ occurs in } Q, e \notin D\}$, where D is a distinguisher consisting of a disjunction of Shoup and non-unique events, and we write $e \notin D$ to say that e does not occur in this disjunction. We have $\text{NonUnique}_{Q,D} = \text{NonUnique}_Q \wedge \neg D$.

We write $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V} Q', D', \text{EvUsed}'$ when \mathcal{D} is a set of distinguishers, \mathcal{D}_{SNU} is a set of Shoup and non-unique events in EvUsed , D and D' are distinguishers consisting of a disjunction of Shoup and non-unique events, the events that occur in Q or in D are in EvUsed , $\text{EvUsed} \subseteq \text{EvUsed}'$, the events that occur in Q' but not in Q are in EvUsed' but not in EvUsed , the events that occur in D' but not in D are in EvUsed' but not in EvUsed , and, for all $\mathcal{A} \in \mathcal{BB}$, for all evaluation contexts C acceptable for Q and Q' with public variables V that do not contain events in EvUsed' , all distinguishers $D_0 \in \mathcal{D} \cup \{D_{\text{false}}\}$ that run in time at most t_{D_0} , all distinguishers D_1 that are disjunctions of events in \mathcal{D}_{SNU} ,

$$\begin{aligned} \Pr^{\mathcal{A}}[C[Q] : (D_0 \vee D_1 \vee D) \wedge \neg \text{NonUnique}_{Q, D_1 \vee D}] \\ \leq \Pr^{\mathcal{A}}[C[Q'] : (D_0 \vee D_1 \vee D') \wedge \neg \text{NonUnique}_{Q', D_1 \vee D'}] + p(\mathcal{A}, C, t_{D_0}) \end{aligned} \quad (1)$$

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 $D_0 \in \mathcal{D} \cup \{D_{\text{false}}\}$ corresponds to the initial query to prove: it is

- a correspondence distinguisher $\neg\varphi$ for a correspondence property (see Section 2.8.3);
- S or $\neg\mathsf{S}$ for (one-session) secrecy (see Section 2.8.1) and bit secrecy (see Section 2.8.2);
- any distinguisher for indistinguishability;
- D_{false} when the initial query has already been proved, and only Shoup and non-unique events remain to be proved.

We need to specify precisely the distinguishers needed for the queries we want to prove, because some game transformations of CryptoVerif rely on that. For instance, **simplify** (Section 5.1.21) removes events that are not used by the queries.

The formula D_1 is a disjunction of Shoup and non-unique events that remain to be proved, both in Q and in Q' . These events are in \mathcal{D}_{SNU} and in EvUsed . The formula D is a disjunction of Shoup and non-unique events that remain to be proved in Q , while the formula D' is a disjunction of Shoup and non-unique events that remain to be proved in Q' . Hence, the events that occur in D and not in D' are events proved while transforming Q into Q' . (“Proving” an event means proving that this event has a negligible probability of occurring, and adding that probability to p .) In contrast, the events that occur in D' and not in D are fresh Shoup and non-unique events introduced during the transformation of Q into Q' , and that will need to be proved later; 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 $EvUsed$.

When there are no Shoup nor non-unique events, we have $D_1 = D = D' = D_{\text{false}}$ and $\text{NonUnique}_{Q,D_1 \vee D} = \text{NonUnique}_{Q',D_1 \vee D'} = D_{\text{false}}$, so the inequality (1) reduces to

$$\Pr^{\mathcal{A}}[C[Q] : D_0] \leq \Pr^{\mathcal{A}}[C[Q'] : D_0] + p(\mathcal{A}, C, t_{D_0})$$

Using $\neg D_0$ instead of D_0 , we obtain

$$1 - \Pr^{\mathcal{A}}[C[Q] : D_0] \leq 1 - \Pr^{\mathcal{A}}[C[Q'] : D_0] + p(\mathcal{A}, C, t_{D_0})$$

so by combining the two, $|\Pr^{\mathcal{A}}[C[Q] : D_0] - \Pr^{\mathcal{A}}[C[Q'] : D_0]| \leq p(\mathcal{A}, C, t_{D_0})$ as in the definition of indistinguishability. In the general case, the inequality (1) differs from this formula because (1) always counts the traces that execute Shoup and non-unique events that remain to be proved (these traces are always included in the probability by $D_1 \vee D$, resp. $D_1 \vee D'$; the probability of these events needs to be bounded), and never counts the traces that execute proved non-unique events (these traces are excluded by $\neg \text{NonUnique}_{Q,D_1 \vee D}$, resp. $\neg \text{NonUnique}_{Q',D_1 \vee D'}$; the probability of these events has already been bounded). We could also exclude traces that execute proved Shoup events, though it is less essential: Shoup events often simply disappear when they are proved, while non-unique events remain in the game. Excluding traces that execute proved non-unique events allows us to exploit that the corresponding find or get is unique in the transformation from Q to Q' : intuitively, the traces in which the find or get is not unique are not counted, so they can be ignored. (Obviously, this point needs to be proved more precisely for each game transformation.)

We need to introduce a distinct event for each $[\text{unique}_e]$ because not all $\text{find}[\text{unique}_e]$ and $\text{get}[\text{unique}_e]$ may be proved unique in the current process and because we need to distinguish the non-unique events that occur in Q from those that occur in the context C . (We note that the attacker \mathcal{A} cannot raise events itself.) However, once a $[\text{unique}_e]$ is proved, its name does not matter: all such events are counted in $\text{NonUnique}_{Q,D_1 \vee D}$, and that remains true in future game transformations, since events are never re-added to D_1 or D . Therefore, we can rename all such events to the same name. So we simply abbreviate $[\text{unique}_e]$ by $[\text{unique}]$ when event e is proved. The context C must not contain the events used in Q or Q' and more generally events in $EvUsed'$.

The formula (1) could also be written

$$\begin{aligned} \Pr^{\mathcal{A}}[C[Q] : (D_0 \wedge \neg \text{NonUnique}_Q) \vee D_1 \vee D] \\ \leq \Pr^{\mathcal{A}}[C[Q'] : (D_0 \wedge \neg \text{NonUnique}_{Q'}) \vee D_1 \vee D'] + p(\mathcal{A}, C, t_{D_0}) \end{aligned}$$

since $\text{NonUnique}_{Q,D_1 \vee D} = \text{NonUnique}_Q \wedge \neg(D_1 \vee D)$. The advantage of the latter formulation is that the dependency in D_1, D, D' is simpler, which we sometimes exploit in the proofs. However, its drawback is that it is less clear for which events the traces are always counted and for which ones they are never counted, because some events appear both in NonUnique_Q and in $D_1 \vee D$. That is why we chose formula (1), to make it clear that the traces that execute events in $D_1 \vee D$ are always counted and the traces that execute events in $\text{NonUnique}_{Q,D_1 \vee D}$ are never counted.

When f is a function from sequences of events to sequences of events and D is a distinguisher, we define the distinguisher $D \circ f$ by $(D \circ f)(\mathcal{E}v) = D(f(\mathcal{E}v))$. In particular, when σ is a renaming of events, the distinguisher $D \circ \sigma^{-1}$ is defined by $(D \circ \sigma^{-1})(\mathcal{E}v) = D(\sigma^{-1}\mathcal{E}v)$. When D is defined as a logical formula, that corresponds to renaming the events in D . For instance, if $D = e_1 \vee \dots \vee e_m$, then $D \circ \sigma^{-1} = \sigma e_1 \vee \dots \vee \sigma e_m$. When \mathcal{D} is a set of distinguishers, $\sigma \mathcal{D} = \{D \circ \sigma^{-1} \mid D \in \mathcal{D}\}$.

We write $\mathcal{D}_{-\mathcal{E}}$ for the set of distinguishers that do not use events in \mathcal{E} .

Lemma 20 1. *Link with indistinguishability:*

- (a) Suppose that \mathcal{D}_{SNU} is a set of Shoup events, D is a distinguisher consisting of a disjunction of Shoup events, Q, Q' do not contain non-unique events, the events that occur in Q, D, \mathcal{D} , or \mathcal{D}_{SNU} are in EvUsed , and the events that occur in Q' also occur in Q . If $Q \approx_p^{V, \mathcal{E}} Q'$, then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_p Q', D, \text{EvUsed}$, for all $\mathcal{D}, \mathcal{D}_{\text{SNU}}, D$ such that $\text{event}(\mathcal{D}) \cap \text{EvUsed} \subseteq \mathcal{E}$, $\mathcal{D}_{\text{SNU}} \subseteq \mathcal{E}$, $\text{event}(D) \subseteq \mathcal{E}$.
- (b) 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 $\mathcal{D}_{-(\text{EvUsed} \setminus \mathcal{E})}, \emptyset : Q, D_{\text{false}}, \text{EvUsed} \xrightarrow{V}_p Q', D_{\text{false}}, \text{EvUsed}$ if and only if $Q \approx_p^{V, \mathcal{E}} Q'$.
2. *Reflexivity:* If \mathcal{D}_{SNU} is a set of Shoup and non-unique events in EvUsed , D is a distinguisher consisting of a disjunction of Shoup and non-unique events, and the events that occur in Q or in D are in EvUsed , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_0 Q, D, \text{EvUsed}$.
3. *Transitivity:* If $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_p Q', D', \text{EvUsed}'$ and $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q', D', \text{EvUsed}' \xrightarrow{V}_{p'} Q'', D'', \text{EvUsed}''$, then we have $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_{p''} Q'', D'', \text{EvUsed}''$, where $p''(\mathcal{A}, C, t_{D_0}) = p(\mathcal{A}, C, t_{D_0}) + p'(\mathcal{A}, C, t_{D_0})$.
4. *Application of a context:* If $\mathcal{D}_{-\text{EvUsed}'}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_p Q', D', \text{EvUsed}'$, σ is a renaming of the events in EvUsed' to events not in EvUsed^+ , C is a context acceptable for σQ and $\sigma Q'$ with public variables V such that $\text{event}(C) \subseteq \text{EvUsed}^+$, and $\mathcal{D}'_{\text{SNU}}$ is a set of Shoup and non-unique events in EvUsed^+ , then we have $\mathcal{D}_{-\sigma \text{EvUsed}'}, \sigma \mathcal{D}_{\text{SNU}} \cup \mathcal{D}'_{\text{SNU}} : C[\sigma Q], D \circ \sigma^{-1}, \sigma \text{EvUsed} \cup \text{EvUsed}^+ \xrightarrow{V}_{p'} C[\sigma Q'], D' \circ \sigma^{-1}, \sigma \text{EvUsed}' \cup \text{EvUsed}^+$, where $p'(\mathcal{A}, C', t_{D_0}) = p(\mathcal{A}, \sigma^{-1}(C'[C[]]), t_{D_0})$, and $V' \subseteq V \cup \text{var}(C)$.
5. *Adding distinguishers:* If $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_p Q', D', \text{EvUsed}'$ and $e \in \mathcal{D}_{\text{SNU}}$, then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D \vee e, \text{EvUsed} \xrightarrow{V}_p Q', D' \vee e, \text{EvUsed}'$.
6. *Removing distinguishers:* If $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_p Q', D', \text{EvUsed}'$, $\mathcal{D}' \subseteq \mathcal{D}$, and $\mathcal{D}'_{\text{SNU}} \subseteq \mathcal{D}_{\text{SNU}}$, then $\mathcal{D}', \mathcal{D}'_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_p Q', D', \text{EvUsed}'$.

Proof Property 1a: Given the hypothesis, $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_p Q', D, \text{EvUsed}$ reduces to: for all attackers $\mathcal{A} \in \mathcal{BB}$, for all evaluation contexts C acceptable for Q and Q' with public variables V that do not contain events in EvUsed , all distinguishers $D_0 \in \mathcal{D} \cup \{D_{\text{false}}\}$ that run in time at most t_{D_0} , and all distinguishers D_1 that are disjunctions of events in \mathcal{D}_{SNU} ,

$$\Pr^{\mathcal{A}}[C[Q] : D_0 \vee D_1 \vee D] \leq \Pr^{\mathcal{A}}[C[Q'] : D_0 \vee D_1 \vee D] + p(\mathcal{A}, C, t_{D_0})$$

We have $\text{event}(D_0 \vee D_1 \vee D) \cap \text{event}(Q, Q') \subseteq \text{event}(D_0 \vee D_1 \vee D) \cap \text{EvUsed} \subseteq \mathcal{E}$. Moreover, $D_0 \vee D_1 \vee D$ can be implemented in the same time as D_0 since evaluating $D_0 \vee D_1 \vee D$ can be done by setting the final result to true as soon as an event in $D_1 \vee D$ is executed, and evaluating D_0 otherwise. This does not take more time than evaluating D_0 . So this inequality is a consequence of $Q \approx_p^{V, \mathcal{E}} Q'$.

Property 1b: Given the hypothesis, $\mathcal{D}_{-(\text{EvUsed} \setminus \mathcal{E})}, \emptyset : Q, D_{\text{false}}, \text{EvUsed} \xrightarrow{V}_p Q', D_{\text{false}}, \text{EvUsed}$ reduces to: for all attackers $\mathcal{A} \in \mathcal{BB}$, for all evaluation contexts C acceptable for Q and Q' with public variables V that do not contain events in EvUsed , and all distinguishers $D_0 \in \mathcal{D}_{-(\text{EvUsed} \setminus \mathcal{E})}$ that run in time at most t_{D_0} ,

$$\Pr^{\mathcal{A}}[C[Q] : D_0] \leq \Pr^{\mathcal{A}}[C[Q'] : D_0] + p(\mathcal{A}, C, t_{D_0}),$$

that is,

$$\Pr^{\mathcal{A}}[C[Q] : D_0] - \Pr^{\mathcal{A}}[C[Q'] : D_0] \leq p(\mathcal{A}, C, t_{D_0}). \quad (2)$$

We have $\text{event}(D_0) \cap \text{event}(Q, Q') \subseteq \text{event}(D_0) \cap \text{EvUsed} \subseteq \mathcal{E}$. Therefore, $Q \approx_p^{V, \mathcal{E}} Q'$ implies $\mathcal{D}_{\neg(\text{EvUsed} \setminus \mathcal{E}), \emptyset} : Q, D_{\text{false}}, \text{EvUsed} \xrightarrow{V}_p Q', D_{\text{false}}, \text{EvUsed}$.

Conversely, assume $\mathcal{D}_{\neg(\text{EvUsed} \setminus \mathcal{E}), \emptyset} : Q, D_{\text{false}}, \text{EvUsed} \xrightarrow{V}_p Q', D_{\text{false}}, \text{EvUsed}$. Let $\mathcal{A} \in \mathcal{BB}$ be an attacker, Let C be an evaluation context C acceptable for Q and Q' with public variables V and D be a distinguisher such that $\text{event}(D) \cap \text{event}(Q, Q') \subseteq \mathcal{E}$. Let σ be a renaming of events in EvUsed to fresh events. Then σC does not contain events in EvUsed . Given a sequence of events $\mathcal{E}v$, let $f(\mathcal{E}v)$ be obtained by removing all events in $\text{EvUsed} \setminus \mathcal{E}$ from $\mathcal{E}v$ and, in the remaining sequence, renaming the events e in σEvUsed to $\sigma^{-1}(e)$. Let $D_0 = D \circ f$. By construction, $D_0 \in \mathcal{D}_{\neg(\text{EvUsed} \setminus \mathcal{E})}$. So by (2), we get

$$\Pr^{\mathcal{A}}[(\sigma C)[Q] : D_0] - \Pr^{\mathcal{A}}[(\sigma C)[Q'] : D_0] \leq p(\mathcal{A}, C, t_{D_0}).$$

Since $\neg D_0 \in \mathcal{D}_{\neg(\text{EvUsed} \setminus \mathcal{E})}$ and runs in the same time as D_0 ,

$$1 - \Pr^{\mathcal{A}}[(\sigma C)[Q] : D_0] - 1 + \Pr^{\mathcal{A}}[(\sigma C)[Q'] : D_0] \leq p(\mathcal{A}, C, t_{D_0})$$

so

$$|\Pr^{\mathcal{A}}[(\sigma C)[Q] : D_0] - \Pr^{\mathcal{A}}[(\sigma C)[Q'] : D_0]| \leq p(\mathcal{A}, C, t_{D_0}).$$

Furthermore, each trace of $(\sigma C)[Q]$ with sequence of events $\mathcal{E}v$ corresponds to a trace of $C[Q]$ with the same probability and a sequence of events $\mathcal{E}v'$ equal to $f(\mathcal{E}v)$ plus some events not used by D , so $D_0(\mathcal{E}v) = D(f(\mathcal{E}v)) = D(\mathcal{E}v')$. Indeed, in a trace of $(\sigma C)[Q]$, if an event $e(\dots)$ is executed in Q , then it is executed by $C[Q]$ as well. If it is in \mathcal{E} , then it is left unchanged by f . If it is not in \mathcal{E} , then it is in $\text{event}(Q) \setminus \mathcal{E} \subseteq \text{EvUsed} \setminus \mathcal{E}$, so it is removed by f ; furthermore, this event is not used by D since $\text{event}(D) \cap \text{event}(Q) \subseteq \mathcal{E}$. If an event $e(\dots)$ is executed in σC , then either $e \in \sigma \text{EvUsed}$, $e'(\dots)$ is executed by $C[Q]$ with $e' = \sigma^{-1}(e)$, and f maps e to e' ; or $e \notin \sigma \text{EvUsed}$, $e \notin \text{EvUsed}$, $e(\dots)$ is executed by $C[Q]$, and f leaves e unchanged. Therefore, $\Pr^{\mathcal{A}}[(\sigma C)[Q] : D_0] = \Pr^{\mathcal{A}}[C[Q] : D]$. We have a similar situation for Q' instead of Q , and D can be implemented in the same time as D_0 , so

$$|\Pr^{\mathcal{A}}[C[Q] : D] - \Pr^{\mathcal{A}}[C[Q'] : D]| \leq p(\mathcal{A}, C, t_D)$$

so $Q \approx_p^{V, \mathcal{E}} Q'$.

Property 2: Obvious.

Property 3: The events in Q or D are in EvUsed . We have $\text{EvUsed} \subseteq \text{EvUsed}' \subseteq \text{EvUsed}''$. 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 $\text{EvUsed}'' \setminus \text{EvUsed}'$; in the latter case, they are in $\text{EvUsed}' \setminus \text{EvUsed}$; so in both cases they are in $\text{EvUsed}'' \setminus \text{EvUsed}$. The same reasoning applies for the events that occur in D'' but not in D .

Let $\mathcal{A} \in \mathcal{BB}$ be an attacker, C be any evaluation context acceptable for Q and Q'' with public variables V that does not contain events in EvUsed'' . After renaming the variables of C that do not occur in Q and Q'' and the tables of C that do not occur in Q and Q'' so that they do not occur in Q' , C is also acceptable for Q' with public variables V . Furthermore, by Property 6, this renaming does not change the probabilities. Since $\text{EvUsed}' \subseteq \text{EvUsed}''$, C does not contain events in EvUsed' .

Let $D_0 \in \mathcal{D} \cup \{D_{\text{false}}\}$ that runs in time at most t_{D_0} . Let D_1 be a disjunction of events of \mathcal{D}_{SNU} . Then we have:

$$\begin{aligned}
& \Pr^{\mathcal{A}}[C[Q] : (D_0 \vee D_1 \vee D)] \wedge \neg \text{NonUnique}_{Q, D_1 \vee D}] \\
& \leq \Pr^{\mathcal{A}}[C[Q'] : (D_0 \vee D_1 \vee D') \wedge \neg \text{NonUnique}_{Q', D_1 \vee D'}] + p(\mathcal{A}, C, t_{D_0}) \\
& \leq \Pr^{\mathcal{A}}[C[Q''] : (D_0 \vee D_1 \vee D'') \wedge \neg \text{NonUnique}_{Q'', D_1 \vee D''}] + p(\mathcal{A}, C, t_{D_0}) + p'(\mathcal{A}, C, t_{D_0}) \\
& \leq \Pr^{\mathcal{A}}[C[Q''] : (D_0 \vee D_1 \vee D'') \wedge \neg \text{NonUnique}_{Q'', D_1 \vee D''}] + p''(\mathcal{A}, C, t_{D_0})
\end{aligned}$$

by definition of p'' . Therefore, we have $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, EvUsed \xrightarrow{V} p'' Q'', D'', EvUsed''$.

Property 4: $\sigma \mathcal{D}_{\text{SNU}}$ is a set of Shoup and non-unique events in $\sigma EvUsed \subseteq \sigma EvUsed \cup EvUsed^+$, so $\sigma \mathcal{D}_{\text{SNU}} \cup \mathcal{D}'_{\text{SNU}}$ is a set of Shoup and non-unique events in $\sigma EvUsed \cup EvUsed^+$. The events that occur in $C[\sigma Q]$ are either in C or in σQ ; the former case, they are in $EvUsed^+$ by hypothesis; in the latter case, they are also in $\sigma EvUsed$, since the events of Q are in $EvUsed$. So the events that occur in $C[\sigma Q]$ are in $\sigma EvUsed \cup EvUsed^+$. The events in $D \circ \sigma^{-1}$ are in $\sigma EvUsed \subseteq \sigma EvUsed \cup EvUsed^+$. We have $EvUsed \subseteq EvUsed'$, so $\sigma EvUsed \cup EvUsed^+ \subseteq \sigma EvUsed' \cup EvUsed^+$. The events that occur in $C[\sigma Q']$ but not in $C[\sigma Q]$ are in $\sigma Q'$ but not in σQ , so they are in $\sigma EvUsed' \setminus \sigma EvUsed$, so in $(\sigma EvUsed' \cup EvUsed^+) \setminus (\sigma EvUsed \cup EvUsed^+)$. Similarly, the events that occur in $D' \circ \sigma^{-1}$ but not in $D \circ \sigma^{-1}$ are in $\sigma EvUsed' \setminus \sigma EvUsed$, so they are in $(\sigma EvUsed' \cup EvUsed^+) \setminus (\sigma EvUsed \cup EvUsed^+)$.

Let C' be any evaluation context acceptable for $C[\sigma Q]$ and $C[\sigma Q']$ with public variables V' that does not contain events in $\sigma EvUsed' \cup EvUsed^+$. We rename the variables of C' not in V' so that they are not in V ; by Property 6, this renaming does not change the probabilities. Then $\sigma^{-1}(C'[C[\]])$ is an evaluation context acceptable for Q and Q' with public variables V . Indeed,

$$\begin{aligned}
\text{var}(\sigma^{-1}(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) \quad \text{since } V' \subseteq V \cup \text{var}(C) \\
&\subseteq V \quad \text{since } \text{var}(C) \cap \text{var}(Q) \subseteq V
\end{aligned}$$

We have similarly $\text{var}(\sigma^{-1}(C'[C[\]])) \cap \text{var}(Q') \subseteq V$. We also have $\text{vardef}(\sigma^{-1}(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 σQ 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 $C[\sigma Q]$ and with public variables V' . Moreover, C and σQ do not use any common table, and C' and $C[\sigma Q]$ do not use any common table so a fortiori C' and σQ do not use any common table. Therefore, $C'[C[\]]$ and σQ do not use any common table, so $\sigma^{-1}(C'[C[\]])$ and Q do not use any common table. Similarly, $\sigma^{-1}(C'[C[\]])$ and Q' do not use any common table. The context $\sigma^{-1}(C'[C[\]])$ does not contain events in $EvUsed'$, since C' and C do not contain events in $\sigma EvUsed'$, because C' does not contain events in $\sigma EvUsed' \cup EvUsed^+$ and the events of C are in $EvUsed^+$ which is disjoint from $\sigma EvUsed'$. (The renamings σ and σ^{-1} are bijections, so for instance σ maps the fresh events introduced by σ to $EvUsed'$ and σ^{-1} maps $EvUsed'$ to the fresh events introduced by σ .)

By using the property $\mathcal{D}_{\neg EvUsed'}, \mathcal{D}_{\text{SNU}} : Q, D, EvUsed \xrightarrow{V} p Q', D', EvUsed'$ with the context $\sigma^{-1}(C'[C[\]])$, we get for all attackers $\mathcal{A} \in \mathcal{B}\mathcal{B}$ and for any distinguishers $D_0 \in \mathcal{D}_{\neg EvUsed'} \cup \{D_{\text{false}}\}$

that runs in time t_{D_0} and D_1 disjunction of events in \mathcal{D}_{SNU} :

$$\begin{aligned} & \Pr^{\mathcal{A}}[\sigma^{-1}C'[\sigma^{-1}C[Q]] : (D_0 \vee D_1 \vee D) \wedge \neg\text{NonUnique}_{Q,D_1 \vee D}] \\ & \leq \Pr^{\mathcal{A}}[\sigma^{-1}C'[\sigma^{-1}C[Q']] : (D_0 \vee D_1 \vee D') \wedge \neg\text{NonUnique}_{Q',D_1 \vee D'}] + p(\mathcal{A}, \sigma^{-1}(C'[C[]]), t_{D_0}) \end{aligned} \quad (3)$$

Let $D'_0 \in \mathcal{D}_{\neg\sigma\text{EvUsed}'} \cup \{D_{\text{false}}\}$ that runs in time at most t_{D_0} , Let D'_1 be a disjunction of events in $\sigma\mathcal{D}_{\text{SNU}} \cup \mathcal{D}'_{\text{SNU}}$. We can write D'_1 under the form $D'_1 = D'_2 \vee D'_3$ where D'_2 is a disjunction of events in $\sigma\mathcal{D}_{\text{SNU}}$ and D'_3 is a disjunction of events in $\mathcal{D}'_{\text{SNU}}$.

By applying (3) to $D_0 = (D'_0 \circ \sigma \vee D'_3 \circ \sigma) \wedge \neg\text{NonUnique}_{\sigma^{-1}C,D'_3 \circ \sigma}$, which uses events not in EvUsed' , and to $D_1 = D'_2 \circ \sigma$, we get for all $\mathcal{A} \in \mathcal{BB}$

$$\begin{aligned} & \Pr^{\mathcal{A}}[\sigma^{-1}C'[\sigma^{-1}C[Q]] : (((D'_0 \circ \sigma \vee D'_3 \circ \sigma) \wedge \neg\text{NonUnique}_{\sigma^{-1}C,D'_3 \circ \sigma}) \vee D'_2 \circ \sigma \vee D)] \\ & \quad \wedge \neg\text{NonUnique}_{Q,D'_2 \circ \sigma \vee D}] \\ & \leq \Pr^{\mathcal{A}}[\sigma^{-1}C'[\sigma^{-1}C[Q']] : (((D'_0 \circ \sigma \vee D'_3 \circ \sigma) \wedge \neg\text{NonUnique}_{\sigma^{-1}C,D'_3 \circ \sigma}) \vee D'_2 \circ \sigma \vee D')] \\ & \quad \wedge \neg\text{NonUnique}_{Q',D'_2 \circ \sigma \vee D'}] + p(\mathcal{A}, \sigma^{-1}(C'[C[]]), t_{D'_0}) \end{aligned}$$

since D_0 can be implemented to run in the same time as D'_0 . By applying σ , we have for all $\mathcal{A} \in \mathcal{BB}$

$$\begin{aligned} & \Pr^{\mathcal{A}}[C'[C[\sigma Q]] : (((D'_0 \vee D'_3) \wedge \neg\text{NonUnique}_{C,D'_3}) \vee D'_2 \vee D \circ \sigma^{-1}) \wedge \neg\text{NonUnique}_{\sigma Q,D'_2 \vee D \circ \sigma^{-1}}] \\ & \leq \Pr^{\mathcal{A}}[C'[C[\sigma Q']] : (((D'_0 \vee D'_3) \wedge \neg\text{NonUnique}_{C,D'_3}) \vee D'_2 \vee D' \circ \sigma^{-1}) \\ & \quad \wedge \neg\text{NonUnique}_{\sigma Q',D'_2 \vee D' \circ \sigma^{-1}}] + p(\mathcal{A}, \sigma^{-1}(C'[C[]]), t_{D'_0}) \end{aligned}$$

Since the events of C are in EvUsed^+ , the events of $\text{NonUnique}_{C,D'_3}$ are disjoint from those in $(D'_2 \vee D \circ \sigma^{-1})$, so $(\neg\text{NonUnique}_{C,D'_3}) \vee (D'_2 \vee D \circ \sigma^{-1}) = \neg\text{NonUnique}_{C,D'_3}$ and similarly $(\neg\text{NonUnique}_{C,D'_3}) \vee (D'_2 \vee D' \circ \sigma^{-1}) = \neg\text{NonUnique}_{C,D'_3}$. So we have for all $\mathcal{A} \in \mathcal{BB}$

$$\begin{aligned} & \Pr^{\mathcal{A}}[C'[C[\sigma Q]] : (D'_0 \vee D'_3 \vee D'_2 \vee D \circ \sigma^{-1}) \wedge \neg\text{NonUnique}_{C,D'_3} \wedge \neg\text{NonUnique}_{\sigma Q,D'_2 \vee D \circ \sigma^{-1}}] \\ & \leq \Pr^{\mathcal{A}}[C'[C[\sigma Q']] : (D'_0 \vee D'_3 \vee D'_2 \vee D' \circ \sigma^{-1}) \wedge \neg\text{NonUnique}_{C,D'_3} \\ & \quad \wedge \neg\text{NonUnique}_{\sigma Q',D'_2 \vee D' \circ \sigma^{-1}}] + p(\mathcal{A}, \sigma^{-1}(C'[C[]]), t_{D'_0}) \end{aligned}$$

that is

$$\begin{aligned} & \Pr^{\mathcal{A}}[C'[C[\sigma Q]] : (D'_0 \vee D'_1 \vee D \circ \sigma^{-1}) \wedge \neg\text{NonUnique}_{C[\sigma Q],D'_1 \vee D \circ \sigma^{-1}}] \\ & \leq \Pr^{\mathcal{A}}[C'[C[\sigma Q']] : (D'_0 \vee D'_1 \vee D' \circ \sigma^{-1}) \wedge \neg\text{NonUnique}_{C[\sigma Q'],D'_1 \vee D' \circ \sigma^{-1}}] \\ & \quad + p(\mathcal{A}, \sigma^{-1}(C'[C[]]), t_{D'_0}) \end{aligned}$$

Properties 5 and 6: Obvious. \square

When CryptoVerif transforms a game G into a game G' , in most cases, we have $G \approx_p^{V,\mathcal{E}} G'$, where p is the probability difference coming from the transformation, and computed by CryptoVerif, which implies $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V} G', D, \text{EvUsed}$, for all $\mathcal{D}, \mathcal{D}_{\text{SNU}}, D$ such that $\text{event}(\mathcal{D}) \cap \text{EvUsed} \subseteq \mathcal{E}$, $\mathcal{D}_{\text{SNU}} \subseteq \mathcal{E}$, $\text{event}(D) \subseteq \mathcal{E}$ by Lemma 20, Property 1a. However, there are exceptions to this situation:

- transformations that exploit the uniqueness of `find[uniquee]` or `get[uniquee]`, which are valid only when event e is not executed. These events are taken into account by `NonUniqueQ,D`.
- transformations that insert events using Shoup's lemma. This is the case of the transformations `insert_event` (see Section 5.1.12) and `insert` (see Section 5.1.13). In this case, we have $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D_{\text{false}}, \text{EvUsed} \xrightarrow{V}_p G', e, \text{EvUsed} \cup \{e\}$ where e is the introduced event. The addition of Shoup events may also be combined with the cryptographic transformation of Section 5.2, for example for specifying the decisional Diffie-Hellman assumption. In general, the cryptographic axioms are of the form

$$\mathcal{D}_{\neg \text{EvUsed}_R}, \emptyset : L, D_{\text{false}}, \emptyset \rightarrow_p R, D_R, \text{EvUsed}_R$$

where L does not contain events, $D_R = \bigvee \{e \mid \text{event_abort } e \text{ occurs in } R\}$ and $\text{EvUsed}_R = \{e \text{ that occur in } R\}$ (both `event_abort` e and `[uniquee]`). By Lemma 20, Property 4, we infer

$$\mathcal{D}_{\neg \sigma \text{EvUsed}_R}, \mathcal{D}_{\text{SNU}} : C[L], D_{\text{false}}, \text{EvUsed} \xrightarrow{V}_{p'} C[\sigma R], D_R \circ \sigma^{-1}, \text{EvUsed} \cup \sigma \text{EvUsed}_R$$

where σ is a renaming of the events in EvUsed_R to events not in EvUsed , C is a context acceptable for L and R such that the events that occur in C are in EvUsed , \mathcal{D}_{SNU} is a set of Shoup and non-unique events in EvUsed , $p'(\mathcal{A}, C', t_{D_0}) = p(\mathcal{A}, \sigma^{-1}(C'[C[]]), t_{D_0})$, and $V \subseteq \text{var}(C)$.

Distinguishers in $\mathcal{D}_{\neg \sigma \text{EvUsed}_R}$ include distinguishers that use events in EvUsed , in particular distinguishers for correspondences in G , as well as distinguishers in $\mathcal{D}_{\neg (\text{EvUsed} \cup \sigma \text{EvUsed}_R)}$ used for secrecy and indistinguishability.

- transformations that prove the absence of some events (up to some probability). For such transformations, we have $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, e, \text{EvUsed} \xrightarrow{V}_p G, D_{\text{false}}, \text{EvUsed}$ where p is an upper bound of the probability of event e in G . For Shoup events, this generally happens when G does not contain e and $p(\mathcal{A}, C, t_{D_0}) = 0$. For non-unique events, p is the probability that the `find[uniquee]` or `get[uniquee]` yields several possible choices; after this step, the event e is in `NonUniqueG,D`, so we can exploit uniqueness of `find[uniquee]` or `get[uniquee]`.

That is why, in general, when `CryptoVerif` transforms a game G into a game G' , we have $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D', \text{EvUsed}'$.

There are still transformations that do not fit in this framework (`guess` and `guess_branch`, because they multiply probabilities, as shown in Sections 5.1.17, 5.1.18, and 5.1.19; `success_simplify` because it needs to compensate probabilities of traces that execute \mathbb{S} with those that execute $\bar{\mathbb{S}}$ to show soundness for secrecy, as shown in Section 5.1.23).

2.8.1 Secrecy

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

Definition 7 ((One-session) secrecy) Let Q be an oracle definition, x a variable, and V a set of variables. Let

$$\begin{aligned} Q_{1\text{-ses.secr.}(x)} &= O_{s0}() := b \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); \\ & (O_s(u_1 : [1, n_1], \dots, u_m : [1, n_m]) := \text{if defined}(x[u_1, \dots, u_m]) \text{ then} \\ & \quad \text{if } b \text{ then return } (x[u_1, \dots, u_m]) \text{ else } y \stackrel{R}{\leftarrow} T; \text{return } (y) \end{aligned}$$

$$\begin{aligned}
 & | O'_s(b' : \text{bool}) := \text{if } b = b' \text{ then event_abort } S \text{ else event_abort } \bar{S} \\
 Q_{\text{Secrecy}(x)} = & O_{s0}() := b \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); \\
 & (\text{foreach } i_s \leq n_s \text{ do } O_s[i_s](u_1 : [1, n_1], \dots, u_m : [1, n_m]) := \\
 & \quad \text{if defined}(x[u_1, \dots, u_m]) \text{ then} \\
 & \quad \text{if } b \text{ then return } (x[u_1, \dots, u_m]) \text{ else} \\
 & \quad \text{find } u'_s = i'_s \leq n_s \text{ suchthat defined}(y[i'_s], u_1[i'_s], \dots, u_m[i'_s]) \wedge \\
 & \quad \quad u_1[i'_s] = u_1 \wedge \dots \wedge u_m[i'_s] = u_m \\
 & \quad \text{then return } (y[u'_s]) \\
 & \quad \text{else } y \stackrel{R}{\leftarrow} T; \text{return } (y) \\
 & | O'_s(b' : \text{bool}) := \text{if } b = b' \text{ then event_abort } S \text{ else event_abort } \bar{S}
 \end{aligned}$$

where $O_{s0}, O_s, O'_s \notin \text{fo}(Q)$, $u_1, \dots, u_m, u'_s, y, b, b' \notin \text{var}(Q) \cup V$, S, \bar{S} do not occur in Q , and $\mathcal{E}(x) = [1, n_1] \times \dots \times [1, n_m] \rightarrow T$.

Let sp be $1\text{-ses.secr.}(x)$ (*one-session secrecy of x*) or $\text{Secrecy}(x)$ (*secrecy of x*). The events used by sp are S and \bar{S} . Let $C_{sp} = [] \mid Q_{sp}$.

Let C be an evaluation context acceptable for $C_{sp}[Q]$ with public variables V ($x \notin V$) that does not contain the events used by sp , and $\mathcal{A} \in \mathcal{BB}$ be an attacker. The advantage of the adversary \mathcal{A}, C against sp in process Q is

$$\text{Adv}_Q^{sp}(\mathcal{A}, C) = \Pr^{\mathcal{A}}[C[C_{sp}[Q]] : S] - \Pr^{\mathcal{A}}[C[C_{sp}[Q]] : \bar{S}]$$

The process Q *satisfies* sp with public variables V ($x \notin V$) up to probability p when, for all attackers $\mathcal{A} \in \mathcal{BB}$, for all evaluation contexts C acceptable for $C_{sp}[Q]$ with public variables V that do not contain the events used by sp , $\text{Adv}_Q^{sp}(\mathcal{A}, C) \leq p(\mathcal{A}, C)$.

Intuitively, when Q satisfies sp , the adversary cannot guess the random bit b , that is, it cannot distinguish whether the test oracle Q_{sp} returns the value of the secret ($b = \text{true}$) or a random number ($b = \text{false}$).

For one-session secrecy, the adversary performs a single test query, modeled by $Q_{1\text{-ses.secr.}(x)}$. In more detail, in $Q_{1\text{-ses.secr.}(x)}$, we choose a random bit b ; the adversary queries oracle O_s with the indices (u_1, \dots, u_m) to perform a test query on $x[u_1, \dots, u_m]$: if $b = \text{true}$, the test query returns $x[u_1, \dots, u_m]$; if $b = \text{false}$, it returns a random value y . Finally, the adversary should guess the bit b : it calls oracle O'_s with its guess b' 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 calls oracle O'_s , we have $\Pr^{\mathcal{A}}[C[C_{sp}[Q]] : \bar{S}] = 1 - \Pr^{\mathcal{A}}[C[C_{sp}[Q]] : S]$, so the advantage of the adversary is $\text{Adv}_Q^{sp}(\mathcal{A}, C) = \Pr^{\mathcal{A}}[C[C_{sp}[Q]] : S] - \Pr^{\mathcal{A}}[C[C_{sp}[Q]] : \bar{S}] = 2\Pr^{\mathcal{A}}[C[C_{sp}[Q]] : S] - 1$, 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 $\Pr^{\mathcal{A}}[C[C_{sp}[Q]] : S] - \Pr^{\mathcal{A}}[C[C_{sp}[Q]] : \bar{S}]$ because, when it is negative, we can obtain the opposite, positive value by considering an adversary that calls O'_s with the guess $\neg b'$ instead of b' .)

For secrecy, the adversary can perform several test queries, modeled by $Q_{\text{Secrecy}(x)}$. This corresponds to the “real-or-random” definition of security [4]. (As shown in [4], 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 $x[u_1, \dots, u_m]$.) The replication bound n_s in $Q_{\text{Secrecy}(x)}$ is chosen large enough so that it does not prevent oracle calls that would otherwise

occur, so n_s does not actually limit the number of test queries. When we return a random value ($b = \text{false}$) and several tests queries are performed on the same indices u_1, \dots, u_m , we must return the same random value. That is why, in this case, we look for previous test queries (`find $u'_s \dots$`) and return the previous value of y in case a previous test query was performed with the same indices. For different indices u_1, \dots, u_m , the returned random values are independent of each other, so the secrecy of x requires that the cells of array x are indistinguishable from independent random values. In contrast, the one-session secrecy of x only requires that all array cells of x are indistinguishable from random values, not that they are independent of each other.

By Invariant 3, the variables defined in conditions of `find` and in patterns and in conditions of `get` have no array accesses. Therefore, the definition above applies only to variables x that are not defined in conditions of `find` nor in patterns nor in conditions of `get`.

Lemma 21 *Let sp be $1\text{-ses.secr.}(x)$ or $\text{Secrecy}(x)$.*

If Q satisfies sp with public variables V up to probability p and C is an acceptable evaluation context for Q with public variables V , then for all $V' \subseteq V \cup \text{var}(C)$, $C[Q]$ satisfies sp with public variables V' up to probability p' such that $p'(\mathcal{A}, C') = p(\mathcal{A}, C'[C])$.

If $Q \approx_p^{V \cup \{x\}, \mathcal{E}} Q'$ and Q satisfies sp with public variables V up to probability p' , then Q' satisfies sp with public variables V up to probability p'' such that $p''(\mathcal{A}, C) = p'(\mathcal{A}, C) + 2 \times p(\mathcal{A}, C[C_{sp}[\]], t_S)$.

Proof Suppose that Q satisfies sp with public variables V ($x \notin V$), $\mathcal{A} \in \mathcal{BB}$ is an attacker and C is an acceptable evaluation context for Q with public variables V . Let $V' \subseteq V \cup \text{var}(C)$. Choose oracles O_{s_0}, O_s, O'_s , variables $u_1, \dots, u_m, u'_s, y, b, b'$, and events $\mathbf{S}, \bar{\mathbf{S}}$ such that they do not occur in $C[Q]$. Let C' be an acceptable evaluation context for $C_{sp}[C[Q]]$ with public variables V' that does not contain \mathbf{S} nor $\bar{\mathbf{S}}$. Then we have

$$\begin{aligned} \text{Adv}_{C[Q]}^{sp}(\mathcal{A}, C') &= \Pr^{\mathcal{A}}[C'[C_{sp}[C[Q]]] : \mathbf{S}] - \Pr^{\mathcal{A}}[C'[C_{sp}[C[Q]]] : \bar{\mathbf{S}}] \\ &= \Pr^{\mathcal{A}}[C'[C[C_{sp}[Q]]] : \mathbf{S}] - \Pr^{\mathcal{A}}[C'[C[C_{sp}[Q]]] : \bar{\mathbf{S}}] \\ &\leq p(\mathcal{A}, C'[C]) \end{aligned}$$

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

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

Suppose that $Q \approx_p^{V \cup \{x\}, \mathcal{E}} Q'$ and Q satisfies sp with public variables V up to probability p' . Let $\mathcal{A} \in \mathcal{BB}$ be an attacker. Let C be an acceptable evaluation context for $C_{sp}[Q']$ with public

variables V that does not contain S nor \bar{S} .

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

since $t_S = t_{\bar{S}}$. Indeed, by renaming the variables and tables of C that do not appear in Q' to variables and tables that also do not occur in Q , C is also an acceptable evaluation context for $C_{sp}[Q]$ with public variables V . Furthermore, by Property 6, this renaming does not change the probabilities. \square

2.8.2 Secrecy for a Bit

Definition 8 (Bit secrecy) Let Q be a process, x a boolean variable defined under no replication, and V a set of variables. Let

$$Q_{\text{bit secr.}(x)} = O_s''(b' : \text{bool}) := \text{if defined}(x) \text{ then if } x = b' \text{ then event_abort } S \text{ else event_abort } \bar{S}$$

where $O_s'' \notin \text{fo}(Q)$, $b' \notin \text{var}(Q) \cup V$, S, \bar{S} do not occur in Q , and $\mathcal{E}(x) = \text{bool}$.

Let sp be $\text{bit secr.}(x)$ (*bit secrecy of x*). The events used by sp are S and \bar{S} . Let $C_{sp} = [] \mid Q_{sp}$. The definitions of $\text{Adv}_Q^{sp}(\mathcal{A}, C)$ and “ Q satisfies sp ” are as in Definition 7.

Intuitively, when Q satisfies sp , the adversary cannot guess the boolean x , that is, it cannot distinguish whether $x = \text{true}$ or $x = \text{false}$. The adversary performs a single test query, modeled by $Q_{\text{bit secr.}(x)}$. This definition is simpler than the definition of (one-session) secrecy for x , because we do not introduce an additional random bit b .

By Invariant 3, the variables defined in conditions of `find` and in patterns and in conditions of `get` have no array accesses. Therefore, the definition above applies only to variables x that are not defined in conditions of `find` nor in patterns nor in conditions of `get`.

Lemma 21 is also valid when $sp = \text{bit secr.}(x)$, with the same statement and proof.

Lemma 22 *If b_0 is a boolean variable defined under no replication and Q preserves the one-session secrecy of b_0 with public variables V up to probability p , then Q preserves the bit secrecy of b_0 with public variables V up to probability $2p$.*

Proof Let $\mathcal{A} \in \mathcal{BB}$ be an attacker. Let C be any acceptable evaluation context for $Q \mid Q_{\text{bit secr.}(b_0)}$ with public variables V . Let

$$C' = C[- \mid O_s''(b'_0 : \text{bool}) := \text{let } () = O_{s_0}() \text{ in let } b''_0 : \text{bool} = O_s() \text{ in let } () = O'_s(b'_0 = b''_0) \text{ in yield}].$$

We execute $C'[Q \mid Q_{1\text{-ses.secr.}(b_0)}]$ with attacker \mathcal{A} .

When $b'_0 = b_0$ ($C[Q \mid Q_{\text{bit secr.}(b_0)}]$ executes S),

- if $b = \text{true}$ (probability $1/2$), then $b''_0 = b_0$, so $b' = (b'_0 = b''_0) = \text{true}$ and S is executed with probability $1/2$
- if $b = \text{false}$ (probability $1/2$), then b''_0 is random, so
 - $b''_0 = b'_0$ with probability $1/4$, so $b' = (b'_0 = b''_0) = \text{true}$ and \bar{S} is executed;

– $b'_0 = \neg b'_0$ with probability $1/4$, so $b' = (b'_0 = b''_0) = \text{false}$ and \mathbf{S} is executed.

When $b'_0 = \neg b_0$ ($C[Q \mid Q_{\text{bit secr.}(b_0)}]$ executes $\bar{\mathbf{S}}$),

- if $b = \text{true}$ (probability $1/2$), then $b'_0 = b_0$, so b' is false and $\bar{\mathbf{S}}$ is executed with probability $1/2$;
- if $b = \text{false}$ (probability $1/2$), then b'_0 is random, so
 - $b'_0 = b'_0$ with probability $1/4$, so $b' = \text{true}$ and $\bar{\mathbf{S}}$ is executed;
 - $b'_0 = \neg b'_0$ with probability $1/4$, so $b' = \text{false}$ and \mathbf{S} is executed.

So

$$\begin{aligned} \Pr^{\mathcal{A}}[C'[Q \mid Q_{1\text{-ses.secr.}(b_0)}] : \mathbf{S}] &= \frac{3}{4} \Pr^{\mathcal{A}}[C[Q \mid Q_{\text{bit secr.}(b_0)}] : \mathbf{S}] + \frac{1}{4} \Pr^{\mathcal{A}}[C[Q \mid Q_{\text{bit secr.}(b_0)}] : \bar{\mathbf{S}}] \\ \Pr^{\mathcal{A}}[C'[Q \mid Q_{1\text{-ses.secr.}(b_0)}] : \bar{\mathbf{S}}] &= \frac{1}{4} \Pr^{\mathcal{A}}[C[Q \mid Q_{\text{bit secr.}(b_0)}] : \mathbf{S}] + \frac{3}{4} \Pr^{\mathcal{A}}[C[Q \mid Q_{\text{bit secr.}(b_0)}] : \bar{\mathbf{S}}] \end{aligned}$$

Finally, we obtain

$$\begin{aligned} &\Pr^{\mathcal{A}}[C'[Q \mid Q_{1\text{-ses.secr.}(b_0)}] : \mathbf{S}] - \Pr^{\mathcal{A}}[C'[Q \mid Q_{1\text{-ses.secr.}(b_0)}] : \bar{\mathbf{S}}] \\ &= \frac{1}{2} (\Pr^{\mathcal{A}}[C[Q \mid Q_{\text{bit secr.}(b_0)}] : \mathbf{S}] - \Pr^{\mathcal{A}}[C[Q \mid Q_{\text{bit secr.}(b_0)}] : \bar{\mathbf{S}}]) \end{aligned}$$

so $\Pr^{\mathcal{A}}[C[Q \mid Q_{\text{bit secr.}(b_0)}] : \mathbf{S}] - \Pr^{\mathcal{A}}[C[Q \mid Q_{\text{bit secr.}(b_0)}] : \bar{\mathbf{S}}] = 2(\Pr^{\mathcal{A}}[C'[Q \mid Q_{1\text{-ses.secr.}(b_0)}] : \mathbf{S}] - \Pr^{\mathcal{A}}[C'[Q \mid Q_{1\text{-ses.secr.}(b_0)}] : \bar{\mathbf{S}}]) \leq 2p(\mathcal{A}, C') = 2p(\mathcal{A}, C)$, neglecting the additional runtime of C' . \square

Intuitively, the factor 2 is necessary, because in the definition of one-session secrecy, even if the adversary knows the secret bit b_0 perfectly, it will not be able to distinguish b_0 from the random bit y in half of the cases, because b_0 and y have the same value.

In the rest of Section 2.8.2, we consider a process

$$Q = O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q'$$

and let $Q_L = Q'\{\text{true}/b_0\}$ and $Q_R = Q'\{\text{false}/b_0\}$, so that Q chooses a random bit b_0 and runs as Q_L when b_0 is true and as Q_R when b_0 is false. We assume that Q_L and Q_R never abort, that is, they contain neither `event_abort` nor `find[uniquee]`. Moreover, they do not use the variable b_0 . We have the following lemmas.

Lemma 23 *We assume that attackers in \mathcal{BB} can at least represent probabilistic Turing machines. If Q preserves the bit secrecy of b_0 with public variables V up to probability p , then $Q_L \approx_p^{V,0} Q_R$ where $p'(\mathcal{A}, C, t_D) = p(\mathcal{A} + C + t_D, C_V)$ for some attacker $\mathcal{A} + C + t_D$ and context C_V such that the attacker $\mathcal{A} + C + t_D$ with context C_V runs in time $t_{\mathcal{A}} + t_C + t_D$ and its other parameters (replication bounds, lengths of bitstrings) are the same as for \mathcal{A} plus C .*

The processes Q_L and Q_R can execute different events without breaking the bit secrecy of b_0 , because the adversary for the bit secrecy of b_0 does not have access to the events executed by Q . Hence, $Q_L \approx_p^V Q_R$ would not hold in general.

Proof Let $\mathcal{A} \in \mathcal{BB}$ be an attacker. Let C be any acceptable evaluation context for Q_L and Q_R with public variables V , and D a distinguisher such that $\text{event}(D) \cap \text{event}(Q_L, Q_R) = \emptyset$. Let O and O'_s be oracles that C does not use.

Let \mathcal{A}' be an attacker that calls oracle O , runs \mathcal{A} and C but stores events executed by C in its internal state instead of actually executing events, computes D on the stored sequence of events executed by C , stores the result in b'_0 , and calls $O'_s(b'_0)$. Such an attacker \mathcal{A}' exists because attackers in \mathcal{BB} can at least represent probabilistic Turing machines. However, this attacker needs to be able to access the public variables in V . To this end, given a variable $x \in V$ with type $x : [1, n_1] \times \dots \times [1, n_m] \rightarrow T$, we define the oracle definition $Q_x = \text{foreach } i \leq n \text{ do } O_x(u_1 : [1, n_1], \dots, u_m : [1, n_m]) := \text{if defined}(x[u_1, \dots, u_m]) \text{ then return } (x[u_1, \dots, u_m])$, with n large enough so that it does not limit the number of queries that the attacker makes. Given a set of public variables $V = \{v_1, \dots, v_m\}$, we then define the context $C_V = Q_{v_1} \mid \dots \mid Q_{v_m} \mid []$. Whenever \mathcal{A}' executes a part of C that depends on some public variable, it uses the corresponding oracle to obtain the values inside the array. With such a construction, and without going into any formal details, the runtime of $t_{\mathcal{A}} + t_C + t_D$ is roughly equal to the runtime of $t_{\mathcal{A}'} + t_{C_V}$. To formalize this statement, some care would need to be taken on how \mathcal{A}' can simulate the find constructs of C by going through the full array by making many queries to oracles in C_V .

When b_0 is true, $Q \mid Q_{\text{bit secr.}(x)}$ run with attacker \mathcal{A}' and context C_V stores in b'_0 the result of $C[Q_L] : D$ run with attacker \mathcal{A} . When $b'_0 = \text{true}$, S is executed. When $b'_0 = \text{false}$, \bar{S} is executed. So

$$\begin{aligned} \Pr^{\mathcal{A}'}[C_V[Q \mid Q_{\text{bit secr.}(x)}] : S/b_0 = \text{true}] &= \Pr^{\mathcal{A}'}[C[Q_L] : D] \\ \Pr^{\mathcal{A}'}[C_V[Q \mid Q_{\text{bit secr.}(x)}] : \bar{S}/b_0 = \text{true}] &= 1 - \Pr^{\mathcal{A}'}[C[Q_L] : D] \end{aligned}$$

When b_0 is false, $Q \mid Q_{\text{bit secr.}(x)}$ run with attacker \mathcal{A}' and context C_V stores in b'_0 the result of $C[Q_R] : D$ run with attacker \mathcal{A} . When $b'_0 = \text{true}$, \bar{S} is executed. When $b'_0 = \text{false}$, S is executed. So

$$\begin{aligned} \Pr^{\mathcal{A}'}[C_V[Q \mid Q_{\text{bit secr.}(x)}] : S/b_0 = \text{false}] &= 1 - \Pr^{\mathcal{A}'}[C[Q_R] : D] \\ \Pr^{\mathcal{A}'}[C_V[Q \mid Q_{\text{bit secr.}(x)}] : \bar{S}/b_0 = \text{false}] &= \Pr^{\mathcal{A}'}[C[Q_R] : D] \end{aligned}$$

Finally, we obtain

$$\begin{aligned} &\Pr^{\mathcal{A}'}[C_V[Q \mid Q_{\text{bit secr.}(x)}] : S] - \Pr^{\mathcal{A}'}[C_V[Q \mid Q_{\text{bit secr.}(x)}] : \bar{S}] \\ &= \frac{1}{2}(\Pr[C[Q_L] : D] + 1 - \Pr^{\mathcal{A}'}[C[Q_R] : D] - (1 - \Pr^{\mathcal{A}'}[C[Q_L] : D]) - \Pr^{\mathcal{A}'}[C[Q_R] : D]) \\ &= \Pr^{\mathcal{A}'}[C[Q_L] : D] - \Pr^{\mathcal{A}'}[C[Q_R] : D] \end{aligned}$$

so $\Pr^{\mathcal{A}'}[C[Q_L] : D] - \Pr^{\mathcal{A}'}[C[Q_R] : D] \leq p'(\mathcal{A}, C, t_D)$. By negating the bit b'_0 , we swap the events S and \bar{S} without changing the probability, so $\Pr^{\mathcal{A}'}[C[Q_R] : D] - \Pr^{\mathcal{A}'}[C[Q_L] : D] \leq p'(\mathcal{A}, C, t_D)$. Therefore, $|\Pr^{\mathcal{A}'}[C[Q_L] : D] - \Pr^{\mathcal{A}'}[C[Q_R] : D]| \leq p'(\mathcal{A}, C, t_D)$. So $Q_L \approx_{p'}^{V, \emptyset} Q_R$. \square

Lemma 23 is the main motivation for the notion of secrecy for a bit: it allows proving indistinguishability between two processes by showing secrecy of bit b_0 . Using this notion instead of one-session secrecy of b_0 avoids losing a factor 2, as shown by Lemma 22.

We use this idea to encode the diff construct originally introduced in ProVerif [30]: given a process Q_1 that contains terms $\text{diff}[M, M']$ and processes $\text{diff}[P, P']$, we define $\text{fst}(Q_1)$ as Q_1 with $\text{diff}[M, M']$ replaced with M and $\text{snd}(Q_1)$ as Q_1 with $\text{diff}[M, M']$ replaced with M' , and similarly for $\text{diff}[P, P']$; the goal is to show that $\text{fst}(Q_1) \approx_p^{V, \emptyset} \text{snd}(Q_1)$ for some p and determine

p . In order to do that, we define Q' as Q_1 with $\text{diff}[M, M']$ replaced with $\text{if_fun}(b_0, M, M')$ when M and M' are simple and with $\text{if } b_0 \text{ then } M \text{ else } M'$ otherwise¹, $\text{diff}[P, P']$ replaced with $\text{if } b_0 \text{ then } P \text{ else } P'$, and $Q = O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q'$. We have $Q_L = Q' \{\text{true}/b_0\} \approx_0^{V, \emptyset} \text{fst}(Q_1)$ and $Q_R = Q' \{\text{false}/b_0\} \approx_0^{V, \emptyset} \text{snd}(Q_1)$, so by Lemma 23, if Q preserves the bit secrecy of b_0 with public variables V up to probability p , then $\text{fst}(Q_1) \approx_{p'}^{V, \emptyset} \text{snd}(Q_1)$ where $p'(\mathcal{A}, C, t_D) = p(\mathcal{A} + C + t_D, C_V)$.

Lemma 24 *If $Q_L \approx_p^{V, \emptyset} Q_R$ then $Q \approx_{p/2}^{V \cup \{b_0\}, \emptyset} O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R$.*

Proof Let $\mathcal{A} \in \mathcal{BB}$ be an attacker. Let C be an evaluation context acceptable for Q and $O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R$ with public variables $V \cup \{b_0\}$ and D be a distinguisher. We have

$$\begin{aligned} & |\Pr^{\mathcal{A}}[C[Q] : D] - \Pr^{\mathcal{A}}[C[O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R] : D]| \\ & \leq \frac{1}{2} |\Pr^{\mathcal{A}}[C[Q] : D/b_0 = \text{true}] - \Pr^{\mathcal{A}}[C[O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R] : D/b_0 = \text{true}]| \\ & \quad + \frac{1}{2} |\Pr^{\mathcal{A}}[C[Q] : D/b_0 = \text{false}] - \Pr^{\mathcal{A}}[C[O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R] : D/b_0 = \text{false}]| \\ & \leq \frac{1}{2} |\Pr^{\mathcal{A}}[C[O() := \text{let } b_0 = \text{true in return } (); Q_L] : D] \\ & \quad - \Pr^{\mathcal{A}}[C[O() := \text{let } b_0 = \text{true in return } (); Q_R] : D]| \\ & \leq \frac{1}{2} p(\mathcal{A}, C, t_D) \end{aligned}$$

Indeed, $\Pr^{\mathcal{A}}[C[O() := \text{let } b_0 = \text{true in return } (); Q_L] : D] = \Pr^{\mathcal{A}}[C''[Q_L] : D]$ (and similarly for Q_R), where $C'' = C[\text{newOracle } \tilde{O}'; (F_{\tilde{O}, \tilde{O}'} \mid \text{newOracle } \tilde{O}; (O() := \text{let } b_0 = \text{true in return } (); F_{\tilde{O}', \tilde{O}}) \mid [])]$, \tilde{O} are the oracles at the root of Q_L and Q_R (they are not used elsewhere by Invariant 7), \tilde{O}' are fresh oracles corresponding to oracles in \tilde{O} , and $F_{\tilde{O}', \tilde{O}}$ forwards all oracle calls on an oracle in \tilde{O}' to the corresponding oracle in \tilde{O} , with the same replication bounds as in Q_L and Q_R . (All calls to oracles in \tilde{O} are forwarded to oracles in \tilde{O}' and then back to oracles in \tilde{O} provided the code $O() := \text{let } b_0 = \text{true in return } ()$ has already been executed. That prevents executing Q_L or Q_R before $O() := \text{let } b_0 = \text{true in return } ()$.) By Property 6, replacing C'' with C as argument of $p(C, t_D)$ does not affect the probability. (We neglect the additional runtime of C'' .) \square

Lemma 25 is the converse of Lemma 23.

Lemma 25 *If $Q_L \approx_p^{V, \emptyset} Q_R$, then Q preserves the bit secrecy of b_0 with public variables V up to probability p' where $p'(\mathcal{A}, C) = p(\mathcal{A}, C[C_{\text{bit secr.}(b_0)}], t_S)$.*

Proof If $Q_L \approx_p^{V, \emptyset} Q_R$, then by Lemma 24, $Q \approx_{p/2}^{V \cup \{b_0\}, \emptyset} O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R$. Moreover, $O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R$ preserves the bit secrecy of b_0 with public variables V up to probability 0. (Since Q_R does not use b_0 , the variable b' in $Q_{\text{bit secr.}(b_0)}$ is independent of b_0 , so a trace that executes S corresponds to a trace of the same probability and that executes \bar{S} by changing the value of b_0 , so $\Pr^{\mathcal{A}}[C[O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R \mid Q_{\text{bit secr.}(b_0)}] : S] =$

¹ $\text{if_fun}(b_0, M, M')$ differs from $\text{if } b_0 \text{ then } M \text{ else } M'$ in that it evaluates both M and M' . Since $\text{diff}[M, M']$ evaluates either M or M' but not both, we translate it into $\text{if } b_0 \text{ then } M \text{ else } M'$ when the evaluation of M or M' may modify the semantic state, e.g. by executing an event or by defining a variable. The evaluation of simple terms does not modify the semantic state.

$\Pr^{\mathcal{A}}[C[O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R \mid Q_{\text{bit secr.}(b_0)}] : \bar{S}]$.) So by Lemma 21 (version for bit secrecy), Q preserves the bit secrecy of b_0 with public variables V up to probability p' . \square

Lemma 26 provides a converse of Lemma 22 when Q has the particular form given above. There is no probability loss in this case.

Lemma 26 *We assume that attackers in \mathcal{BB} can at least represent probabilistic Turing machines. If Q preserves the bit secrecy of b_0 with public variables V up to probability p , then Q preserves the one-session secrecy of b_0 with public variables V up to probability p .*

Proof If Q preserves the bit secrecy of b_0 with public variables V up to probability p , then by Lemma 23, $Q_L \approx_{p'}^{V, \emptyset} Q_R$ where $p'(\mathcal{A}, C, t_D) = p(\mathcal{A} + C + t_D, C_V)$. By Lemma 24, $Q \approx_{p'/2}^{V \cup \{b_0\}, \emptyset} O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R$.

Moreover, $O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R$ preserves the one-session secrecy of b_0 with public variables V up to probability 0. Indeed, since Q_R does not use b_0 , b_0 can in fact be chosen in the test query in $Q_{1\text{-ses.secr.}(b_0)}$, so that test query always returns a random boolean, independently of the value of the variable b of $Q_{1\text{-ses.secr.}(b_0)}$. Therefore, the variable b' is independent of b , so a trace that executes S corresponds to a trace of the same probability and that executes \bar{S} by changing the value of b , so $\Pr^{\mathcal{A}}[C[O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R \mid Q_{1\text{-ses.secr.}(b_0)}] : S] = \Pr^{\mathcal{A}}[C[O() := b_0 \stackrel{R}{\leftarrow} \text{bool}; \text{return } (); Q_R \mid Q_{1\text{-ses.secr.}(b_0)}] : \bar{S}]$.

So by Lemma 21 (version for one-session secrecy), Q preserves the one-session secrecy of b_0 with public variables V up to probability p'' such that $p''(\mathcal{A}, C) = p'(\mathcal{A}, C[C_{1\text{-ses.secr.}(b_0)}], t_S) = p(\mathcal{A} + C[C_{1\text{-ses.secr.}(b_0)}] + t_S, C_V)$ which is about $p(\mathcal{A}, C)$ by Property 6, neglecting the additional runtime of the context. \square

2.8.3 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 formulas $\psi \Rightarrow \phi$, which may contain events. We use the following logical formulas:

| $\phi ::=$ | 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 formulas must contain only variables x without array indices and function applications, and their variables are assumed to be distinct from variables of processes. Formulas denoted by ψ are conjunctions of events. In a correspondence $\psi \Rightarrow \phi$, the variables of ψ are universally quantified; those of ϕ that do not occur in ψ are existentially quantified. Formally:

Definition 9 The semantics of the correspondence $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$, also written $\tilde{x} : \tilde{T}, \tilde{y} : \tilde{T}'; \psi \Rightarrow \phi$ in a less explicit syntax, is $\llbracket \forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket = \llbracket \tilde{x} : \tilde{T}, \tilde{y} : \tilde{T}'; \psi \Rightarrow \phi \rrbracket = \forall \tilde{x} \in \tilde{T}, (\psi \Rightarrow \exists \tilde{y} \in \tilde{T}', \phi)$, where $\tilde{x} = \text{var}(\psi)$ and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$.

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. Conjunction, disjunction, implication, existential and universal quantifications are defined as usual. More formally, we write $\rho, \mathcal{E}v \vdash \varphi$ when the sequence of events $\mathcal{E}v$ satisfies the formula φ , in the environment ρ that maps variables to their values. We define $\rho, \mathcal{E}v \vdash \varphi$ as follows:

$\rho, \mathcal{E}v \vdash M$ if and only if $\rho, M \Downarrow \text{true}$
 $\rho, \mathcal{E}v \vdash \text{event}(e(M_1, \dots, M_m))$ if and only if
 for all $j \leq m$, $\rho, M_j \Downarrow a_j$ and $e(a_1, \dots, a_m) \in \mathcal{E}v$
 $\rho, \mathcal{E}v \vdash \varphi_1 \wedge \varphi_2$ if and only if $\rho, \mathcal{E}v \vdash \varphi_1$ and $\rho, \mathcal{E}v \vdash \varphi_2$
 $\rho, \mathcal{E}v \vdash \varphi_1 \vee \varphi_2$ if and only if $\rho, \mathcal{E}v \vdash \varphi_1$ or $\rho, \mathcal{E}v \vdash \varphi_2$
 $\rho, \mathcal{E}v \vdash \varphi_1 \Rightarrow \varphi_2$ if and only if $\rho, \mathcal{E}v \vdash \varphi_1$ implies $\rho, \mathcal{E}v \vdash \varphi_2$
 $\rho, \mathcal{E}v \vdash \exists x \in T, \varphi$ if and only if there exists $a \in T$ such that $\rho[x \mapsto a], \mathcal{E}v \vdash \varphi$
 $\rho, \mathcal{E}v \vdash \forall x \in T, \varphi$ if and only if for every $a \in T$, we have $\rho[x \mapsto a], \mathcal{E}v \vdash \varphi$

When φ is a closed formula, we write $\mathcal{E}v \vdash \varphi$ for $\rho, \mathcal{E}v \vdash \varphi$ where ρ is the empty function.

Definition 10 The sequence of events $\mathcal{E}v$ satisfies the correspondence φ if and only if $\mathcal{E}v \vdash \varphi$.

Definition 11 We define a distinguisher $D(\mathcal{E}v) = \text{true}$ if and only if $\mathcal{E}v \vdash \varphi$, and we denote this distinguisher D simply by φ .

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

The process Q satisfies the correspondence φ 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 φ and all $\mathcal{A} \in \mathcal{BB}$, $\text{Adv}_Q^\varphi(\mathcal{A}, C) \leq p(\mathcal{A}, C)$.

When sp is a correspondence φ , we define $C_{sp} = []$ and the events used by sp are the events that occur in the formula φ . Therefore, the definition of “ Q satisfies the correspondence φ ” matches the definition of “ Q satisfies sp ” given in Definition 7.

A process satisfies φ up to probability p when the probability that it generates a sequence of events $\mathcal{E}v$ that does not satisfy φ is at most $p(\mathcal{A}, C)$, in the presence of an adversary represented by the attacker \mathcal{A} and the context C .

Example 2 The semantics of the correspondence

$$\forall x : pkey, y : host, z : nonce; \text{event}(e_B(x, y, z)) \Rightarrow \text{event}(e_A(x, y, z)) \quad (4)$$

is

$$\forall x \in pkey, \forall y \in host, \forall z \in nonce, \text{event}(e_B(x, y, z)) \Rightarrow \text{event}(e_A(x, y, z)) \quad (5)$$

It 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 semantics of the correspondence

$$\begin{aligned} \forall x : T; \text{event}(e_1(x)) \wedge \text{event}(e_2(x)) \Rightarrow \\ \exists y : T'; \text{event}(e_3(x)) \vee (\text{event}(e_4(x, y)) \wedge \text{event}(e_5(x, y))) \end{aligned}$$

is

$$\begin{aligned} \forall x \in T, \text{event}(e_1(x)) \wedge \text{event}(e_2(x)) \Rightarrow \\ \exists y \in T', \text{event}(e_3(x)) \vee (\text{event}(e_4(x, y)) \wedge \text{event}(e_5(x, y))) \end{aligned}$$

It 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 formulas, we extend the grammar of formulas 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, to be able to indicate at which step an event has been executed (that is, at which index it appears in $\mathcal{E}v$): $\text{event}(e(\widetilde{M}))@_\tau$ means that event $e(\widetilde{M})$ has been executed at step τ . Formally:

$\rho, \mathcal{E}v \vdash \text{event}(e(M_1, \dots, M_m))@M_0$ if and only if
for all $j \leq m$, $\rho, M_j \Downarrow a_j$, $a_0 \neq \perp$, and $e(a_1, \dots, a_m) = \mathcal{E}v(a_0)$

With this definition, we have:

Definition 12 The semantics of the correspondence $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$, also written $\tilde{x} : \tilde{T}, \tilde{y} : \tilde{T}'; \psi \Rightarrow \phi$ in a less explicit syntax, is

$$\begin{aligned} \llbracket \forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket &= \llbracket \tilde{x} : \tilde{T}, \tilde{y} : \tilde{T}'; \psi \Rightarrow \phi \rrbracket = \exists f_1, \dots, f_k \in \mathbb{N}^m \times \prod \tilde{T} \rightarrow \mathbb{N} \cup \{\perp\}, \\ &\text{Inj}(I, f_1) \wedge \dots \wedge \text{Inj}(I, f_k) \wedge \forall \tau_1, \dots, \tau_m \in \mathbb{N}, \forall \tilde{x} \in \tilde{T}, (\psi^\tau \Rightarrow \exists \tilde{y} \in \tilde{T}', \phi^\tau), \end{aligned}$$

where $\tilde{x} = \text{var}(\psi)$, $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$, $\psi = F_1 \wedge \dots \wedge F_m$, $\psi^\tau = F_1^\tau \wedge \dots \wedge F_m^\tau$, $F_j^\tau = \text{event}(e(\widetilde{M}))@_\tau$ if $F_j = \text{event}(e(\widetilde{M}))$ or $F_j = \text{inj-event}(e(\widetilde{M}))$, $I = \{j \mid F_j = \text{inj-event}(\dots)\}$, ϕ^τ is obtained from ϕ by replacing each injective event $\text{inj-event}(e(\widetilde{M}))$ with $\text{event}(e(\widetilde{M}))@_{f_j(\tau_1, \dots, \tau_m, \tilde{x})}$ using a distinct function f_j for each injective event in ϕ^τ , and $\text{Inj}(I, f)$ if and only if $f(\tau_1, \dots, \tau_m, \tilde{x}) = f(\tau'_1, \dots, \tau'_m, \tilde{x}') \neq \perp \Rightarrow \forall j \in I, \tau_j = \tau'_j$.

In ψ^τ and ϕ^τ , events are labeled with their associated execution step, τ_j for the events in ψ^τ and $f_j(\tau_1, \dots, \tau_m, \tilde{x})$ for the injective events in ϕ^τ . Therefore, the functions f_j map the execution steps of events in ψ , τ_1, \dots, τ_m , and the values of the variables in ψ , \tilde{x} , to the associated execution steps of injective events in ϕ . (The result \perp corresponds to the case in which the event in ϕ is not executed: in case of disjunctions, not all events in ϕ are required to be executed.) The correspondence is injective when these functions f_j are injective in their arguments that correspond to injective events in ψ . The indices of injective events in ψ are collected in the set I , and injectivity is guaranteed by $\text{Inj}(I, f_j)$, which means that, ignoring the result \perp , f_j is injective in its arguments of indices in I .

Definition 11 is unchanged for injective correspondences.

Example 3 The semantics of the correspondence

$$\forall x : pkey, y : host, z : nonce; \text{inj-event}(e_B(x, y, z)) \Rightarrow \text{inj-event}(e_A(x, y, z)) \quad (6)$$

is

$$\begin{aligned} \exists f \in \mathbb{N} \times pkey \times host \times nonce \rightarrow \mathbb{N} \cup \{\perp\}, \text{Inj}(\{1\}, f) \wedge \\ \forall \tau \in \mathbb{N}, \forall x \in pkey, \forall y \in host, \forall z \in nonce, \\ \text{event}(e_B(x, y, z))@_\tau \Rightarrow \text{event}(e_A(x, y, z))@_{f(\tau, x, y, z)} \end{aligned} \quad (7)$$

It 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, f is a function that maps the execution step τ of $e_B(x, y, z)$ and the variables x, y, z to the execution step of $e_A(x, y, z)$. (This step is never \perp .) This function is injective in its first argument, the step τ , so if there are n executions of $e_B(x, y, z)$, at steps τ_1, \dots, τ_n , then there are at least n executions of $e_A(x, y, z)$, at steps $f(\tau_1, x, y, z), \dots, f(\tau_n, x, y, z)$ and these steps are distinct by injectivity of f in its first argument.

The semantics of the correspondence

$$\forall x : T; \text{event}(e_1(x)) \wedge \text{inj-event}(e_2(x)) \Rightarrow \exists y : T'; \text{inj-event}(e_3(x)) \vee (\text{inj-event}(e_4(x, y)) \wedge \text{inj-event}(e_5(x, y)))$$

is

$$\begin{aligned} & \exists f_1, f_2, f_3 \in \mathbb{N}^2 \times T \rightarrow \mathbb{N} \cup \{\perp\}, \text{Inj}(\{2\}, f_1) \wedge \text{Inj}(\{2\}, f_2) \wedge \text{Inj}(\{2\}, f_3) \wedge \\ & \forall \tau_1, \tau_2 \in \mathbb{N}, \forall x \in T, \text{event}(e_1(x))@_{\tau_1} \wedge \text{event}(e_2(x))@_{\tau_2} \Rightarrow \exists y \in T', \text{event}(e_3(x))@_{f_1(\tau_1, \tau_2, x)} \vee \\ & (\text{event}(e_4(x, y))@_{f_2(\tau_1, \tau_2, x)} \wedge \text{event}(e_5(x, y))@_{f_3(\tau_1, \tau_2, x)}) \end{aligned}$$

It 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 functions f_1, f_2 , and f_3 map the execution steps τ_1 and τ_2 of e_1 and e_2 and the variable x to the execution steps of e_3, e_4 , and e_5 respectively. Ignoring the result \perp , they are injective in their second argument, which corresponds to the execution step of the injective event e_2 .

When no injective event occurs in $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$, Definition 12 reduces to the definition of non-injective correspondences: there are no functions f_j , $\phi^\tau = \phi$, and $\text{event}(e(\tilde{M}))@_\tau$ holds for some τ if and only if $\text{event}(e(\tilde{M}))$ holds, so ψ^τ holds for some τ_1, \dots, τ_m if and only if ψ holds.

Well-formedness condition When we consider a correspondence $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$, with $\tilde{x} = \text{var}(\psi)$ and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$, we should have

$$\forall \tilde{x} \in \tilde{T}, \forall \tilde{x}' \in \tilde{T}, \forall \tilde{y} \in \tilde{T}', \psi = \psi\{\tilde{x}'/\tilde{x}\} \Rightarrow \phi = \phi\{\tilde{x}'/\tilde{x}\} \quad (8)$$

where \tilde{x}' are fresh variables, and the equality of terms is the equality of their values, but disjunctions, conjunctions, and events are considered syntactically. This condition guarantees that, given an execution of events in ψ , the formula to verify $\exists \tilde{y} \in \tilde{T}'; \phi^\tau$ is uniquely determined. It avoids pathological correspondences such as

$$\forall x : T; \text{event}(e(f(x))) \Rightarrow \text{event}(e'(x)) \quad (9)$$

with $f(a) = f(b) = c$, for which $\text{event}(e(c))$ corresponds to both $\text{event}(e(f(a)))$ and $\text{event}(e(f(b)))$, so when event $e(c)$ is executed, x can take both values a and b , so (9) requires the execution of events $e'(a)$ and $e'(b)$. An even more pathological case is when $f(x) = c$ for all x : in this case, when event $e(c)$ is executed, (9) requires the execution of event $e'(x)$ for all $x \in T$, which is impossible when T is infinite. However, the condition allows the correspondence (9) when f is injective, so x is uniquely determined, and it also allows the correspondences

$$\forall x : T; \text{event}(e(f(x))) \Rightarrow \text{false} \quad (10)$$

and

$$\forall x : T; \text{event}(e(f(x))) \Rightarrow \text{event}(e'(f(x))) \quad (11)$$

for any function f : (10) requires that event $e(y)$ is never executed with y in the image of f , and (11) requires that $e'(y)$ is executed when $e(y)$ is executed with y in the image of f .

CryptoVerif displays a warning when it does not manage to prove the well-formedness condition (8).

Property

Lemma 27 *If Q satisfies a correspondence φ 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 φ , then for all $V' \subseteq V \cup \text{var}(C)$, $C[Q]$ satisfies a correspondence φ with public variables V' up to probability p' such that $p'(\mathcal{A}, C') = p(\mathcal{A}, C'[C])$.*

If $Q \approx_p^{V, \mathcal{E}} Q'$, Q satisfies a correspondence φ with public variables V up to probability p' , and \mathcal{E} contains all events in φ , then Q' satisfies φ with public variables V up to probability p'' such that $p''(\mathcal{A}, C) = p'(\mathcal{A}, C) + p(\mathcal{A}, C, t_\varphi)$.

Proof Suppose that Q satisfies a correspondence φ with public variables V , $\mathcal{A} \in \mathcal{BB}$ is an attacker, and C is an acceptable evaluation context for Q with public variables V that does not contain events used in φ . 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 φ . We rename the variables of C' not in V' so that they are not in V ; by Property 6, this renaming does not change the probabilities. We have

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

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 $Q \approx_p^{V, \mathcal{E}} Q'$, Q satisfies a correspondence φ with public variables V up to probability p' , and \mathcal{E} contains all events in φ . Let $\mathcal{A} \in \mathcal{BB}$ be an attacker. Let C be an evaluation context acceptable for Q' with public variables V that does not contain events used by φ . We have

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

Indeed, by renaming the variables and tables of C that do not appear in Q' to variables and tables that also do not occur in Q , C is also an acceptable evaluation context for Q with public variables V . Furthermore, by Property 6, this renaming does not change the probabilities. \square

Reachability secrecy Reachability secrecy aims to show that the adversary cannot compute the secret value. This notion is standard in the symbolic model, but less common than the notion of secrecy as “the adversary cannot distinguish the secret from a random value” (Section 2.8.1) in the computational model. It is still used, e.g. in the property of one-wayness or in the computational Diffie-Hellman assumption.

This notion makes sense only when the secret value is of a large type. Otherwise, the adversary would have a non-negligible probability of finding the secret value just by random guessing.

This notion is in fact encoded as a correspondence property. We distinguish two variants. One-session reachability secrecy of x means that the adversary cannot compute any cell of array x , even if it has access to the public variables in V . Reachability secrecy of x means that the adversary cannot compute any cell of array x , even if it has access to the other cells of x and to the public variables in V .

Definition 13 ((One-session) reachability secrecy) Let Q be a process, x a variable, and V a set of variables. Let

$$\begin{aligned}
Q_{1\text{-ses.reach.secr.}(x)} &= \text{foreach } i_t \leq n_t \text{ do } O_s[i_t](x' : T, u_1 : [1, n_1], \dots, u_m : [1, n_m]) := \\
&\quad \text{if defined}(x[u_1, \dots, u_m]) \wedge x' = x[u_1, \dots, u_m] \text{ then event adv_has_x} \\
Q_{\text{Reach.secr.}(x)} &= \text{foreach } i_r \leq n_r \text{ do } O_r[i_r](u'_1 : [1, n_1], \dots, u'_m : [1, n_m]) := \\
&\quad \text{if defined}(x[u'_1, \dots, u'_m]) \text{ then} \\
&\quad \text{let } reveal : \text{bool} = \text{true} \text{ in return } (x[u'_1, \dots, u'_m]) \\
&| \text{foreach } i_t \leq n_t \text{ do } O_s[i_t](x' : T, u_1 : [1, n_1], \dots, u_m : [1, n_m]) := \\
&\quad \text{find } i \leq i_r \text{ suchthat defined}(reveal[i], u'_1[i], \dots, u'_m[i]) \wedge \\
&\quad \quad u'_1[i] = u_1 \wedge \dots \wedge u'_m[i] = u_m \text{ then yield else} \\
&\quad \text{if defined}(x[u_1, \dots, u_m]) \wedge x' = x[u_1, \dots, u_m] \text{ then event adv_has_x}
\end{aligned}$$

where $O_s, O_r \notin \text{fo}(Q)$, $x', u_1, \dots, u_m, u'_1, \dots, u'_m, reveal \notin \text{var}(Q) \cup V$, adv_has_x does not occur in Q , and $\mathcal{E}(x) = [1, n_1] \times \dots \times [1, n_m] \rightarrow T$.

The process Q satisfies one-session reachability secrecy of x with public variables V ($x \notin V$) up to probability p if and only if the process $Q \mid Q_{1\text{-ses.reach.secr.}(x)}$ satisfies the correspondence $\text{event}(\text{adv_has_x}) \Rightarrow \text{false}$ with public variables V up to probability p .

The process Q satisfies reachability secrecy of x with public variables V ($x \notin V$) up to probability p if and only if $Q \mid Q_{\text{Reach.secr.}(x)}$ satisfies the correspondence $\text{event}(\text{adv_has_x}) \Rightarrow \text{false}$ with public variables V up to probability p .

The process $Q_{1\text{-ses.reach.secr.}(x)}$ waits for a call to oracle $O_s[i_t]$ with a candidate value x' and indices u_1, \dots, u_m . If $x[u_1, \dots, u_m]$ is defined and equal to x' , the adversary managed to compute $x[u_1, \dots, u_m]$, hence to break one-session reachability secrecy. In this case, we execute event adv_has_x , and our goal will be to bound the probability of this event, by showing the correspondence $\text{event}(\text{adv_has_x}) \Rightarrow \text{false}$.

The process $Q_{\text{Reach.secr.}(x)}$ additionally provides a reveal query: by calling oracle $O_r[i_r]$ with indices u'_1, \dots, u'_m , the adversary can obtain the value of $x[u'_1, \dots, u'_m]$ if it is defined. Obviously, the adversary breaks reachability secrecy if it computes $x[u_1, \dots, u_m]$ without having first made a successful reveal query on the indices u_1, \dots, u_m . The absence of such a reveal query is verified by $\text{find } i \leq i_r \dots$ before executing event adv_has_x .

The bounds on the number of queries (n_t, n_r) are chosen large enough that they do not limit the adversary.

2.8.4 Computation of Advantages

Definition 14 Let sp be a security property: sp is $1\text{-ses.secr.}(x)$, $\text{Secrecy}(x)$, $\text{bit secr.}(x)$, or a trace property, represented by any distinguisher that does not use S , \bar{S} , nor non-unique events. (Trace properties include correspondences φ , as well as true , the property that is always true.) Let D be a disjunction of Shoup and non-unique events that does not contain S nor \bar{S} . Let C be an evaluation context acceptable for Q with any public variables and \mathcal{A} an attacker in \mathcal{BB} .

When sp is a trace property, we define $V_{sp} = \emptyset$ and

$$\text{Adv}_Q(\mathcal{A}, C, sp, D) = \Pr^{\mathcal{A}}[C[Q] : (\neg sp \vee D) \wedge \neg \text{NonUnique}_{Q,D}] \quad (12)$$

When sp is $1\text{-ses.secr.}(x)$, $\text{Secrecy}(x)$, or $\text{bit secr.}(x)$, we define $V_{sp} = \{x\}$ and

$$\text{Adv}_Q(\mathcal{A}, C, sp, D) = \Pr^{\mathcal{A}}[C[Q] : S \vee D] - \Pr^{\mathcal{A}}[C[Q] : \bar{S} \vee \text{NonUnique}_{Q,D}] \quad (13)$$

We write $\text{Bound}_Q(V, sp, D, p)$ when $V_{sp} \subseteq V$ and for all $\mathcal{A} \in \mathcal{BB}$, for all evaluation contexts C' acceptable for $C_{sp}[Q]$ with public variables $V \setminus V_{sp}$ that does not contain events used by sp or D nor non-unique events in Q , we have $\text{Adv}_Q(\mathcal{A}, C, sp, D) \leq p(\mathcal{A}, C)$ for $C = C'[C_{sp}[]]$.

The events S and \bar{S} are only executed by `event_abort`. In (13), we could write $(S \vee D) \wedge \neg \text{NonUnique}_{Q,D}$ instead of $S \vee D$, to be more similar to (12). That would be equivalent because the game immediately aborts after executing S as well as events in D and in $\text{NonUnique}_{Q,D}$, so only one of these events is executed. In (13), we expect that $C = C'[C_{sp}[]]$ for some context C' . This is what happens in the definition of $\text{Bound}_Q(V, sp, D, p)$. In that definition, C is a context acceptable for Q with public variables V .

Lemma 28 1. In the initial game Q , let $D_U = \bigvee\{e \mid [\text{unique}_e] \text{ occurs in } Q\} = \text{NonUnique}_Q$.

If $\text{Bound}_Q(V, sp, D_U, p)$, then Q satisfies property sp with public variables $V \setminus V_{sp}$ up to probability p' where $p'(\mathcal{A}, C') = p(\mathcal{A}, C'[C_{sp}[]])$.

2. If $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_p Q', D', \text{EvUsed}'$, the events S and \bar{S} are not in EvUsed' , $\text{Bound}_{Q'}(V, sp, D', p')$, and

- either sp is a trace property, $\neg sp \in \mathcal{D}$, and $p''(\mathcal{A}, C) = p(\mathcal{A}, C, t_{\neg sp}) + p'(\mathcal{A}, C)$;
- or sp is $1\text{-ses.secr.}(x)$, $\text{Secrecy}(x)$, or $\text{bit secr.}(x)$, $\{S, \bar{S}\} \subseteq \mathcal{D}$, and $p''(\mathcal{A}, C) = 2p(\mathcal{A}, C, t_S) + p'(\mathcal{A}, C)$

then $\text{Bound}_Q(V, sp, D, p'')$.

3. Let D and D' be disjunctions of Shoup and non-unique events that do not contain S nor \bar{S} .

- If sp is a trace property, $\text{Bound}_Q(V, sp, D, p)$, and $\text{Bound}_Q(V, \text{true}, D', p')$, then we have $\text{Bound}_Q(V, sp, D \vee D', p + p')$.
- If sp is $1\text{-ses.secr.}(x)$, $\text{Secrecy}(x)$, or $\text{bit secr.}(x)$, $D'_{\text{NU}} = \bigvee\{e \mid e \text{ occurs in } D' \text{ and } e \text{ is a non-unique event}\}$, $\text{Bound}_Q(V, sp, D, p)$, $\text{Bound}_Q(V, \text{true}, D', p')$, and $\text{Bound}_Q(V, \text{true}, D'_{\text{NU}}, p'')$, then $\text{Bound}_Q(V, sp, D \vee D', p + p' + p'')$.

4. If $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, \text{EvUsed} \xrightarrow{V}_p Q', D', \text{EvUsed}'$, the distinguisher D'' is a disjunction of Shoup and non-unique events in EvUsed that does not contain S nor \bar{S} , and $\text{Bound}_Q(V, \text{true}, D'', p')$, then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D \vee D'', \text{EvUsed} \xrightarrow{V}_{p+p'} Q', D', \text{EvUsed}'$.

Proof Property 1: We have $\text{NonUnique}_{Q, D_U} = D_{\text{false}}$. Let \mathcal{A} be an attacker in \mathcal{BB} and C' be an evaluation context acceptable for $C_{sp}[Q]$ with public variables $V \setminus V_{sp}$ that does not contain events used by sp , and $C = C'[C_{sp}[]]$. Since $\text{Bound}_Q(V, sp, D_U, p)$, we have $\text{Adv}_Q(\mathcal{A}, C, sp, D_U) \leq p(\mathcal{A}, C)$.

- In case sp is a trace property, $\text{Adv}_Q^{sp}(\mathcal{A}, C') = \Pr^{\mathcal{A}}[C'[Q] : \neg sp] \leq \Pr^{\mathcal{A}}[C'[Q] : \neg sp \vee D_U] = \Pr^{\mathcal{A}}[C'[C_{sp}[Q]] : \neg sp \vee D_U] = \text{Adv}_Q(\mathcal{A}, C'[C_{sp}[]], sp, D_U)$.
- In case sp is 1-ses.secr. (x) , Secrecy (x) , or bit secr. (x) , $\text{Adv}_Q^{sp}(\mathcal{A}, C') = \Pr^{\mathcal{A}}[C'[C_{sp}[Q]] : \mathbb{S}] - \Pr^{\mathcal{A}}[C'[C_{sp}[Q]] : \bar{\mathbb{S}}] \leq \Pr^{\mathcal{A}}[C'[C_{sp}[Q]] : \mathbb{S} \vee D_U] - \Pr^{\mathcal{A}}[C'[C_{sp}[Q]] : \bar{\mathbb{S}}] = \text{Adv}_Q(\mathcal{A}, C'[C_{sp}[]], sp, D_U)$.

In both cases, $\text{Adv}_Q^{sp}(\mathcal{A}, C') \leq \text{Adv}_Q(\mathcal{A}, C'[C_{sp}[]], sp, D_U) = \text{Adv}_Q(\mathcal{A}, C, sp, D_U) \leq p(\mathcal{A}, C) = p'(\mathcal{A}, C')$, so Q satisfies property sp with public variables $V \setminus V_{sp}$ up to probability p' .

Property 2, case sp is a trace property: Let \mathcal{A} be an attacker in \mathcal{BB} and C be an evaluation context acceptable for Q with public variables V that does not contain events used by sp nor D nor non-unique events of Q . Let C' be obtained by renaming the variables and tables of C that do not occur in Q to variables and tables that also do not occur in Q' , and by renaming the events of C so that they are not in $EvUsed'$. The context C' is then acceptable for Q and Q' with public variables V and does not contain events in $EvUsed'$. Since $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, EvUsed \xrightarrow{V}_p Q', D', EvUsed'$ and $\neg sp \in \mathcal{D}$, we have by taking $D_1 = D_{\text{false}}$,

$$\begin{aligned}
\text{Adv}_Q(\mathcal{A}, C, sp, D) &= \Pr^{\mathcal{A}}[C[Q] : (\neg sp \vee D) \wedge \neg \text{NonUnique}_{Q, D}] \\
&= \Pr^{\mathcal{A}}[C'[Q] : (\neg sp \vee D) \wedge \neg \text{NonUnique}_{Q, D}] \\
&\quad \text{since the renaming of events does not affect the events of the distinguisher} \\
&\leq p(\mathcal{A}, C', t_{\neg sp}) + \Pr^{\mathcal{A}}[C'[Q'] : (\neg sp \vee D') \wedge \neg \text{NonUnique}_{Q', D'}] \\
&\leq p(\mathcal{A}, C', t_{\neg sp}) + \text{Adv}_{Q'}(\mathcal{A}, C', sp, D') \\
&\leq p(\mathcal{A}, C', t_{\neg sp}) + p'(\mathcal{A}, C') \quad \text{since } \text{Bound}_{Q'}(V, sp, D', p') \\
&\leq p(\mathcal{A}, C, t_{\neg sp}) + p'(\mathcal{A}, C) \\
&\quad \text{since the renaming does not modify the probability formulas by Property 6} \\
&\leq p''(\mathcal{A}, C)
\end{aligned}$$

so $\text{Bound}_Q(V, sp, D, p'')$.

Property 2, case sp is 1-ses.secr. (x) , Secrecy (x) , or bit secr. (x) : Let \mathcal{A} be an attacker in \mathcal{BB} and C' be an evaluation context acceptable for $C_{sp}[Q]$ with public variables $V \setminus V_{sp}$ that does not contain \mathbb{S} , $\bar{\mathbb{S}}$, events used by D , nor non-unique events of Q . Let $C = C'[C_{sp}[]]$. Let C'' be obtained by renaming the variables and tables of C that do not occur in Q to variables and tables that also do not occur in Q' , and by renaming the events of C so that they are not in $EvUsed'$. (The events \mathbb{S} and $\bar{\mathbb{S}}$ are left unchanged by this renaming; they are not in $EvUsed'$.) The context C'' is then acceptable for Q and Q' with public variables V and does not contain events in $EvUsed'$. Then we have

$$\begin{aligned}
\text{Adv}_Q(\mathcal{A}, C, sp, D) &= \Pr^{\mathcal{A}}[C[Q] : \mathbb{S} \vee D] - \Pr^{\mathcal{A}}[C[Q] : \bar{\mathbb{S}} \vee \text{NonUnique}_{Q, D}] \\
&= \Pr^{\mathcal{A}}[C''[Q] : \mathbb{S} \vee D] - \Pr^{\mathcal{A}}[C''[Q] : \bar{\mathbb{S}} \vee \text{NonUnique}_{Q, D}] \\
&\quad \text{since the renaming of events does not affect the events of the distinguisher} \\
&= \Pr^{\mathcal{A}}[C''[Q] : (\mathbb{S} \vee D) \wedge \neg \text{NonUnique}_{Q, D}] \\
&\quad - \Pr^{\mathcal{A}}[C''[Q] : (\bar{\mathbb{S}} \wedge \neg D) \vee \text{NonUnique}_{Q, D}]
\end{aligned}$$

since $S \vee D$ is equivalent to $(S \vee D) \wedge \neg \text{NonUnique}_{Q,D}$ and \bar{S} is equivalent to $\bar{S} \wedge \neg D$: the game aborts immediately after executing S , \bar{S} , and the events in D and $\text{NonUnique}_{Q,D}$ so only one of them can be executed

$$\begin{aligned}
&= \Pr^{\mathcal{A}}[C''[Q] : (S \vee D) \wedge \neg \text{NonUnique}_{Q,D}] \\
&\quad + \Pr^{\mathcal{A}}[C''[Q] : (\neg \bar{S} \vee D) \wedge \neg \text{NonUnique}_{Q,D}] - 1 \\
&\leq \Pr^{\mathcal{A}}[C''[Q'] : (S \vee D') \wedge \neg \text{NonUnique}_{Q',D'}] + p(\mathcal{A}, C'', t_S) \\
&\quad + \Pr^{\mathcal{A}}[C''[Q'] : (\neg \bar{S} \vee D') \wedge \neg \text{NonUnique}_{Q',D'}] + p(\mathcal{A}, C'', t_S) - 1 \\
&\hspace{10em} \text{since } \neg \bar{S} \text{ can be implemented in the same time as } S \\
&\leq \Pr^{\mathcal{A}}[C''[Q'] : (S \vee D') \wedge \neg \text{NonUnique}_{Q',D'}] + p(\mathcal{A}, C'', t_S) \\
&\quad - \Pr^{\mathcal{A}}[C''[Q'] : (\bar{S} \wedge \neg D') \vee \text{NonUnique}_{Q',D'}] + p(\mathcal{A}, C'', t_S) \\
&\leq \Pr^{\mathcal{A}}[C''[Q'] : S \vee D'] - \Pr^{\mathcal{A}}[C''[Q'] : \bar{S} \vee \text{NonUnique}_{Q',D'}] + 2p(\mathcal{A}, C'', t_S)
\end{aligned}$$

since $S \vee D'$ is also equivalent to $(S \vee D') \wedge \neg \text{NonUnique}_{Q,D'}$ and \bar{S} is also equivalent to $\bar{S} \wedge \neg D'$

$$\begin{aligned}
&\leq \text{Adv}_{Q'}(\mathcal{A}, C'', sp, D') + 2p(\mathcal{A}, C'', t_S) \\
&\leq p'(\mathcal{A}, C'') + 2p(\mathcal{A}, C'', t_S) \hspace{10em} \text{since } \text{Bound}_{Q'}(V, sp, D', p') \\
&\leq p'(\mathcal{A}, C) + 2p(\mathcal{A}, C, t_S) \\
&\hspace{10em} \text{since the renaming does not modify the probability formulas by Property 6} \\
&\leq p''(\mathcal{A}, C)
\end{aligned}$$

so $\text{Bound}_Q(V, sp, D, p'')$.

Property 3, case sp is a trace property: Let \mathcal{A} be an attacker in \mathcal{BB} and C be an evaluation context acceptable for Q with public variables V that does not contain events used by sp , D , D' , nor non-unique events of Q . We have $\text{NonUnique}_{Q,D} = \text{NonUnique}_Q \wedge \neg D$. Therefore

$$\begin{aligned}
\text{Adv}_Q(\mathcal{A}, C, sp, D) &= \Pr^{\mathcal{A}}[C[Q] : (\neg sp \vee D) \wedge \neg \text{NonUnique}_{Q,D}] \\
&= \Pr^{\mathcal{A}}[C[Q] : (\neg sp \vee D) \wedge \neg (\text{NonUnique}_Q \wedge \neg D)] \\
&= \Pr^{\mathcal{A}}[C[Q] : (\neg sp \vee D) \wedge (\neg \text{NonUnique}_Q \vee D)] \\
&= \Pr^{\mathcal{A}}[C[Q] : (\neg sp \wedge \neg \text{NonUnique}_Q) \vee D]
\end{aligned}$$

In particular, $\text{Adv}_Q(\mathcal{A}, C, \text{true}, D) = \Pr^{\mathcal{A}}[C[Q] : D]$. Hence

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

since $\text{Bound}_Q(V, sp, D, p)$ and $\text{Bound}_Q(V, \text{true}, D', p')$. So $\text{Bound}_Q(V, sp, D \vee D', p + p')$.

Property 3, case sp is 1-ses.secr.(x), Secrecy(x), or bit secr.(x): Let \mathcal{A} be an attacker in \mathcal{BB} and C' be an evaluation context acceptable for $C_{sp}[Q]$ with public variables $V \setminus V_{sp}$ that does not contain S , \bar{S} , events used by D , D' , nor non-unique events of Q . Let $C = C'[C_{sp}[[]]$. Since $\text{NonUnique}_{Q,D} = \text{NonUnique}_Q \wedge \neg D$, we have

$$\text{Adv}_Q(\mathcal{A}, C, sp, D) = \Pr^{\mathcal{A}}[C[Q] : S \vee D] - \Pr^{\mathcal{A}}[C[Q] : \bar{S} \vee (\text{NonUnique}_Q \wedge \neg D)]$$

Hence

$$\begin{aligned}
\text{Adv}_Q(\mathcal{A}, C, sp, D \vee D') &= \Pr^{\mathcal{A}}[C[Q] : \mathbb{S} \vee D \vee D'] \\
&\quad - \Pr^{\mathcal{A}}[C[Q] : \bar{\mathbb{S}} \vee (\text{NonUnique}_Q \wedge \neg D \wedge \neg D')] \\
&\leq \Pr^{\mathcal{A}}[C[Q] : \mathbb{S} \vee D] + \Pr^{\mathcal{A}}[C[Q] : D'] \\
&\quad - \Pr^{\mathcal{A}}[C[Q] : \bar{\mathbb{S}} \vee (\text{NonUnique}_Q \wedge \neg D)] + \Pr^{\mathcal{A}}[C[Q] : D'_{\text{NU}}] \\
&\leq \text{Adv}_Q(\mathcal{A}, C, sp, D) + \text{Adv}_Q(\mathcal{A}, C, \text{true}, D') + \text{Adv}_Q(\mathcal{A}, C, \text{true}, D'_{\text{NU}}) \\
&\leq p(\mathcal{A}, C) + p'(\mathcal{A}, C) + p''(\mathcal{A}, C)
\end{aligned}$$

since $\text{Bound}_Q(V, sp, D, p)$, $\text{Bound}_Q(V, \text{true}, D', p')$, $\text{Bound}_Q(V, \text{true}, D'_{\text{NU}}, p'')$, and C is also an evaluation context acceptable for Q with public variables V . So $\text{Bound}_Q(V, sp, D \vee D', p+p'+p'')$.

The inequality $-\Pr^{\mathcal{A}}[C[Q] : \bar{\mathbb{S}} \vee (\text{NonUnique}_Q \wedge \neg D \wedge \neg D')] \leq -\Pr^{\mathcal{A}}[C[Q] : \bar{\mathbb{S}} \vee (\text{NonUnique}_Q \wedge \neg D)] + \Pr^{\mathcal{A}}[C[Q] : D'_{\text{NU}}]$ used above is justified as follows:


$$\begin{aligned}
&\Pr^{\mathcal{A}}[C[Q] : \bar{\mathbb{S}} \vee (\text{NonUnique}_Q \wedge \neg D)] \\
&\leq \Pr^{\mathcal{A}}[C[Q] : \bar{\mathbb{S}} \vee (\text{NonUnique}_Q \wedge \neg D \wedge \neg D') \vee (\text{NonUnique}_Q \wedge D')] \\
&\leq \Pr^{\mathcal{A}}[C[Q] : \bar{\mathbb{S}} \vee (\text{NonUnique}_Q \wedge \neg D \wedge \neg D')] + \Pr^{\mathcal{A}}[C[Q] : \text{NonUnique}_Q \wedge D'] \\
&\leq \Pr^{\mathcal{A}}[C[Q] : \bar{\mathbb{S}} \vee (\text{NonUnique}_Q \wedge \neg D \wedge \neg D')] + \Pr^{\mathcal{A}}[C[Q] : D'_{\text{NU}}]
\end{aligned}$$

Property 4: The distinguisher $D \vee D''$ is a disjunction of Shoup and non-unique events in EvUsed . Let \mathcal{A} be an attacker in \mathcal{BB} and C be any evaluation context acceptable for Q with public variables V that does not contain events in EvUsed' . Let $D_0 \in \mathcal{D} \cup \{D_{\text{false}}\}$. Let D_1 be a disjunction of events in \mathcal{D}_{SNU} . We have


$$\begin{aligned}
&\Pr^{\mathcal{A}}[C[Q] : (D_0 \vee D_1 \vee D \vee D'') \wedge \neg \text{NonUnique}_{Q, D_1 \vee D \vee D''}] \\
&= \Pr^{\mathcal{A}}[C[Q] : (D_0 \wedge \neg \text{NonUnique}_Q) \vee D_1 \vee D \vee D''] \\
&\quad \text{as in Property 3, case } sp \text{ is a trace property} \\
&\leq \Pr^{\mathcal{A}}[C[Q] : (D_0 \wedge \neg \text{NonUnique}_Q) \vee D_1 \vee D] + \Pr^{\mathcal{A}}[C[Q] : D''] \\
&\leq \Pr^{\mathcal{A}}[C[Q] : (D_0 \vee D_1 \vee D) \wedge \neg \text{NonUnique}_{Q, D_1 \vee D}] + \text{Adv}_Q(\mathcal{A}, C, \text{true}, D'') \\
&\text{since } \text{Adv}_Q(\mathcal{A}, C, \text{true}, D'') = \Pr^{\mathcal{A}}[C[Q] : D''] \text{ (see Property 3, case } sp \text{ is a trace property)} \\
&\leq \Pr^{\mathcal{A}}[C[Q'] : (D_0 \vee D_1 \vee D') \wedge \neg \text{NonUnique}_{Q, D_1 \vee D'}] + p(\mathcal{A}, C, t_{D_0}) + p'(\mathcal{A}, C)
\end{aligned}$$

since $\text{Bound}_Q(V, \text{true}, D'', p')$. (Note that the context C does not contain events used by D'' nor non-unique events of Q .) \square

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 from $\text{Bound}_Q(V, sp, D_U, p)$. (Using the distinguisher D_U can also be understood by saying that we consider that the adversary wins if some non-unique event is executed, that is, if a `find` or `get` declared `unique` by the user actually has several possible choices. That allows the implementation to make any choice when a `find[uniquee]` or `get[uniquee]` has several possible choices: the security proof remains valid. In particular, a `find[uniquee]` or `get[uniquee]` can be implemented by always choosing the first found element.) 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 breaking sp and of D on one side and of D' on the other side. We bound these two



(a) edge, source, and target nodes



(b) node and outgoing edges

Figure 15: Structure of a proof tree

probabilities by $\text{Bound}_Q(V, sp, D, p)$ and $\text{Bound}_Q(V, \text{true}, D', p')$ separately, then obtain a bound $\text{Bound}_Q(V, sp, D \vee D', p'')$ by computing a sum. (When we deal with secrecy and $D'_{\text{NU}} \neq D'_{\text{false}}$, the probability of D'_{NU} can be bounded by looking at the proof for D' .)

More formally, consider the following cases, using Lemma 28, Property 1:

- If we want to prove that Q_0 satisfies the correspondence φ with public variables V , then we let $sp = \varphi$.
- If we want to prove that Q_0 satisfies the (one-session or bit) secrecy of x with public variables V' ($x \notin V'$), then we let sp be 1-ses.secr. (x) , Secrecy (x) , or bit secr. (x) and $V = V' \cup \{x\}$.

In both cases, we show $\text{Bound}_{Q_0}(V, sp, D_U, p)$. The proof produced by CryptoVerif can be represented as a tree whose nodes are labeled with quintuples $(Q, V, sp', D, EvUsed)$ and whose edges are labeled with triples (sp'', D'', p'') , where Q is the current game, V is the set of public variables, sp' and sp'' are either the initial property to prove sp or true, D and D'' are disjunctions of Shoup and non-unique events, $EvUsed$ is the set of events used so far, and p'' is a probability formula. The edges have a single source node, but may have 0, 1, or several target nodes. We associate to each node labeled with $(Q, V, sp', D, EvUsed)$ a property $\text{Bound}_Q(V, sp', D, p)$, and to each edge labeled with (sp'', D'', p'') with source node labeled with $(Q, V, sp', D, EvUsed)$ a property $\text{Bound}_Q(V, sp'', D'', p')$. CryptoVerif computes the probabilities p, p' such that these properties hold from the leaves of the tree to its root, as explained next.

The root of the tree is labeled with $(Q_0, V, sp, D_U, EvUsed_0)$ such that $EvUsed_0$ is the set containing the events used by correspondences to prove or that occur in Q_0 .

When an edge labeled with (sp'', D'', p'') has a source node labeled with $(Q, V, -, -, EvUsed)$ and target nodes labeled with $(Q_j, V_j, sp'_j, D_j, EvUsed_j)$ for $j \in \{1, \dots, l\}$ (Figure 15(a)), the bound associated to the edge can be computed from the bound associated to the target nodes by the probability formula p'' that labels the edge: if $\text{Bound}_{Q_j}(V_j, sp'_j, D_j, p_j)$ for $j \in \{1, \dots, l\}$, then $\text{Bound}_Q(V, sp'', D'', p)$ where $p(A, C) = p''(A, C, p_1(C), \dots, p_l(C))$. This situation corresponds to a game transformation that transforms game Q into games Q_j ($j \in \{1, \dots, l\}$). We can distinguish several cases depending on where the edge comes from:

- For most game transformations, the edge has a single target node ($l = 1$), $V_1 = V$, $sp'_1 = sp''$, and the game transformation transforms Q into Q_1 and satisfies $\mathcal{D}_1, \emptyset : Q, D'', EvUsed \xrightarrow{V}_{p'} Q_1, D_1, EvUsed_1$ where $\mathcal{D}_1 = \{\neg sp''\}$ when sp'' is a trace property, $\mathcal{D}_1 = \{S, \bar{S}\}$ when sp'' is 1-ses.secr.(x), Secrecy(x), or bit secr.(x), and $\mathcal{D}_1 = \emptyset$ when $sp'' = \text{true}$. (During the building of the proof tree, we have $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, EvUsed \xrightarrow{V}_{p'} Q_1, D', EvUsed_1$, where the distinguishers \mathcal{D} correspond to the active queries not for introduced Shoup and non-unique events, so $\mathcal{D}_1 \subseteq \mathcal{D}$, \mathcal{D}_{SNU} are the active queries for Shoup and non-unique events both before and after this step so D'' and D_1 are disjunctions of events in \mathcal{D}_{SNU} , $D = D'' \wedge \neg D_1$ are the Shoup/non-unique events proved at this step, $D' = D_1 \wedge \neg D''$ are the Shoup/non-unique events introduced at this step. By applying several times Lemma 20, Property 5, we obtain $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D'', EvUsed \xrightarrow{V}_{p'} Q_1, D_1, EvUsed_1$. By Lemma 20, Property 6, we obtain $\mathcal{D}_1, \emptyset : Q, D'', EvUsed \xrightarrow{V}_{p'} Q_1, D_1, EvUsed_1$.) The bound is inferred by Lemma 28, Property 2.

If sp'' is a trace property and $\text{Bound}_{Q_1}(V, sp'', D_1, p_1)$, then $\text{Bound}_Q(V, sp'', D'', p)$ where $p(\mathcal{A}, C) = p'(\mathcal{A}, C, t_{\neg sp''}) + p_1(\mathcal{A}, C)$, so we can define $p''(\mathcal{A}, C, p_t) = p'(\mathcal{A}, C, t_{\neg sp''}) + p_t$.

If sp'' is 1-ses.secr.(x), Secrecy(x), or bit secr.(x) and $\text{Bound}_{Q_1}(V, sp'', D_1, p_1)$, then we have $\text{Bound}_Q(V, sp'', D'', p)$ where $p(\mathcal{A}, C) = 2p'(\mathcal{A}, C, t_S) + p_1(\mathcal{A}, C)$, so we can define $p''(\mathcal{A}, C, p_t) = 2p'(\mathcal{A}, C, t_S) + p_t$.

To unify these two cases, we define

$$\text{pstd}_{p'}^{sp''}(\mathcal{A}, C, p_t) = \begin{cases} p'(\mathcal{A}, C, t_{\neg sp''}) + p_t & \text{when } sp'' \text{ is a trace property} \\ 2p'(\mathcal{A}, C, t_S) + p_t & \text{when } sp'' \text{ is 1-ses.secr.}(x), \text{ Secrecy}(x) \text{ or bit secr.}(x) \end{cases}$$

so that, when an edge comes from a transformation $\mathcal{D}_1, \emptyset : Q, D'', EvUsed \xrightarrow{V}_{p'} Q_1, D_1, EvUsed_1$ and the considered security property is sp'' , the edge can be labeled with the probability formula $\text{pstd}_{p'}^{sp''}$. If $\text{Bound}_{Q_1}(V, sp'', D_1, p_1)$, then $\text{Bound}_Q(V, sp'', D'', p)$ where $p(\mathcal{A}, C) = \text{pstd}_{p'}^{sp''}(\mathcal{A}, C, p_1(C))$.

- When a query is proved (by the command **success**, Section 4), the edge has no target node ($l = 0$), and we simply obtain $\text{Bound}_Q(V, sp'', D'', p'')$. The probability p'' that labels the edge is determined by the **success** command (Proposition 1, 2, or 3).

These propositions require that $D'' = D_{\text{false}}$. This is obtained by splitting the properties to prove one property at a time (with one edge for each property starting from the source node), yielding bounds of the form $\text{Bound}_Q(V, sp'', D_{\text{false}}, p'')$ or $\text{Bound}_Q(V, \text{true}, e, p'')$. For the first form, p'' can immediately be computed from Proposition 1, 2, or 3. For the second form, when e is a non-unique event, we use Section 4.2.2 and when e is a Shoup event, we notice that $\text{Adv}_Q(\mathcal{A}, C, \text{true}, e) = \Pr^A[C[Q] : e \wedge \neg \text{NonUnique}_{Q,e}] = \Pr^A[C[Q] : e \wedge \neg \text{NonUnique}_Q] = \text{Adv}_Q(\mathcal{A}, C, [\text{event}(e) \Rightarrow \text{false}], D_{\text{false}})$, so we can use $\text{Bound}_Q(V, [\text{event}(e) \Rightarrow \text{false}], D_{\text{false}}, p'')$ instead.

The **success** command is the only one that removes an event from D'' , which then happens only when we evaluate $\text{Bound}_Q(V, \text{true}, e, p'')$, so $D'' = e$, $sp'' = \text{true}$, $l = 0$.

- In case of other transformations such as the **guess** transformation, the relation between bounds that defines p'' is given directly in the soundness lemma for the transformation (Lemma 59, 60, or 61). In particular, for the transformation **guess.branch**, the edge has as

many target nodes as there are branches in the guessed instruction. For the transformation **guess** i , Lemma 59 requires $D'' = D_{\text{false}}$, which can be achieved as for **success** above.

When a node labeled with $(Q, V, sp', D, EvUsed)$ has outgoing edges labeled respectively $(sp'_1, D'_1, -), \dots, (sp'_l, D'_l, -)$ (Figure 15(b)), then $D = D'_1 \vee \dots \vee D'_l$ (D is a disjunction of the form $e_1 \vee \dots \vee e_m$, D'_1, \dots, D'_l are disjunctions that form a partition of the disjuncts of D), there exists $j_0 \leq l$ such that $sp'_{j_0} = sp'$ and for all $j \neq j_0$, $sp'_j = \text{true}$. The bound associated to the node is computed from the bound associated to the edges by Lemma 28, Property 3:

- If sp' is a trace property, then we have $\text{Bound}_Q(V, sp', D, p'_1 + \dots + p'_l)$, where $\text{Bound}_Q(V, sp'_j, D'_j, p'_j)$ for all $j \in \{1, \dots, l\}$.
- If sp' is $1\text{-ses.secr.}(x)$, $\text{Secrecy}(x)$, or $\text{bit secr.}(x)$, then we have $\text{Bound}_Q(V, sp', D, \sum_{j=1}^l p'_j + \sum_{j=1, \dots, l; j \neq j_0} p''_j)$, where $\text{Bound}_Q(V, sp'_j, D'_j, p'_j)$ for all $j \in \{1, \dots, l\}$, $\text{Bound}_Q(V, sp'_j, D'_{\text{NU}j}, p''_j)$ for all $j \in \{1, \dots, l\}$ with $j \neq j_0$, and $D'_{\text{NU}j}$ is obtained from D'_j by keeping only the non-unique events.

To prove $\text{Bound}_Q(V, sp'_j, D'_{\text{NU}j}, p''_j)$, we build a proof tree with root $Q, V, sp'_j, D'_{\text{NU}j}, EvUsed$ from the subtree $Q, V, sp', D, EvUsed \xrightarrow{sp'_j, D'_j, p'_j} \dots$ by

- replacing the root $Q, V, sp', D, EvUsed$ with $Q, V, sp'_j, D'_j, EvUsed$.
- removing all Shoup events of D'_j from distinguishers that label nodes and edges. In particular, the root $Q, V, sp'_j, D'_j, EvUsed$ then becomes $Q, V, sp'_j, D'_{\text{NU}j}, EvUsed$.
- removing subtrees that start with an edge labeled $sp'_j, D_{\text{false}}, p$ for some p .

The proof steps remain valid: all proof steps are a fortiori valid when we ignore some Shoup events. That can be verified for each transformation using its soundness lemma. For instance, for proof steps that come from usual transformations that satisfy property preservation with introduction of events and have a Shoup event e of D'_j both before and after, we can avoid adding that event e when we derive $\mathcal{D}_1, \emptyset : Q, D'_j, EvUsed \xrightarrow{V}_{p'} Q_1, D_1, EvUsed_1$ from $\mathcal{D}, \mathcal{D}_{\text{SNU}} : Q, D, EvUsed \xrightarrow{V}_{p'} Q_1, D', EvUsed_1$. For proof steps that prove a Shoup event e of D'_j (via **success**), that proof step starts with just event e , so it is simply removed. We can then apply the previous reasoning to that proof tree.

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. Only the final proof step is distinct for each query. 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 16; they are explained below.

The bound associated to the leaves of the tree is computed by **success**; the bound associated to an edge is computed from the bounds associated to its target nodes, and the bound associated to a node is computed from the bounds associated to its outgoing edges. We can then compute the bounds associated to all nodes of the tree, by induction from the leaves to the root. At the root, we obtain a bound $\text{Bound}_{Q_0}(V, sp, D_U, p)$ that yields the desired result.

Lemma 28 allows us to obtain more precise probability bounds than the standard computation of probabilities generally done by cryptographers, when we use Shoup's lemma [68]. 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




Figure 16: Examples of proof trees

all attackers $\mathcal{A} \in \mathcal{BB}$ and all contexts C acceptable for G and G' (with any public variables) and all distinguishers D , $|\Pr^{\mathcal{A}}[C[G] : D] - \Pr^{\mathcal{A}}[C[G'] : D]| \leq \Pr^{\mathcal{A}}[C[G'] : e]$. Hence,

$$\Pr^{\mathcal{A}}[C[G] : D] \leq \Pr^{\mathcal{A}}[C[G'] : e] + \Pr^{\mathcal{A}}[C[G'] : D].$$

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

$$\Pr^{\mathcal{A}}[C[G] : D] \leq \Pr^{\mathcal{A}}[C[G'] : D \vee e].$$

For example, suppose that we want to bound the probability of event e_0 in G_0 : we define $sp = \llbracket \text{event}(e_0) \Rightarrow \text{false} \rrbracket = \neg e_0$. 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 $\{e_0\}, \emptyset : G_0, D_{\text{false}}, \text{EvUsed}_0 \rightarrow_0 G_1, e, \text{EvUsed}_1$; then we transform G_1 into G_2 , so that $G_1 \approx_p^{\{e_0, e\}} G_2$, so we have $\{e_0\}, \emptyset : G_1, e, \text{EvUsed}_1 \rightarrow_p G_2, e, \text{EvUsed}_1$ by Lemma 20, Property 1a; and G_2 executes neither e_0 nor e . We suppose for simplicity that no $\llbracket \text{unique}_{e'} \rrbracket$ occurs, so that $\text{NonUnique}_{G_i, D}$ is always false. The corresponding proof tree is given in Figure 16(a).

- Since e_0 does not occur in G_2 , we have $\text{Adv}_{G_2}(\mathcal{A}, C, sp, D_{\text{false}}) = 0$ for all $\mathcal{A} \in \mathcal{BB}$ and all evaluation contexts C acceptable for G_2 with public variables V that do not contain event e_0 . So $\text{Bound}_{G_2}(V, sp, D_{\text{false}}, 0)$. Similarly, $\text{Bound}_{G_2}(V, \text{true}, e, 0)$. These two properties are represented in the proof tree by the two edges outgoing from node $G_2, V, sp, e, \{e_0, e\}$.
- By Lemma 28, Property 3, $\text{Bound}_{G_2}(V, sp, e, p'_3)$ where $p'_3(C) = 0$.
- Since $\{e_0\}, \emptyset : G_1, e, \text{EvUsed}_1 \rightarrow_p G_2, e, \text{EvUsed}_1$, by Lemma 28, Property 2, we obtain $\text{Bound}_{G_1}(V, sp, e, p'_2)$ where $p'_2(\mathcal{A}, C) = p(\mathcal{A}, C, t_{\neg sp}) + p'_3(\mathcal{A}, C)$. This is represented in the proof tree by the edge from node $G_1, V, sp, e, \{e_0, e\}$ to node $G_2, V, sp, e, \{e_0, e\}$ labeled with $sp, e, \text{pstd}_p^{\text{sp}}$.




Figure 17: Removing subtrees

- Since $\{e_0\}, \emptyset : G_0, D_{\text{false}}, \text{EvUsed}_0 \rightarrow_0 G_1, e, \text{EvUsed}_1$, by Lemma 28, Property 2, we obtain $\text{Bound}_{G_0}(V, sp, D_{\text{false}}, p'_1)$ where $p'_1(\mathcal{A}, C) = 0 + p'_2(\mathcal{A}, C)$. This is represented in the proof tree by the edge from node $G_0, V, sp, D_{\text{false}}, \{e_0\}$ to node $G_1, V, sp, e, \{e_0, e\}$ labeled with $sp, D_{\text{false}}, \text{pstd}_{p'_1}^{sp}$.
- Finally, by Lemma 28, Property 1, we conclude that G_0 satisfies sp with public variables V up to probability p'_1 , where $p'_1(\mathcal{A}, C) = p'_2(\mathcal{A}, C) = p(\mathcal{A}, C, t_{\neg sp}) + p'_3(\mathcal{A}, C) = p(\mathcal{A}, C, t_{e_0})$, which means that $\Pr^{\mathcal{A}}[C[G_0] : e_0] \leq p(\mathcal{A}, C, t_{e_0})$ for all evaluation contexts C acceptable for G_0 with public variables V that do not contain event e_0 .

Let C be an evaluation context acceptable for G_0 with public variables V that does not contain event e_0 and \mathcal{A} an attacker in \mathcal{BB} . Since we suppose for simplicity that no $[\text{unique}_{e'}]$ occurs, so that $\text{NonUnique}_{G_i, D}$ is always false, we have $\text{Adv}_{G_i}(\mathcal{A}, C, sp, D) = \Pr^{\mathcal{A}}[C[G_i] : \neg sp \vee D]$, so we can write the previous computation simply using probabilities:

$$\begin{aligned}
 \Pr^{\mathcal{A}}[C[G_0] : e_0] &\leq \Pr^{\mathcal{A}}[C[G_1] : e_0 \vee e] && \text{since } \{e_0\}, \emptyset : G_0, D_{\text{false}}, \text{EvUsed}_0 \rightarrow_0 G_1, e, \text{EvUsed}_1 \\
 &\leq p(\mathcal{A}, C, t_{e_0}) + \Pr^{\mathcal{A}}[C[G_2] : e_0 \vee e] && \text{since } \{e_0\}, \emptyset : G_1, e, \text{EvUsed}_1 \rightarrow_p G_2, e, \text{EvUsed}_1 \\
 &\leq p(\mathcal{A}, C, t_{e_0}) && \text{since } G_2 \text{ executes neither } e_0 \text{ nor } e.
 \end{aligned}$$

In contrast, the standard computation of probabilities yields

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

The runtime t_D of D is essentially the same for e_0, e , and $e_0 \vee e$, so $\Pr^{\mathcal{A}}[C[G_0] : e_0] \leq p(\mathcal{A}, C, t_D)$ by Lemma 28, while $\Pr^{\mathcal{A}}[C[G_0] : e_0] \leq 2p(\mathcal{A}, C, t_D)$ 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 28 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 28 only when the proof uses different sequences of games to bound the probabilities of the events, as in Figure 16(b).

Consider a proof tree that consists of a main branch that is a sequence of applications of transformations that satisfy property preservation with introduction of events (properties of the

form $\mathcal{D}, \emptyset : Q_i, D'_i, EvUsed_i \xrightarrow{V}_{p_i} Q_{i+1}, D_{i+1}, EvUsed_{i+1}$), and side branches that may use any transformation to bound the probability of Shoup and non-unique events. All nodes on the main branch use the same security property sp , while nodes on side branches use true as security property. Such a proof tree happens when we prove indistinguishability properties, where sp is any distinguisher used in the definition of indistinguishability (see Section 2.8.5). In particular, the transformations **guess**, **guess_branch**, and **success_simplify** are not allowed in the main branch but may be used in side branches. Lemma 28, Property 4 allows one to transform such a proof tree into a single property of the form $\mathcal{D}, \emptyset : Q, D, EvUsed \xrightarrow{V}_p Q', D', EvUsed'$. Indeed, the main branch starts from the root $Q, V, sp, D, EvUsed$ to a leaf $Q', V, sp, D', EvUsed$, with additional subtrees starting from various nodes on this branch. Each node on the main branch is labeled $Q_i, V, sp, D_i, EvUsed_i$, and the edge of the main branch that starts from $Q_i, V, sp, D_i, EvUsed_i$ is labeled $sp, D'_i, \text{pstd}_{p_i}^{sp}$ (see Figure 17). Suppose an additional subtree starts from a node $Q_i, V, sp, D_i, EvUsed_i$ with an edge labeled true, D''_i, \dots . This subtree yields a bound $\text{Bound}_{Q_i}(V, \text{true}, D''_i, p'_i)$. The edge of the main branch from $Q_i, V, sp, D_i, EvUsed_i$ yields a property $\mathcal{D}, \emptyset : Q_i, D'_i, EvUsed_i \xrightarrow{V}_{p_i} Q_{i+1}, D_{i+1}, EvUsed_{i+1}$. By Lemma 28, Property 4, the additional subtree can then be removed from the proof tree by replacing $sp, D'_i, \text{pstd}_{p_i}^{sp}$ with $sp, D'_i \vee D''_i, \text{pstd}_{p_i+p'_i}^{sp}$ as label of the edge of the main branch starting from the node $Q_i, V, sp, D_i, EvUsed_i$. By repeating this operation, we remove all additional subtrees, obtaining a proof tree that consists of a single branch with nodes labeled $Q_i, V, sp, D_i, EvUsed_i$, such that the edge that starts from $Q_i, V, sp, D_i, EvUsed_i$ is labeled $sp, D_i, \text{pstd}_{p_i}^{sp}$. This yields a sequence of properties $\mathcal{D}, \emptyset : Q_i, D_i, EvUsed_i \xrightarrow{V}_{p_i} Q_{i+1}, D_{i+1}, EvUsed_{i+1}$, which yields a single such property $\mathcal{D}, \emptyset : Q, D, EvUsed \xrightarrow{V}_p Q', D', EvUsed'$ by transitivity (Lemma 20, Property 3).

2.8.5 Proof of Indistinguishability

To prove indistinguishability between two games G_0 and G_1 , CryptoVerif finds a game G_2 such that $\mathcal{D}, \emptyset : G_0, D_{U_0}, EvUsed \xrightarrow{V}_p G_2, D_2, EvUsed'$ and $\mathcal{D}, \emptyset : G_1, D_{U_1}, EvUsed_1 \xrightarrow{V}_{p', D+} G_2, D_2, EvUsed'_1$ where \mathcal{D} is the set of all distinguishers, $EvUsed = \text{event}(G_0)$, $EvUsed_1 = \text{event}(G_1)$, $D_{U_0} = \bigvee \{e \mid [\text{unique}_e] \text{ occurs in } G_0\}$ and $D_{U_1} = \bigvee \{e \mid [\text{unique}_e] \text{ occurs in } G_1\}$. The active queries D_2 are also required to be the same in both sequences of games. (In general, CryptoVerif builds proof trees; they can be transformed into the properties above by Lemma 28, Property 4 as explained above. Only transformations that satisfy property preservation with introduction of events are allowed in the sequence of games that proves indistinguishability. The transformations **guess**, **guess_branch**, and **success_simplify** are not allowed in that sequence, but are allowed in side branches that bound the probability of introduced events.) So for all attackers $\mathcal{A} \in \mathcal{BB}$, all evaluation contexts C acceptable for G_0 and G_2 with public variables V that do not contain events $EvUsed'$, and all distinguishers $D_0 \in \mathcal{D}$ that run in time at most t_{D_0} ,

$$\Pr^{\mathcal{A}}[C[G_0] : D_0 \vee D_{U_0}] \leq \Pr^{\mathcal{A}}[C[G_2] : (D_0 \vee D_2) \wedge \neg \text{NonUnique}_{G_2, D_2}] + p(\mathcal{A}, C, t_{D_0}) \quad (14)$$

since $\text{NonUnique}_{G_0, D_{U_0}} = D_{\text{false}}$, and for all attackers $\mathcal{A} \in \mathcal{BB}$, all evaluation contexts C acceptable for G_1 and G_2 with public variables V that do not contain events in $EvUsed'_1$, and all distinguishers $D_0 \in \mathcal{D}$ that run in time at most t_{D_0} ,

$$\Pr^{\mathcal{A}}[C[G_1] : D_0 \vee D_{U_1}] \leq \Pr^{\mathcal{A}}[C[G_2] : (D_0 \vee D_2) \wedge \neg \text{NonUnique}_{G_2, D_2}] + p'(\mathcal{A}, C, t_{D_0}). \quad (15)$$

Let \mathcal{A} be an attacker in \mathcal{BB} and C be an evaluation context acceptable for G_0 and G_1 with public variables V . After renaming the variables of C that do not occur in G_0 and G_1 and the tables of C that do not occur in G_0 and G_1 so that they do not occur in G_2 , C is also acceptable

for G_2 with public variables V . Furthermore, by Property 6, this renaming does not change the probabilities. Let $D_0 \in \mathcal{D}$ be a distinguisher that runs in time at most t_{D_0} . We rename the events of C in $EvUsed'$ or $EvUsed'_1$ to some fresh events, and modify D_0 so that it considers the renamed events as if they were the original events. That does not change the probability $|\Pr^{\mathcal{A}}[C[G_0] : D_0] - \Pr^{\mathcal{A}}[C[G_1] : D_0]|$, and guarantees that C does not contain events in $EvUsed'$ nor in $EvUsed'_1$. So

$$\begin{aligned} \Pr^{\mathcal{A}}[C[G_0] : D_0] &\leq \Pr^{\mathcal{A}}[C[G_0] : D_0 \vee D_{U0}] \\ &\leq \Pr^{\mathcal{A}}[C[G_2] : (D_0 \vee D_2) \wedge \neg \text{NonUnique}_{G_2, D_2}] + p(\mathcal{A}, C, t_{D_0}) \\ &\leq \Pr^{\mathcal{A}}[C[G_2] : ((D_0 \wedge \neg D_2) \vee D_2) \wedge \neg \text{NonUnique}_{G_2, D_2}] + p(\mathcal{A}, C, t_{D_0}) \\ &\leq \Pr^{\mathcal{A}}[C[G_2] : (D_0 \wedge \neg D_2) \wedge \neg \text{NonUnique}_{G_2, D_2}] \\ &\quad + \Pr^{\mathcal{A}}[C[G_2] : D_2 \wedge \neg \text{NonUnique}_{G_2, D_2}] + p(\mathcal{A}, C, t_{D_0}) \\ &\leq \Pr^{\mathcal{A}}[C[G_2] : D_0 \wedge \neg D_2] + \Pr^{\mathcal{A}}[C[G_2] : D_2 \wedge \neg \text{NonUnique}_{G_2, D_2}] + p(\mathcal{A}, C, t_{D_0}) \end{aligned}$$

By applying (15) to $\neg D_0$, which is also in \mathcal{D} and runs in the same time as D_0 , we have

$$\begin{aligned} 1 - \Pr^{\mathcal{A}}[C[G_1] : D_0] &= \Pr^{\mathcal{A}}[C[G_1] : \neg D_0] \\ &\leq \Pr^{\mathcal{A}}[C[G_1] : \neg D_0 \vee D_{U1}] \\ &\leq \Pr^{\mathcal{A}}[C[G_2] : (\neg D_0 \vee D_2) \wedge \neg \text{NonUnique}_{G_2, D_2}] + p'(\mathcal{A}, C, t_{D_0}) \\ &\leq \Pr^{\mathcal{A}}[C[G_2] : \neg D_0 \vee D_2] + p'(\mathcal{A}, C, t_{D_0}) \\ &\leq 1 - \Pr^{\mathcal{A}}[C[G_2] : D_0 \wedge \neg D_2] + p'(\mathcal{A}, C, t_{D_0}) \end{aligned}$$

so

$$-\Pr^{\mathcal{A}}[C[G_1] : D_0] \leq -\Pr^{\mathcal{A}}[C[G_2] : D_0 \wedge \neg D_2] + p'(\mathcal{A}, C, t_{D_0})$$

so

$$\begin{aligned} \Pr^{\mathcal{A}}[C[G_0] : D_0] - \Pr^{\mathcal{A}}[C[G_1] : D_0] \\ \leq \Pr^{\mathcal{A}}[C[G_2] : D_2 \wedge \neg \text{NonUnique}_{G_2, D_2}] + p(\mathcal{A}, C, t_{D_0}) + p'(\mathcal{A}, C, t_{D_0}) \end{aligned}$$

By applying the formula above to $\neg D_0$, which runs in the same time as D_0 , we have

$$\begin{aligned} \Pr^{\mathcal{A}}[C[G_1] : D_0] - \Pr^{\mathcal{A}}[C[G_0] : D_0] \\ \leq \Pr^{\mathcal{A}}[C[G_2] : D_2 \wedge \neg \text{NonUnique}_{G_2, D_2}] + p(\mathcal{A}, C, t_{D_0}) + p'(\mathcal{A}, C, t_{D_0}) \end{aligned}$$

so

$$\begin{aligned} |\Pr^{\mathcal{A}}[C[G_0] : D_0] - \Pr^{\mathcal{A}}[C[G_1] : D_0]| \\ \leq \Pr^{\mathcal{A}}[C[G_2] : D_2 \wedge \neg \text{NonUnique}_{G_2, D_2}] + p(\mathcal{A}, C, t_{D_0}) + p'(\mathcal{A}, C, t_{D_0}) \end{aligned}$$

so $G_0 \approx_{p''}^V G_1$ where $p''(\mathcal{A}, C, t_{D_0}) = \Pr^{\mathcal{A}}[C[G_2] : D_2 \wedge \neg \text{NonUnique}_{G_2, D_2}] + p(\mathcal{A}, C, t_{D_0}) + p'(\mathcal{A}, C, t_{D_0})$.

2.8.6 Proof of query_equiv

Decisional case (query_equiv without [computational] annotation) The situation is similar to the proof of indistinguishability, but the property we want to prove is $\mathcal{D}_{\neg EvUsed_1, \emptyset} : G_0, D_{\text{false}}, \emptyset \xrightarrow{V}_p G_1, D_1, EvUsed_1$ where G_0 contains no events, $D_1 = e_1 \vee \dots \vee e_m$, e_1, \dots, e_m

are the Shoup events occurring in G_1 , e'_1, \dots, e'_l are the non-unique events occurring in G_1 , $EvUsed_1 = \{e_1, \dots, e_m, e'_1, \dots, e'_l\}$, $V = \emptyset$.

We need to show that, for all attackers $\mathcal{A} \in \mathcal{BB}$, all evaluation contexts C acceptable for G_0 and G_1 without public variables that do not contain events in $EvUsed_1$ and all distinguishers $D_0 \in \mathcal{D}_{\neg EvUsed_1}$,

$$\Pr^{\mathcal{A}}[C[G_0] : D_0] \leq \Pr^{\mathcal{A}}[C[G_1] : (D_0 \vee D_1) \wedge \neg \text{NonUnique}_{G_1, D_1}] + p(\mathcal{A}, C, t_{D_0})$$

CryptoVerif finds a game G_2 such that $\mathcal{D}_{\neg EvUsed'_0, \emptyset} : G_0, D_{\text{false}}, \emptyset \xrightarrow{V}_{p_0} G_2, D_2, EvUsed'_0$ and $\mathcal{D}_{\neg EvUsed'_1, \emptyset} : G_1, D_{U1}, EvUsed_1 \xrightarrow{V}_{p_1} G_2, D'_2, EvUsed'_1$ where $EvUsed'_1 = EvUsed_1 \setminus \{e_1, \dots, e_m\}$, $D_{U1} = e'_1 \vee \dots \vee e'_l = \text{NonUnique}_{G_1, D_1}$, and $D_2 = D_1 \vee D'_2$. (The events e_1, \dots, e_m must be preserved by the second proof, hence we allow distinguishers in $\mathcal{D}_{\neg EvUsed'_1}$ to use these events. The events e_1, \dots, e_m will be introduced in the first proof, and the active queries in G_2 are also required to match so $D_2 = D_1 \vee D'_2$.) So for all attackers $\mathcal{A} \in \mathcal{BB}$, all evaluation contexts C acceptable for G_0 and G_2 without public variables that do not contain events in $EvUsed'_0$, and all distinguishers $D_0 \in \mathcal{D}_{\neg EvUsed'_0}$ that run in time at most t_{D_0} ,

$$\Pr^{\mathcal{A}}[C[G_0] : D_0] \leq \Pr^{\mathcal{A}}[C[G_2] : (D_0 \vee D_2) \wedge \neg \text{NonUnique}_{G_2, D_2}] + p_0(\mathcal{A}, C, t_{D_0})$$

and for all attackers $\mathcal{A} \in \mathcal{BB}$, for all evaluation contexts C acceptable for G_1 and G_2 without public variables that do not contain events in $EvUsed'_1$, and all distinguishers $D_0 \in \mathcal{D}_{\neg EvUsed'_1}$ that run in time at most t_{D_0} ,

$$\Pr^{\mathcal{A}}[C[G_1] : D_0 \vee D_{U1}] \leq \Pr^{\mathcal{A}}[C[G_2] : (D_0 \vee D'_2) \wedge \neg \text{NonUnique}_{G_2, D'_2}] + p_1(\mathcal{A}, C, t_{D_0})$$

In the last equation, we replace D_0 with $\neg D_0 \wedge \neg D_1$ for $D_0 \in \mathcal{D}_{\neg EvUsed'_1}$, yielding

$$\begin{aligned} \Pr^{\mathcal{A}}[C[G_1] : (\neg D_0 \wedge \neg D_1) \vee D_{U1}] \\ \leq \Pr^{\mathcal{A}}[C[G_2] : ((\neg D_0 \wedge \neg D_1) \vee D'_2) \wedge \neg \text{NonUnique}_{G_2, D'_2}] + p_1(\mathcal{A}, C, t_{D_0}) \end{aligned}$$

so

$$\begin{aligned} \Pr^{\mathcal{A}}[C[G_2] : ((D_0 \vee D_1) \wedge \neg D'_2) \vee \text{NonUnique}_{G_2, D'_2}] \\ \leq \Pr^{\mathcal{A}}[C[G_1] : (D_0 \vee D_1) \wedge \neg D_{U1}] + p_1(\mathcal{A}, C, t_{D_0}) \end{aligned}$$

Then we get

$$\begin{aligned} \Pr^{\mathcal{A}}[C[G_0] : D_0] &\leq \Pr^{\mathcal{A}}[C[G_2] : (D_0 \vee D_2) \wedge \neg \text{NonUnique}_{G_2, D_2}] + p_0(\mathcal{A}, C, t_{D_0}) \\ &\leq \Pr^{\mathcal{A}}[C[G_2] : (D_0 \vee D_1 \vee D'_2) \wedge \neg \text{NonUnique}_{G_2, D'_2}] + p_0(\mathcal{A}, C, t_{D_0}) \\ &\leq \Pr^{\mathcal{A}}[C[G_2] : ((D_0 \vee D_1) \wedge \neg D'_2) \vee \text{NonUnique}_{G_2, D'_2}] \\ &\quad + \Pr^{\mathcal{A}}[C[G_2] : D'_2 \wedge \neg \text{NonUnique}_{G_2, D'_2}] + p_0(\mathcal{A}, C, t_{D_0}) \\ &\leq \Pr^{\mathcal{A}}[C[G_1] : (D_0 \vee D_1) \wedge \neg D_{U1}] \\ &\quad + p_1(\mathcal{A}, C, t_{D_0}) + \Pr^{\mathcal{A}}[C[G_2] : D'_2 \wedge \neg \text{NonUnique}_{G_2, D'_2}] + p_0(\mathcal{A}, C, t_{D_0}) \\ &\leq \Pr^{\mathcal{A}}[C[G_1] : (D_0 \vee D_1) \wedge \neg \text{NonUnique}_{G_1, D_1}] \\ &\quad + p_1(\mathcal{A}, C, t_{D_0}) + \Pr^{\mathcal{A}}[C[G_2] : D'_2 \wedge \neg \text{NonUnique}_{G_2, D'_2}] + p_0(\mathcal{A}, C, t_{D_0}) \end{aligned}$$

so we get the desired result with

$$p(\mathcal{A}, C, t_{D_0}) = p_1(\mathcal{A}, C, t_{D_0}) + \Pr^{\mathcal{A}}[C[G_2] : D'_2 \wedge \neg \text{NonUnique}_{G_2, D'_2}] + p_0(\mathcal{A}, C, t_{D_0}).$$

Computational case (query_equiv with [computational] annotation) As in the decisional case, we want to prove $\mathcal{D}_{\neg EvUsed_1, \emptyset} : G_0, D_{\text{false}}, \emptyset \xrightarrow{V}_p G_1, D_1, EvUsed_1$ where G_0 contains no events, $D_1 = e_1 \vee \dots \vee e_m$, e_1, \dots, e_m are the Shoup events occurring in G_1 , e'_1, \dots, e'_l are the non-unique events occurring in G_1 , $EvUsed_1 = \{e_1, \dots, e_m, e'_1, \dots, e'_l\}$, $V = \emptyset$. Additionally, we want to show that the random values of G_0 and G_1 marked [unchanged] can be used in events (different from e_1, \dots, e_m since e_1, \dots, e_m have no arguments) in the game transformed using this assumption. That corresponds to adding oracles that execute the same arbitrary events using [unchanged] random values to both G_0 and G_1 . We write G'_0 and G'_1 for the games G_0 and G_1 respectively with additional events and show $\mathcal{D}_{\neg EvUsed_1, \emptyset} : G'_0, D_{\text{false}}, EvUsed_0 \xrightarrow{V}_p G'_1, D_1, EvUsed_1 \cup EvUsed_0$ where $EvUsed_0$ contains these additional events. These additional events can be observed by the adversary, so they are allowed in distinguishers in $\mathcal{D}_{\neg EvUsed_1}$.

Let $D_{U1} = e'_1 \vee \dots \vee e'_l = \text{NonUnique}_{G_1, D_1}$.

Let us write $O_0(\tilde{r}, \widetilde{args})$ (resp. $O_1(\tilde{r}, \widetilde{args})$) for the result of oracle O in game G_0 (resp. G_1) with randomness \tilde{r} and arguments \widetilde{args} .

In order to establish this property, we show that there exists a mapping ϕ of the randomness, such that if random value variable r has value v in G_0 , then it has value $\phi_r(v)$ in G_1 , ϕ_r is the identity when the variable r is marked [unchanged], ϕ_r preserves the probability distribution of variable r , and we define a game G_2 in which oracle O with randomness \tilde{r} and arguments \widetilde{args} returns

$$\begin{aligned} &\text{let } x_0 = O_0(\tilde{r}, \widetilde{args}) \text{ in} \\ &\text{let } x_1 = O_1(\phi(\tilde{r}), \widetilde{args}) \text{ in} \\ &\text{if } x_0 = x_1 \text{ then } x_0 \text{ else event_abort distinguisher} \end{aligned}$$

We bound

$$\begin{aligned} p(\mathcal{A}, C) &= \Pr^{\mathcal{A}}[C[G_2] : \text{distinguish} \vee \text{NonUnique}_{G_2, D_{\text{false}}}] \\ &= \Pr^{\mathcal{A}}[C[G_2] : \text{distinguish} \vee D_{U1}] \\ &= \text{Adv}_{G_2}(\mathcal{A}, C, \text{distinguish} \Rightarrow \text{false}, D_{U1}) \end{aligned}$$

From this bound, we infer the desired property.

We consider a game G_3 in which oracle O returns

$$\begin{aligned} &\text{let } x_0 = O_0(\tilde{r}, \widetilde{args}) \text{ in} \\ &\text{let } x_1 = O_1(\phi(\tilde{r}), \widetilde{args}) \text{ in} \\ &x_0 \end{aligned}$$

We define G'_2 as G_2 with the same additional events as in G'_0 and G'_1 , and G'_3 as G_3 with the same additional events as in G'_0 and G'_1 . The game G'_3 behaves as G'_0 except that it executes a Shoup event e_i or a non-unique event when G_1 does, so we have, for any $\mathcal{A} \in \mathcal{BB}$ and any evaluation context C acceptable for G'_0 and G'_3 without public variables, and any distinguisher D_0 ,

$$\Pr^{\mathcal{A}}[C[G'_0] : D_0] \leq \Pr^{\mathcal{A}}[C[G'_3] : D_0 \vee D_1 \vee D_{U1}]$$

Moreover, G'_3 behaves as G'_1 except when G'_2 executes event distinguish, that is, when G_2 executes event distinguish (the additional events introduced in G'_2 are not needed to evaluate the probability of distinguish since we do not consider their probability), so for all D ,

$$|\Pr^{\mathcal{A}}[C[G'_3] : D] - \Pr^{\mathcal{A}}[C[G'_1] : D]| \leq \Pr^{\mathcal{A}}[C[G_2] : \text{distinguish}] \quad (16)$$

Let C be any evaluation context acceptable for G'_0 and G'_2 without public variables and \mathcal{A} be an attacker in \mathcal{BB} . By renaming x_0 and x_1 to variables not in C , C is also acceptable for G'_0 and G'_3 without public variables. Let D_0 be any distinguisher. With $D = D_0 \vee D_1 \vee D_{U1}$ in (16), we obtain

$$\begin{aligned} \Pr^{\mathcal{A}}[C[G'_0] : D_0] &\leq \Pr^{\mathcal{A}}[C[G'_3] : D_0 \vee D_1 \vee D_{U1}] \\ &\leq \Pr^{\mathcal{A}}[C[G'_1] : D_0 \vee D_1 \vee D_{U1}] + \Pr^{\mathcal{A}}[C[G_2] : \text{distinguish}] \\ &\leq \Pr^{\mathcal{A}}[C[G'_1] : (D_0 \vee D_1) \wedge \neg \text{NonUnique}_{G_1, D_1}] + \Pr^{\mathcal{A}}[C[G'_1] : D_{U1}] \\ &\quad + \Pr^{\mathcal{A}}[C[G_2] : \text{distinguish}] \\ &\leq \Pr^{\mathcal{A}}[C[G'_1] : (D_0 \vee D_1) \wedge \neg \text{NonUnique}_{G_1, D_1}] + p(\mathcal{A}, C) \end{aligned}$$

since G'_1 and G_2 execute events D_{U1} in the same cases, and D_{U1} and distinguish are mutually exclusive. Therefore, we have $\mathcal{D}_{\neg \text{EvUsed}_1}, \emptyset : G'_0, D_{\text{false}}, \text{EvUsed}_0 \xrightarrow{V \rightarrow p} G'_1, D_1, \text{EvUsed}_1 \cup \text{EvUsed}_0$.

Currently, *CryptoVerif* can prove `query_equiv` only when the mapping ϕ is the identity for all variables. Other cases can be proved manually and used as assumptions in `equiv` statements.

3 Collecting True Facts

In this section, we consider only processes that satisfy Properties 3 and 4. We can assume without loss of generality that the adversary context also satisfies these properties: tables (`insert` and `get`) can be removed by encoding them using `find` by transformation `expand_tables` (Section 5.1.2) and variables defined in conditions of `find` can be renamed to have distinct names by transformation `auto_SArename` (Section 5.1.1).

Given a configuration $\text{Conf} = E, \sigma, N, \mathcal{T}, \mu\mathcal{E}v$ or $\text{Conf} = E, St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v$, we denote by E_{Conf} the environment E in configuration Conf . We denote by $E_{\text{Tr} \preceq \text{Conf}}$ the union of $E_{\text{Conf}'}$ for all configurations $\text{Conf}' \preceq_{\text{Tr}} \text{Conf}$ in Tr . It is a set of mappings $x[\tilde{a}] \mapsto b$. At this stage, it may include conflicting mappings $x[\tilde{a}] \mapsto b$ and $x[\tilde{a}] \mapsto b'$ with $b \neq b'$. We prove below (Lemma 29) that this situation never happens. The notation $E_{\text{Tr} \preceq \text{Conf}}$ is useful because the environment computed in the semantics does not keep the values of variables defined in conditions of `find` after these conditions are evaluated. Considering the union of all environments of previous configurations allows us to recover the values of these variables, and to use them in the facts that we collect. Given a configuration $\text{Conf} = E, \sigma, N, \mathcal{T}, \mu\mathcal{E}v$ or $\text{Conf} = E, (\sigma, P) :: St, \mathcal{Q}, Or, \mathcal{T}, \mu\mathcal{E}v$, we denote by σ_{Conf} the mapping sequence for replication indices σ in the configuration Conf and by $\mu\mathcal{E}v_{\text{Conf}}$ the sequence of events $\mu\mathcal{E}v$ in configuration Conf .

Let us define *Defined* as in Section 2.4.3, except that

$$\begin{aligned} \text{Defined}(\sigma, \text{find}[\text{unique?}] (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(\tilde{M}_j) \wedge M_j \text{ then } N_j) \text{ else } N) = \\ \left(\bigoplus_{j=1}^m \bigoplus_{\tilde{a} \leq \tilde{n}_j} \text{Defined}(\sigma[\tilde{i}_j \mapsto \tilde{a}], M_j) \right) \uplus \max \left(\max_{j=1}^m \left(\{\tilde{u}_j[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, N_j) \right), \text{Defined}(\sigma, N) \right) \\ \text{Defined}(\sigma, \text{find}[\text{unique?}] (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(\tilde{M}_j) \wedge M_j \text{ then } P_j) \text{ else } P) = \\ \left(\bigoplus_{j=1}^m \bigoplus_{\tilde{a} \leq \tilde{n}_j} \text{Defined}(\sigma[\tilde{i}_j \mapsto \tilde{a}], M_j) \right) \uplus \max \left(\max_{j=1}^m \left(\{\tilde{u}_j[\sigma(\tilde{i})]\} \uplus \text{Defined}(\sigma, P_j) \right), \text{Defined}(\sigma, P) \right) \end{aligned}$$

so that the variables defined in conditions of `find` are now considered as defined forever, and not

temporarily during the evaluation of the considered condition. We also define

$$\text{Defined}(Tr \preceq Conf) = \text{Dom}(E_{Tr \preceq Conf}) \uplus \text{Defined}^{\text{Fut}}(Conf).$$

Lemma 29 *Let Q_0 be a process that satisfies Properties 3 and 4. Let Tr be a trace of Q_0 for any $A \in \mathcal{BB}$ and $Conf$ be a configuration in the derivation of Tr . Then the following properties hold:*

1. $\text{Defined}(Tr \preceq Conf)$ does not contain duplicate elements.
2. Each variable is defined at most once for each value of its array indices in Tr .
3. $E_{Tr \preceq Conf}$ contains at most one binding for each $x[\tilde{a}]$.

Proof sketch The proof is similar to the proof of Lemma 9. We first show as in Lemma 9 that, for all program points μ in Q_0 , if $\text{Dom}(\sigma) = I_\mu$ are the current replication indices at μ and the process or term Q at μ satisfies Invariant 1, then all elements of $\text{Defined}(\sigma, Q)$ are of the form $x[\tilde{a}]$ where $x \in \text{vardef}(Q)$ and $\text{Im}(\sigma)$ is a prefix of \tilde{a} .

Next, we show that, for all program points μ , if $\text{Dom}(\sigma) = I_\mu$ are the current replication indices at μ and the process or term Q at μ satisfies Invariant 1, then $\text{Defined}(\sigma, Q)$ does not contain duplicate elements. The proof proceeds by induction on Q . All multiset unions in the computation of $\text{Defined}(\sigma, Q)$ are disjoint unions by the property above, because either they use different extensions of σ (cases of replication and of conditions of find) or they use disjoint variable definitions or subprocesses or subterms in the same branch of find or if, which must define different variables by Invariant 1 and by Property 4.

We show by induction on the derivations that, if $Conf \xrightarrow{p}_t Conf'$, then $\text{Defined}(Tr \preceq Conf) \supseteq \text{Defined}(Tr \preceq Conf')$ and for all semantic configurations $Conf''$ in the derivation of $Conf \xrightarrow{p}_t Conf'$, $\text{Defined}(Tr \preceq Conf) \supseteq \text{Defined}(Tr \preceq Conf'')$, and similarly with \rightsquigarrow instead of \xrightarrow{p}_t .

The first result follows: since Q_0 satisfies Invariant 1, $\text{Defined}(\sigma_0, Q_0)$ does not contain duplicate elements, where σ_0 is the empty mapping sequence. Let $Conf_0 = \{(\sigma_0, Q_0)\}$, $(\text{fo}(Q_0), \emptyset)$, $Conf_1 = \text{reduce}(\{(\sigma_0, Q_0)\}, (\text{fo}(Q_0), \emptyset))$, $Conf_2 = \text{initConfig}(Q_0, \mathcal{A})$, and $Conf_3$ be any other configuration of Tr . Then $\text{Defined}(Tr \preceq Conf_0)$, $\text{Defined}(Tr \preceq Conf_1)$, $\text{Defined}(Tr \preceq Conf_2)$, and therefore $\text{Defined}(Tr \preceq Conf_3)$ do not contain duplicate elements.

Let us prove the second result. In order to derive a contradiction, assume that two transitions $Conf_1 \xrightarrow{p_1}_{t_1} Conf'_1$ and $Conf_2 \xrightarrow{p_2}_{t_2} Conf'_2$ inside Tr define the same variable $x[\tilde{a}]$.

- First case: one transition happens before the other, for instance $Conf'_1 \preceq_{Tr} Conf_2$. (The case $Conf'_2 \preceq_{Tr} Conf_1$ is symmetric.) Since $Conf_1 \xrightarrow{p_1}_{t_1} Conf'_1$ defines $x[\tilde{a}]$, we have $x[\tilde{a}] \in \text{Dom}(E_{Conf'_1})$, so $x[\tilde{a}] \in \text{Dom}(E_{Tr \preceq Conf_2})$. Moreover, since $Conf_2 \xrightarrow{p_2}_{t_2} Conf'_2$ defines $x[\tilde{a}]$, we have $x[\tilde{a}] \in \text{Defined}^{\text{Fut}}(Conf_2)$, by inspecting all rules that add elements to the environment. Therefore $\text{Defined}(Tr \preceq Conf_2) = \text{Dom}(E_{Tr \preceq Conf_2}) \uplus \text{Defined}^{\text{Fut}}(Conf_2)$ contains twice $x[\tilde{a}]$. Contradiction.
- Second case: the transitions cannot be ordered. By definition of \preceq_{Tr} , this can happen only when a semantic rule uses several derivations for its assumptions, which happens only in rules for find. (Recall that get is excluded by Property 3.) Therefore, there exists k_1 and k_2 such that $Conf_1 \xrightarrow{p_1}_{t_1} Conf'_1$ is in the derivation of $Conf_{0,k} = E, \sigma[\tilde{i}_{j_k} \mapsto \tilde{a}_k], D_{j_k} \wedge M_{j_k}, \mathcal{T}, \mu \mathcal{E}v \xrightarrow{p_k}_{t_k} Conf'_{0,k} = E_k, \sigma_k, r_k, \mathcal{T}, \mu \mathcal{E}v$ with $v_k = (j_k, \tilde{a}_k)$ for $k = k_1$ and $Conf_2 \xrightarrow{p_2}_{t_2} Conf'_2$ is in that derivation for $k = k_2$, with $k_1 \neq k_2$. We have $x[\tilde{a}] \in \text{Defined}(Tr \preceq Conf'_{0,k_1}) \subseteq \text{Defined}(Tr \preceq Conf_{0,k_1}) = \text{Dom}(E_{Tr \preceq Conf_{0,k_1}}) \cup$

$Defined(\sigma[\widetilde{i}_{j_{k_1}} \mapsto \widetilde{a}_{k_1}], M_{j_{k_1}})$. Moreover, $Conf_{0,k} \preceq Conf_1$, so $\text{Dom}(E_{Tr \preceq Conf_{0,k_1}}) \subseteq \text{Dom}(E_{Tr \preceq Conf_1})$. Since $Conf_1 \xrightarrow{P_1}_{t_1} Conf'_1$ defines $x[\widetilde{a}]$, we have $x[\widetilde{a}] \in Defined^{\text{Fut}}(Conf_1)$, by inspecting all rules that add elements to the environment. Since $Defined(Tr \preceq Conf_1) = \text{Dom}(E_{Tr \preceq Conf_1}) \uplus Defined^{\text{Fut}}(Conf_1)$ does not contain duplicate elements, we have $x[\widetilde{a}] \notin \text{Dom}(E_{Tr \preceq Conf_1})$, so $x[\widetilde{a}] \notin \text{Dom}(E_{Tr \preceq Conf_{0,k_1}})$. Hence $x[\widetilde{a}] \in Defined(\sigma[\widetilde{i}_{j_{k_1}} \mapsto \widetilde{a}_{k_1}], M_{j_{k_1}})$. Similarly, $x[\widetilde{a}] \in Defined(\sigma[\widetilde{i}_{j_{k_2}} \mapsto \widetilde{a}_{k_2}], M_{j_{k_2}})$. Let us show that the sets $Defined(\sigma[\widetilde{i}_{j_{k_1}} \mapsto \widetilde{a}_{k_1}], M_{j_{k_1}})$ and $Defined(\sigma[\widetilde{i}_{j_{k_2}} \mapsto \widetilde{a}_{k_2}], M_{j_{k_2}})$ are disjoint. We have $v_{k_1} \neq v_{k_2}$, so either $j_{k_1} \neq j_{k_2}$ and in this case these sets are disjoint because $M_{j_{k_1}}$ and $M_{j_{k_2}}$ define different variables by Property 4, or $j_{k_1} = j_{k_2}$ and $\widetilde{a}_{k_1} \neq \widetilde{a}_{k_2}$ and in this case these sets are disjoint because they use different extensions of σ . Since these sets are disjoint, they cannot both contain $x[\widetilde{a}]$. Contradiction.

The last result is an immediate consequence of the second one. \square

Lemma 30 *Let Q_0 be a process that satisfies Properties 3 and 4. Let Tr be a trace of Q_0 for any $\mathcal{A} \in \mathcal{BB}$ and $Conf$ be a configuration in the derivation of Tr . We have $E_{Tr \preceq Conf} = E_{Conf}[x[\widetilde{a}] \mapsto b$ for some variables x defined in a condition of find and some indices \widetilde{a} and values b].*

Proof sketch By induction on the derivation. \square

The previous lemma shows that the only difference between E_{Conf} and $E_{Tr \preceq Conf}$ is that variables defined in conditions of find are added to $E_{Tr \preceq Conf}$. These variables have no array accesses, so they do not appear in defined conditions of find. Therefore, these defined conditions yield the same result whether they are evaluated in E_{Conf} or in $E_{Tr \preceq Conf}$.

We consider the following facts:

- The boolean term M means that M evaluates to true.
- $\text{defined}(M)$ means that M is defined (all array accesses in M are defined).
- $\text{event}(e(\widetilde{M}))$ means that event $e(\widetilde{M})$ has been executed.
- $\text{event}(e(\widetilde{M}))@_\tau$ means that event $e(\widetilde{M})$ has been executed at step τ (index in the sequence of events $\mu \mathcal{E} v$).
- $M_1 : \text{event}(e(\widetilde{M}))$ means that event $e(\widetilde{M})$ has been executed with pair (program point, replication indices) equal to M_1 .
- $M_1 : \text{event}(e(\widetilde{M}))@_\tau$ means that event $e(\widetilde{M})$ has been executed at step τ with pair (program point, replication indices) equal to M_1 .
- $\text{programpoint}(\mu, \widetilde{M})$ means that program point μ has been executed with replication indices equal to \widetilde{M} .
- $\text{programpoint}(\mathcal{S}_1, \widetilde{M}_1) \preceq \dots \preceq \text{programpoint}(\mathcal{S}_m, \widetilde{M}_m)$ means that, for $j \leq m$, some program point $\mu_j \in \mathcal{S}_j$ has been executed with replication indices equal to \widetilde{M}_j , and furthermore these program points have been executed in the order of increasing j .
- $\text{lastdefprogrampoint}(\mu, \widetilde{M})$ means that program point μ has been executed with replication indices equal to \widetilde{M} and the values of variables and replication indices are unchanged since that program point (that is, no variable definition nor oracle call, return, or yield that changes the replication indices was executed since that program point).

Given an environment E mapping process variables to their values, an environment ρ mapping replication indices and non-process variables of the formula φ to their values, and a sequence of events $\mu\mathcal{E}v$, we define $E, \rho, \mu\mathcal{E}v \vdash \varphi$, meaning that $E, \rho, \mu\mathcal{E}v$ satisfy φ , as follows:

- $E, \rho, \mu\mathcal{E}v \vdash M$ if and only if $E, \rho, M \Downarrow \text{true}$.
- $E, \rho, \mu\mathcal{E}v \vdash \text{defined}(M)$ if and only if $E, \rho, M \Downarrow a$ for some a .
- $E, \rho, \mu\mathcal{E}v \vdash \text{event}(e(\widetilde{M}))$ if and only if $E, \rho, \widetilde{M} \Downarrow \widetilde{a}$ and $(\mu, \widetilde{a}') : e(\widetilde{a}) \in \mu\mathcal{E}v$ for some μ and \widetilde{a}' .
- $E, \rho, \mu\mathcal{E}v \vdash \text{event}(e(\widetilde{M}))@M_0$ if and only if $E, \rho, \widetilde{M} \Downarrow \widetilde{a}$, $E, \rho, M_0 \Downarrow a_0$, and $\mu\mathcal{E}v(a_0) = (\mu, \widetilde{a}') : e(\widetilde{a})$ for some μ and \widetilde{a}' .
- $E, \rho, \mu\mathcal{E}v \vdash M_1 : \text{event}(e(\widetilde{M}))$ if and only if $E, \rho, \widetilde{M} \Downarrow \widetilde{a}$, $E, \rho, M_1 \Downarrow (\mu, \widetilde{a}')$ and $(\mu, \widetilde{a}') : e(\widetilde{a}) \in \mu\mathcal{E}v$.
- $E, \rho, \mu\mathcal{E}v \vdash M_1 : \text{event}(e(\widetilde{M}))@M_0$ if and only if $E, \rho, \widetilde{M} \Downarrow \widetilde{a}$, $E, \rho, M_0 \Downarrow a_0$, $E, \rho, M_1 \Downarrow (\mu, \widetilde{a}')$, and $\mu\mathcal{E}v(a_0) = (\mu, \widetilde{a}') : e(\widetilde{a})$.

Logical connectives are defined as usual. When φ does not contain events, $\mu\mathcal{E}v$ can be omitted, writing $E, \rho \vdash \varphi$.

Let \mathcal{A} be an attacker in \mathcal{BB} and Tr be a trace of Q_0 for attacker \mathcal{A} . Let $Conf = E, \sigma, N, \mathcal{T}, \mu\mathcal{E}v$ or $Conf = E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$ be a configuration that occurs in the derivation of Tr . We define $Tr \preceq Conf, \rho \vdash \varphi$, meaning that the prefix of Tr until $Conf$ satisfies the formula φ with environment ρ (giving values of non-process variables of φ) as follows:

- $Tr \preceq Conf, \rho \vdash F$ if and only if $E_{Tr \preceq Conf}, \sigma_{Conf} \cup \rho, \mu\mathcal{E}v_{Conf} \vdash F$, when F is a term M , a defined fact $\text{defined}(M)$, or an event $\text{event}(e(\widetilde{M}))$, $\text{event}(e(\widetilde{M}))@M_0$, $M_1 : \text{event}(e(\widetilde{M}))$, or $M_1 : \text{event}(e(\widetilde{M}))@M_0$.
- $Tr \preceq Conf, \rho \vdash \text{programpoint}(\mathcal{S}_1, \widetilde{M}_1) \preceq \dots \preceq \text{programpoint}(\mathcal{S}_m, \widetilde{M}_m)$ if and only if, for all $j \in \{1, \dots, m\}$, there exists $Conf_j$ at program point $\mu_j \in \mathcal{S}_j$ in Tr such that $E_{Tr \preceq Conf}, \sigma_{Conf} \cup \rho, \widetilde{M}_j \Downarrow \text{Im}(\sigma_{Conf_j})$ and $Conf_1 \preceq_{Tr} \dots \preceq_{Tr} Conf_m \preceq_{Tr} Conf$.
The fact $\text{programpoint}(\mu, \widetilde{M})$ is actually a particular case of $\text{programpoint}(\mathcal{S}_1, \widetilde{M}_1) \preceq \dots \preceq \text{programpoint}(\mathcal{S}_m, \widetilde{M}_m)$ with $m = 1$ and $\mathcal{S}_1 = \{\mu\}$. Hence, we have $Tr \preceq Conf, \rho \vdash \text{programpoint}(\mu, \widetilde{M})$ if and only if there is a configuration $Conf'$ at program point μ in Tr such that $Conf' \preceq_{Tr} Conf$ and $E_{Tr \preceq Conf}, \sigma_{Conf} \cup \rho, \widetilde{M} \Downarrow \text{Im}(\sigma_{Conf'})$.
- $Tr \preceq Conf, \rho \vdash \text{lastdefprogrampoint}(\mu, \widetilde{M})$ if and only if there is a configuration $Conf'$ at program point μ in Tr such that $Conf' \preceq_{Tr} Conf$, $E_{Tr \preceq Conf'} = E_{Tr \preceq Conf}$, $\sigma_{Conf'} = \sigma_{Conf}$, and $E_{Tr \preceq Conf}, \sigma_{Conf} \cup \rho, \widetilde{M} \Downarrow \text{Im}(\sigma_{Conf'})$.

Logical connectives are defined as usual. Most facts are evaluated in the environment $E_{Tr \preceq Conf}$ and the mapping sequence σ_{Conf} . Events are evaluated using the sequence of events $\mu\mathcal{E}v_{Conf}$ in $Conf$, but correspond to an execution of the event at some point before $Conf$ in the trace.

We define that a trace Tr satisfies a logical formula φ with environment ρ (giving values of non-process variables of φ), denoted $Tr, \rho \vdash \varphi$ as $Tr \preceq Conf, \rho \vdash \varphi$, where Tr ends with $Conf$. Along the same line, we define $E_{Tr} = E_{Tr \preceq Conf}$ and $\mu\mathcal{E}v_{Tr} = \mu\mathcal{E}v_{Conf}$ where Tr ends with $Conf$.

When the formula φ does not contain free non-process variables, we may write $Tr \vdash \varphi$ instead of $Tr, \rho \vdash \varphi$ since the environment ρ is useless. When \mathcal{F} is a set of formulas (in particular, of

facts), we write $\bigwedge \mathcal{F}$ for $\bigwedge_{F \in \mathcal{F}} F$ and $\bigvee \mathcal{F}$ for $\bigvee_{F \in \mathcal{F}} F$. We also write $Tr, \rho \vdash \mathcal{F}$ when for all $\varphi \in \mathcal{F}$, $Tr, \rho \vdash \varphi$. This is equivalent to $Tr, \rho \vdash \bigwedge \mathcal{F}$. We use similar notations for prefixes $Tr \preceq Conf$ instead of traces Tr .

Additionally, we define the following facts:

- $\llbracket \text{elsefind}((i_1 \leq n_1, \dots, i_m \leq n_m), (M_1, \dots, M_l), M) \rrbracket = \forall i_1 \in [1, n_1], \dots, \forall i_m \in [1, n_m], \neg(\text{defined}(M_1) \wedge \dots \wedge \text{defined}(M_l) \wedge M)$.
- $\llbracket \text{elselet}(\tilde{x} : \tilde{T}, N, M) \rrbracket = \forall \tilde{x} \in \tilde{T}, N \neq M$.

Extension: *The following description is based on the first version of the collection of known facts. Many improvements and extensions have been implemented. In particular, the following description often uses P instead of μ for program points; the syntax does not allow $\stackrel{R}{\leftarrow}$, if, let, find in terms; events are not handled.*

3.1 User-defined Rewrite Rules

The user can give two kinds of information:

- claims of the form $\forall x_1 : T_1, \dots, \forall x_m : T_m, M$ which mean that for all environments E , if for all $j \leq m$, $E(x_j) \in T_j$, then $E, M \Downarrow \text{true}$.

Such claims must be well-typed, that is, $\{x_1 \mapsto T_1, \dots, x_m \mapsto T_m\} \vdash M : \text{bool}$.

They are translated into rewrite rules as follows:

- If M is of the form $M_1 = M_2$ and $\text{vardef}(M_2) \subseteq \text{vardef}(M_1)$, we generate the rewrite rule $\forall x_1 : T_1, \dots, \forall x_m : T_m, M_1 \rightarrow M_2$.
- If M is of the form $M_1 \neq M_2$, we generate the rewrite rules $\forall x_1 : T_1, \dots, \forall x_m : T_m, (M_1 = M_2) \rightarrow \text{false}$, $\forall x_1 : T_1, \dots, \forall x_m : T_m, (M_1 \neq M_2) \rightarrow \text{true}$. (Such rules are used for instance to express that different constants are different.)
- Otherwise, we generate the rewrite rule $\forall x_1 : T_1, \dots, \forall x_m : T_m, M \rightarrow \text{true}$.

The term M reduces into M' by the rewrite rule $\forall x_1 : T_1, \dots, \forall x_m : T_m, M_1 \rightarrow M_2$ if and only if $M = C[\sigma M_1]$, $M' = C[\sigma M_2]$, where C is a term context and σ is a substitution that maps x_j to any term of type T_j for all $j \leq m$.

- claims of the form $y_1 \stackrel{R}{\leftarrow} T'_1, \dots, y_l \stackrel{R}{\leftarrow} T'_l, \forall x_1 : T_1, \dots, \forall x_m : T_m, M_1 \approx_p M_2$ with $\text{vardef}(M_2) \subseteq \text{vardef}(M_1)$. Informally, these claims mean that M_1 and M_2 evaluate to the same bitstring except in cases of probability at most p , provided that y_1, \dots, y_l are chosen randomly with uniform probability and independently among T'_1, \dots, T'_l respectively, and that x_1, \dots, x_m are of type T_1, \dots, T_m . (x_1, \dots, x_m may depend on y_1, \dots, y_l .) Formally, these claims are defined as:

$$\begin{aligned} & \Pr[E(y_1) \stackrel{R}{\leftarrow} T'_1; \dots; E(y_l) \stackrel{R}{\leftarrow} T'_l; \\ & \quad (E(x_1), \dots, E(x_m)) \leftarrow \mathcal{A}(E(y_1), \dots, E(y_l)); \\ & \quad E, M_1 \Downarrow a; E, M_2 \Downarrow a' : a \neq a'] \leq p(\mathcal{A}) \end{aligned}$$

where \mathcal{A} is an attacker in \mathcal{BB} .

The above claim must be well-typed, that is, $\{x_1 \mapsto T_1, \dots, x_m \mapsto T_m, y_1 \mapsto T'_1, \dots, y_l \mapsto T'_l\} \vdash M_1 = M_2$.

This claim is translated into the rewrite rule $y_1 \stackrel{R}{\leftarrow} T'_1, \dots, y_l \stackrel{R}{\leftarrow} T'_l, \forall x_1 : T_1, \dots, \forall x_m : T_m, M_1 \rightarrow M_2$.

The prover has built-in rewrite rules for defining boolean functions:

$$\begin{aligned}
& \neg \text{true} \rightarrow \text{false} & \neg \text{false} \rightarrow \text{true} & \forall x : \text{bool}, \neg(\neg x) \rightarrow x \\
& \forall x : T, \forall y : T, \neg(x = y) \rightarrow x \neq y \\
& \forall x : T, \forall y : T, \neg(x \neq y) \rightarrow x = y \\
& \forall x : T, x = x \rightarrow \text{true} & \forall x : T, x \neq x \rightarrow \text{false} \\
& \forall x : \text{bool}, \forall y : \text{bool}, \neg(x \wedge y) \rightarrow (\neg x) \vee (\neg y) \\
& \forall x : \text{bool}, \forall y : \text{bool}, \neg(x \vee y) \rightarrow (\neg x) \wedge (\neg y) \\
& \forall x : \text{bool}, x \wedge \text{true} \rightarrow x & \forall x : \text{bool}, x \wedge \text{false} \rightarrow \text{false} \\
& \forall x : \text{bool}, x \vee \text{true} \rightarrow \text{true} & \forall x : \text{bool}, x \vee \text{false} \rightarrow x \\
& \forall x : T, \forall y : T, \text{if_fun}(\text{true}, x, y) \rightarrow x & \forall x : T, \forall y : T, \text{if_fun}(\text{false}, x, y) \rightarrow y \\
& \forall x : \text{bool}, \forall y : T, \text{if_fun}(x, y, y) \rightarrow y \\
& \forall x_1 : T_1, \dots, \forall x_m : T_m, \forall x : \text{bool}, \forall y : T_k, \forall z : T_k, \\
& \quad f(x_1, \dots, x_{k-1}, \text{if_fun}(x, y, z), x_{k+1}, \dots, x_m) \rightarrow \\
& \quad \text{if_fun}(x, f(x_1, \dots, x_{k-1}, y, x_{k+1}, \dots, x_m), f(x_1, \dots, x_{k-1}, z, x_{k+1}, \dots, x_m)) \\
& \quad \text{when } f : T_1 \times \dots \times T_m \rightarrow T \text{ has option } \mathbf{autoSwapIf}
\end{aligned}$$

The prover also has support for commutative function symbols, that is, binary function symbols $f : T \times T \rightarrow T'$ such that for all $x, y \in T$, $f(x, y) = f(y, x)$. For such symbols, all equality and matching tests are performed modulo commutativity. The functions \wedge , \vee , $=$, and \neq are commutative. So, for instance, the rewrite rules above may also be used to rewrite $\text{true} \wedge M$ into M , $\text{false} \wedge M$ into false , $\text{true} \vee M$ into true , and $\text{false} \vee M$ into M . Used-defined functions may also be declared commutative; xor is an example of such a commutative function.

Example 4 For example, considering MAC and encryption schemes as in Definitions 2 and 3 respectively, we have:

$$\begin{aligned}
& \forall k : T_{mk}, \forall m : \text{bitstring}, \\
& \quad \text{verify}(m, k, \text{mac}(m, k)) = \text{true}
\end{aligned} \tag{mac}$$

$$\begin{aligned}
& \forall m : \text{bitstring}; \forall k : T_k, \forall r : T_r, \\
& \quad \text{dec}(\text{enc}(m, k, r), k) = i_{\perp}(m)
\end{aligned} \tag{enc}$$

We express the poly-injectivity of the function k2b of Example 1 by

$$\begin{aligned}
& \forall x : T_k, \forall y : T_k, (\text{k2b}(x) = \text{k2b}(y)) = (x = y) \\
& \forall x : T_k, \text{k2b}^{-1}(\text{k2b}(x)) = x
\end{aligned} \tag{k2b}$$

where k2b^{-1} is a function symbol that denotes the inverse of k2b. We have similar formulas for i_{\perp} .

3.2 Collecting True Facts from a Game

We use *facts* to represent properties that hold at certain program points in processes. We consider two kinds of facts: $\text{defined}(M)$ means that M is defined, and a term M means that M is true (the boolean term M evaluates to true). In this section, we show how to compute a set of facts \mathcal{F}_{μ} that are guaranteed to hold at the program point μ of the game.

The function `collectFacts` collects facts that hold at each program point of the game. More precisely, for each program point μ in the game, it computes a set \mathcal{F}_μ of facts that hold at that program point. The function `collectFacts` also computes a set \mathcal{D} containing pairs $(x[\tilde{i}], \mu)$ where $x[\tilde{i}]$ has been defined immediately above program point μ . (If there are several definitions of x , there is one such pair for each definition of x .) Finally, for oracle bodies P , `collectFacts`(P) returns a set of facts that will hold when the return is executed and stores this set in $\mathcal{F}_P^{\text{Fut}}$. (The superscript `Fut` stands for *future*, since these facts do not hold yet at P , but will hold in the future.)

The function `collectFacts` is defined in Figure 18. It is initially called by `collectFacts`(Q_0). It takes into account that $x[\tilde{i}]$ may be defined by an oracle definition, a random choice, a let, or a find and updates \mathcal{D} accordingly. Furthermore, when we execute `let` $x[\tilde{i}] : T = M$ in P' , $x[\tilde{i}] = M$ holds in P' and $x[\tilde{i}]$ is defined in P' . When we execute `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` `defined`(M_{j1}, \dots, M_{jl_j}) $\wedge M_j$ `then` P_j) `else` P' , $u_{j1}[\tilde{i}], \dots, u_{jm_j}[\tilde{i}], \sigma M_{j1}, \dots, \sigma M_{jl_j}$ are defined in P_j , σM_j holds in P_j where $\sigma = \{u_{j1}[\tilde{i}]/i_{j1}, \dots, u_{jm_j}[\tilde{i}]/i_{jm_j}\}$, and $\neg M_j$ holds in P' when $m_j = l_j = 0$.

After calling `collectFacts`(Q_0), we complete the computed sets \mathcal{F}_P (where P may be an oracle definition or body) by adding facts that come from processes above P :

$$\mathcal{F}_P \leftarrow \mathcal{F}_P \cup \mathcal{F}_{P'} \text{ if } P \text{ is immediately under } P'$$

We also add facts that we can deduce from facts `defined`(M). Precisely, if `defined`(M) $\in \mathcal{F}_P$ and $x[M_1, \dots, M_m]$ is a subterm of M , then we take into account facts that are known to be true at the definitions of x by adding them to \mathcal{F}_P as follows:

$$\mathcal{F}_P \leftarrow \mathcal{F}_P \cup \left(\bigcap_{(x[i_1, \dots, i_m], P') \in \mathcal{D}} \begin{cases} \sigma(\mathcal{F}_{P'} \cup (\mathcal{F}_{P'}^{\text{Fut}} \cap \mathcal{F}_P)) & \text{if } P \text{ is under } P' \\ \sigma(\mathcal{F}_{P'} \cup \mathcal{F}_{P'}^{\text{Fut}}) & \text{otherwise} \end{cases} \right)$$

where $\sigma = \{M_1/i_1, \dots, M_m/i_m\}$. Indeed, if `defined`(M) $\in \mathcal{F}_P$ and $x[M_1, \dots, M_m]$ is a subterm of M , then $x[M_1, \dots, M_m]$ is defined at P , so some definition of $x[M_1, \dots, M_m]$, immediately above the process P' , must have been executed before reaching P , so the facts that hold at P' also hold at P , with a suitable substitution of indices: we have $\sigma\mathcal{F}_{P'}$, that is, $\mathcal{F}_{P'}\{M_1/i_1, \dots, M_m/i_m\}$. Moreover, if the occurrence P is not syntactically under the occurrence P' , then the code of P' must have been executed until the next return before yielding control to some other code and reaching P , so in fact $\sigma(\mathcal{F}_{P'} \cup \mathcal{F}_{P'}^{\text{Fut}})$ hold. If P is syntactically under P' , it is possible that the code of P' has been executed until reaching P instead of until reaching the next return or yield, so we have only $\sigma(\mathcal{F}_{P'} \cup (\mathcal{F}_{P'}^{\text{Fut}} \cap \mathcal{F}_P))$. If there are several definitions of x , we do not know which one has been executed, so we only add to \mathcal{F}_P the facts that hold in all cases, by taking the intersection on all definitions of x .

This operation may add new `defined` facts to \mathcal{F}_P , so it is executed until a fixpoint is reached, except that, in order to avoid infinite loops, we do not execute this step for definitions `defined`(M) in which M contains nested occurrences of the same symbol (such as $x[\dots x[\dots]]$).

We also consider an additional fact that serves in expressing that the condition part of a `find` failed. Precisely, the fact `elsefind`(($i_1 \leq n_1, \dots, i_m \leq n_m$), (M_1, \dots, M_l), M) means that for all $i_1 \in [1, n_1], \dots, i_m \in [1, n_m]$, the terms M_1, \dots, M_l are not all `defined` or M is false. The function `collectElseFind` described in Figure 19 collects `elsefind` facts that hold at each occurrence. The function `collectElseFind`(P, \mathcal{F}) is called when \mathcal{F} is the set of true `elsefind` facts at occurrence P . It sets the value of $\mathcal{F}_P^{\text{ElseFind}}$ to \mathcal{F} .

```

collectFacts(Q) =
  if Q = Q1 | Q2 then collectFacts(Q1); collectFacts(Q2)
  if Q = foreach i ≤ n do Q' then collectFacts(Q')
  if Q = newOracle O; Q' then collectFacts(Q')
  if Q = O[ $\tilde{i}$ ](x[ $\tilde{i}$ ] : T) := P then
     $\mathcal{F}_P = \{\text{defined}(x[\tilde{i}])\}; \mathcal{F}_P^{\text{Fut}} = \text{collectFacts}(P)$ 
     $\mathcal{D} = \mathcal{D} \cup \{(x[\tilde{i}], P)\}$ 
collectFacts(P) =
  if P = return (N); Q then collectFacts(Q); return  $\emptyset$ 
  if P = x[ $\tilde{i}$ ]  $\stackrel{R}{\leftarrow}$  T; P' then
     $\mathcal{F}_{P'} = \{\text{defined}(x[\tilde{i}])\}; \mathcal{F}_{P'}^{\text{Fut}} = \text{collectFacts}(P')$ 
     $\mathcal{D} = \mathcal{D} \cup \{(x[\tilde{i}], P')\}; \text{return } \mathcal{F}_{P'} \cup \mathcal{F}_{P'}^{\text{Fut}}$ 
  if P = let x[ $\tilde{i}$ ] : T = M in P' then
     $\mathcal{F}_{P'} = \{\text{defined}(x[\tilde{i}]), x[\tilde{i}] = M\}$ 
     $\mathcal{F}_{P'}^{\text{Fut}} = \text{collectFacts}(P')$ 
     $\mathcal{D} = \mathcal{D} \cup \{(x[\tilde{i}], P')\}; \text{return } \mathcal{F}_{P'} \cup \mathcal{F}_{P'}^{\text{Fut}}$ 
  if P = 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 defined( $M_{j1}, \dots, M_{jl_j}$ )  $\wedge$   $M_j$  then  $P_j$ ) else P'
  then
    for each j ≤ m,
       $\mathcal{F}_{P_j} = \{\text{defined}(u_{j1}[\tilde{i}]), \dots, \text{defined}(u_{jm_j}[\tilde{i}]), \text{defined}(\sigma M_{j1}), \dots, \text{defined}(\sigma M_{jl_j}), \sigma M_j\}$ 
      where  $\sigma = \{u_{j1}[\tilde{i}]/i_{j1}, \dots, u_{jm_j}[\tilde{i}]/i_{jm_j}\}$ 
       $\mathcal{F}_{P_j}^{\text{Fut}} = \text{collectFacts}(P_j);$ 
       $\mathcal{D} = \mathcal{D} \cup \{(u_{j1}[\tilde{i}], P_j), \dots, (u_{jm_j}[\tilde{i}], P_j)\}$ 
     $\mathcal{F}_{P'} = \{\neg M_j \mid m_j = l_j = 0\}; \mathcal{F}_{P'}^{\text{Fut}} = \text{collectFacts}(P')$ 
  return ( $\mathcal{F}_{P'} \cup \mathcal{F}_{P'}^{\text{Fut}}$ )  $\cap \bigcap_{j=1}^m (\mathcal{F}_{P_j} \cup \mathcal{F}_{P_j}^{\text{Fut}})$ 

```

Figure 18: The function collectFacts

```

collectElseFind(Q) =
  if Q = Q1 | Q2 then collectElseFind(Q1); collectElseFind(Q2)
  if Q = foreach i ≤ n do Q' then collectElseFind(Q')
  if Q = newOracle O; Q' then collectElseFind(Q')
  if Q = O[ $\tilde{i}$ ](x[ $\tilde{i}$ ] : T) := P then collectElseFind(P, ∅)

collectElseFind(P,  $\mathcal{F}$ ) =
   $\mathcal{F}_P^{\text{ElseFind}} = \mathcal{F}$ 
  if P = return (N); Q then collectElseFind(Q)
  if P = x[ $\tilde{i}$ ]  $\stackrel{R}{\leftarrow}$  T; P' or P = let x[ $\tilde{i}$ ] : T = M in P' then
     $\mathcal{F}' = \{ \text{elsefind}((\tilde{i}' \leq \tilde{n}), (M'_1, \dots, M'_l), M') \in \mathcal{F} \mid x \text{ does not occur in } M'_1, \dots, M'_l \}$ 
    collectElseFind(P',  $\mathcal{F}'$ )
  if P = 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 defined( $M_{j1}, \dots, M_{jl_j} \wedge M_j$  then  $P_j$ ) else P')
  then
    for each j ≤ m,
       $\mathcal{F}'_j = \{ \text{elsefind}((\tilde{i}' \leq \tilde{n}), (M'_1, \dots, M'_l), M') \in \mathcal{F}$ 
        |  $u_{j1}, \dots, u_{jm_j}$  do not occur in  $M'_1, \dots, M'_l \}$ 
      collectElseFind(Pj,  $\mathcal{F}'_j$ )
       $\sigma_j = \{ i_1/i_{j1}, \dots, i_{m_j}/i_{jm_j} \}$ 
      collectElseFind(P',  $\mathcal{F} \cup$ 
        {  $\text{elsefind}((i_1 \leq n_{j1}, \dots, i_{m_j} \leq n_{jm_j}), \sigma_j(M_{j1}, \dots, M_{jl_j}), \sigma_j M_j) \mid j \in \{1, \dots, m\}$  })

```

Figure 19: The function collectElseFind

- In the case of random choices, assignments, and then branches of `find`, it takes into account that a variable x or u_{j1}, \dots, u_{jm_j} is newly defined. Hence `elsefind` facts that claim that one of these variables is not defined are removed.
- In the case of the `else` branch of a `find`, it adds the new `elsefind` facts that hold when the conditions of the `find` fail. These conditions express that each `then` branch of the `find` fails by a `elsefind` fact. To construct this fact, we replace (by applying σ_j) the indices i_{j1}, \dots, i_{jm_j} with fresh indices i_1, \dots, i_{m_j} , respectively.
- In the case of a `return`, any code may be executed before the oracle definition under it, so any variable may be defined by that code, and all `elsefind` facts are removed. That is why the function `collectElseFind` for oracle definitions has no \mathcal{F} argument (this argument would always be empty) and calls `collectElseFind(P, \emptyset)` for processes P that follow an oracle definition. The situation is similar for an oracle call inside an oracle body.

The `elsefind` facts can be used to add new facts to the facts \mathcal{F}_P . Indeed, if \mathcal{F}_P implies that M_1, \dots, M_l are defined for some values of i_1, \dots, i_m , then the fact `elsefind` $((i_1 \leq n_1, \dots, i_m \leq n_m), (M_1, \dots, M_l), M)$ implies that M is false for these values of i_1, \dots, i_m . Precisely, we define

$$\begin{aligned} \text{convert_elsefind}(\mathcal{F}^{\text{ElseFind}}, \mathcal{F}) = \{ & \neg\sigma M \mid \\ & \text{elsefind}((i_1 \leq n_1, \dots, i_m \leq n_m), (M_1, \dots, M_l), M) \in \mathcal{F}^{\text{ElseFind}}, \text{Dom}(\sigma) = \{i_1, \dots, i_m\}, \\ & \text{for each } j \in \{1, \dots, l\}, \sigma M_j \text{ is a subterm of } M'_j \text{ and } \text{defined}(M'_j) \in \mathcal{F} \} \end{aligned}$$

The possible images of σ are found by exploring the set of `defined` facts in \mathcal{F}_P . We execute:

$$\mathcal{F}_P \leftarrow \mathcal{F}_P \cup \text{convert_elsefind}(\mathcal{F}_P^{\text{ElseFind}}, \mathcal{F}_P)$$

In the implementation, an additional fact expresses that a pattern-matching failed: in the `else` branch of `let N = M in P else Q`, we know that $\forall x_1, \dots, x_l, N \neq M$, where x_1, \dots, x_l are the variables bound in the pattern N . In this report, we consider that the pattern-matching is encoded using assignments and tests, so we do not consider this fact further.

Furthermore, when the previous update of \mathcal{F}_P adds facts, we again complete the computed sets \mathcal{F}_P by adding facts that come from processes above P :

$$\mathcal{F}_P \leftarrow \mathcal{F}_P \cup \mathcal{F}_{P'} \text{ if } P \text{ is immediately under } P'$$

We could also iterate the addition of consequences of `defined` facts. (However, for simplicity, the current implementation does not perform such an iteration.)

We have `programpoint` $(\mu, I_\mu) \in \mathcal{F}_\mu$.

Lemma 31 *Let C be an evaluation context acceptable for Q_0 , \mathcal{A} an attacker in \mathcal{BB} , Tr be a trace of $C[Q_0]$ for \mathcal{A} , μ be a program point in Q_0 , and \mathcal{F}_μ be computed in Q_0 . If a configuration $Conf$ is at program point μ in Tr , then $Tr \preceq Conf \vdash \mathcal{F}_\mu$.*

When a single variable can be defined at multiple distinct program points, the basic version presented previously of the collected facts only contains the intersection of what holds at all the possible definitions points. A more refined version of \mathcal{F}_μ has also been implemented, that allows distinguishing cases depending on the multiple possible program points at which a variable is defined, generating several $\mathcal{F}_{\mu,c}$ for the various cases c . This more refined version is easily leveraged within proofs, simply by checking that a fact holds in all the possible cases over c . Note that more refined versions could be made by increasing the possible set of cases, but this would increase the overall complexity of the proof search. For this version, we have:

Lemma 32 *Let C be an evaluation context acceptable for Q_0 , \mathcal{A} an attacker in \mathcal{BB} , Tr be a trace of $C[Q_0]$ for \mathcal{A} , μ be a program point in Q_0 , and $\mathcal{F}_{\mu,c}$ be computed in Q_0 . If a configuration $Conf$ is at program point μ in Tr , then there exists c such that $Tr \preceq Conf \vdash \mathcal{F}_{\mu,c}$.*

\mathcal{F}_μ can be seen as a particular case of $\mathcal{F}_{\mu,c}$ by considering a single case c .

Corollary 3 *Let C be an evaluation context acceptable for Q_0 , \mathcal{A} an attacker in \mathcal{BB} , Tr be a trace of $C[Q_0]$ for \mathcal{A} , μ be a program point in Q_0 , and \mathcal{F}_μ (resp. $\mathcal{F}_{\mu,c}$) be computed in Q_0 . Let $Conf$ be a configuration at program point μ in Tr . Let θ be a renaming of I_μ to fresh replication indices and $\rho = \{\theta I_\mu \mapsto \sigma_{Conf} I_\mu\}$. Let $Conf'$ be a term or oracle body configuration in Tr such that $Conf \preceq_{Tr} Conf'$.*

We have $Tr \preceq Conf', \rho \vdash \theta \mathcal{F}_\mu$ and there exists c such that $Tr \preceq Conf', \rho \vdash \theta \mathcal{F}_{\mu,c}$.

In particular, $Tr, \rho \vdash \theta \mathcal{F}_\mu$ and there exists c such that $Tr, \rho \vdash \theta \mathcal{F}_{\mu,c}$.

Proof By Lemma 31, $Tr \preceq Conf \vdash \mathcal{F}_\mu$. By Lemma 32, there exists c such that $Tr \preceq Conf \vdash \mathcal{F}_{\mu,c}$. Let $\mathcal{F} = \mathcal{F}_\mu$ (resp. $\mathcal{F} = \mathcal{F}_{\mu,c}$) such that $Tr \preceq Conf \vdash \mathcal{F}$. By definition of ρ , we have $Tr \preceq Conf, \rho \vdash \theta \mathcal{F}$. Since $Conf \preceq_{Tr} Conf'$, the environment $E_{Tr \preceq Conf'}$ is an extension of $E_{Tr \preceq Conf}$, so the terms and defined facts in $\theta \mathcal{F}$ are preserved when considering $Tr \preceq Conf'$ instead of $Tr \preceq Conf$. (They do not use σ_{Conf} , resp. $\sigma_{Conf'}$, by the renaming θ .) Moreover, by Lemma 3, $\mu \mathcal{E}v_{Conf'}$ is an extension of $\mu \mathcal{E}v_{Conf}$, so the events are also preserved. By definition of $Tr \preceq Conf, \rho \vdash \text{programpoint}(\mathcal{S}_1, \widetilde{M}_1) \preceq \dots \preceq \text{programpoint}(\mathcal{S}_m, \widetilde{M}_m)$, the sequences of program points $\text{programpoint}(\mathcal{S}_1, \widetilde{M}_1) \preceq \dots \preceq \text{programpoint}(\mathcal{S}_m, \widetilde{M}_m)$ are also preserved. Since \mathcal{F} is a set of facts containing only terms, defined facts, events, and sequences of program points, we conclude that $Tr \preceq Conf', \rho \vdash \theta \mathcal{F}$.

The last point is obtained by choosing $Conf'$ to be the last configuration of Tr . \square

Lemma 33 *Let C be an evaluation context acceptable for Q_0 , \mathcal{A} an attacker in \mathcal{BB} , $Tr = \text{initConfig}(C[Q_0], \mathcal{A}) \xrightarrow{p_t} \dots \xrightarrow{p'_t} E, (\sigma, P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v$ be a trace for \mathcal{A} that does not execute any non-unique event of Q_0 with $P = \text{return}(a); Q$ for some a and Q or $P = \text{yield}$ or $P = \text{abort}$, μ be a program point in Q_0 , and $\mathcal{F}_\mu^{\text{Fut}}$ be computed in Q_0 . If the configuration $Conf$ is at program point μ in Tr and no executed process in the configurations between the configuration at the end of reduction step that contains $Conf$ (included) and $E, (\sigma, P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v$ (excluded) is of the form $\text{return}(a); Q$ for some a and Q or yield (when $Conf$ is a process configuration with process $\text{return}(a); Q$ for some a and Q or yield , we have $Conf = E, (\sigma, P) :: St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v$), then $Tr \vdash \mathcal{F}_\mu^{\text{Fut}}$.*

3.3 Local Dependency Analysis

For each program point P and each variable x , the local dependency analysis tries to find which variables and terms depend on $x[\tilde{i}]$ at program point P , where \tilde{i} denotes the current replication indices at the definition of x . It simplifies the game on-the-fly when possible.

For each occurrence of a process P and each variable x such that a random choice $x \stackrel{R}{\leftarrow} T$ occurs above P and T is a large type, we compute a set of terms $\text{indep}_P(x)$ that are independent of $x[\tilde{i}]$ where \tilde{i} denotes the current replication indices at the definition of x .

For each occurrence of a process P and each variable x such that a random choice $x \stackrel{R}{\leftarrow} T$ occurs above P and T is a large type, we also compute $\text{depend}_P(x)$ which can be either \top (I don't know) or a set of pairs (y, M) where $y[\tilde{i}]$ depends on $x[\tilde{i}]$ by assignments, and M is a term defining $y[\tilde{i}]$ as a function of $x[\tilde{i}]$. (The tuple \tilde{i} denotes the current replication indices at the definition of x and of y .)

$$\begin{aligned}
\text{depAnal}(Q, \text{indep}) = & \\
& \forall y, \text{depend}_Q(y) = \top; \text{indep}_Q = \text{indep} \\
& \text{if } Q = Q_1 \mid Q_2 \text{ then } \text{depAnal}(Q_1, \text{indep}); \text{depAnal}(Q_2, \text{indep}) \\
& \text{if } Q = \text{foreach } i \leq n \text{ do } Q' \text{ then } \text{depAnal}(Q', \text{indep}) \\
& \text{if } Q = \text{newOracle } O; Q' \text{ then } \text{depAnal}(Q', \text{indep}) \\
& \text{if } Q = O[\tilde{i}](x[\tilde{i}] : T) := P \text{ then } \text{depAnal}(P, \{\forall y, y \mapsto \top\}, \text{indep})
\end{aligned}$$

Figure 20: Local dependency analysis (1)

We define “ M characterizes a part of $x[\tilde{i}]$ at P ” as follows. Let α be defined by $\alpha(f(M_1, \dots, M_m)) = f(\alpha M_1, \dots, \alpha M_m)$; $\alpha(i) = i$ where i is a replication index; $\alpha(M') = M'$ when $M' \in \text{indep}_P(x)$; $\alpha(y[M_1, \dots, M_m]) = y[\alpha M_1, \dots, \alpha M_m]$ when $y \neq x$ and y either is defined only by random choices or $\text{depend}_P(x) \neq \top$ and $(y, M') \notin \text{depend}_P(x)$ for any M' ; $\alpha(y[M_1, \dots, M_m]) = y'[\alpha M_1, \dots, \alpha M_m]$ where y' is a fresh variable, otherwise. We write $y' = \alpha y$ in this case. We say that M characterizes a part of $x[\tilde{i}]$ at P when $\alpha M = M$ implies $f_1(\dots f_k((\alpha x)[\tilde{i}])) = f_1(\dots f_k(x[\tilde{i}]))$ for some uniform functions f_1, \dots, f_k , where $x[\tilde{i}]$ is a subterm of M , $(\alpha x)[\tilde{i}]$ is a subterm of αM , T' is the type of the result of f_1 (or of x when $k = 0$), and T' is a large type. In that case, the value of M uniquely determines the value of $f_1(\dots f_k(x[\tilde{i}]))$. This property is shown by a simple rewriting prover, as in the global dependency analysis.

We denote by $\text{subterms}(M)$ the set of subterms of the term M .

We say that M does not depend on x at P when M is built by function applications from terms in $\text{indep}_P(x)$, replications indices, and terms $y[M_1, \dots, M_m]$ such that M_1, \dots, M_m do not depend on x at P , $y \neq x$, and either y is defined only by random choices or $\text{depend}_P(x) \neq \top$ and $y \neq y'$ for all $(y', M') \in \text{depend}_P(x)$. Since terms in $\text{indep}_P(x)$ do not depend on $x[\tilde{i}]$ and when $\text{depend}_P(x) \neq \top$, variables not in the first component of $\text{depend}_P(x)$ do not depend on $x[\tilde{i}]$, the conditions above guarantee that M does not depend on $x[\tilde{i}]$, where \tilde{i} are the current replication indices at the definition of x .

When $\text{depend} \neq \top$, we denote by $M\text{depend}$ the term obtained from M by replacing $y[\tilde{i}]$ with M' for each $(y, M') \in \text{depend}$, where \tilde{i} denotes the replication indices at the definition of y .

We define simplifyTerm such that $\text{simplifyTerm}(M, P)$ is a simplified version of M , equal to M except in cases of negligible probability. The term $\text{simplifyTerm}(M, P)$ is defined as follows:

- Case 1: M is $M_1 = M_2$. For each x , we proceed as follows. If $\text{depend}_P(x) = \top$, let $M_0 = M_1$; otherwise, let $M_0 = M_1\text{depend}_P(x)$. Let M'_0 and M'_2 be obtained respectively from M_0 and M_2 by replacing all array indices that depend on x at P with fresh replication indices. If M'_0 characterizes a part of $x[\tilde{i}]$ at P , and M'_2 does not depend on x at P , then $\text{simplifyTerm}(M, P) = \text{false}$. Indeed, M is equal to false up to negligible probability in this case. We have similar cases swapping M_1 and M_2 or when M is $M_1 \neq M_2$. (In the latter case, $\text{simplifyTerm}(M, P) = \text{true}$.)
- Case 2: M is $M_1 \wedge M_2$. Let $M'_1 = \text{simplifyTerm}(M_1, P)$ and $M'_2 = \text{simplifyTerm}(M_2, P)$. If M'_1 or M'_2 are false, we return false. If M'_1 is true, we return M'_2 . If M'_2 is true, we return M'_1 . Otherwise, we return $M'_1 \wedge M'_2$. We have similar cases when M is $M_1 \vee M_2$ or $\neg M_1$.
- In all other cases, $\text{simplifyTerm}(M, P) = M$.

The local dependency analysis is defined in Figures 20 and 21. The function depAnal is initially called with $\text{depAnal}(Q_0, \emptyset)$ where \emptyset designates the function defined nowhere.

$\text{depAnal}(P, \text{depend}, \text{indep}) =$
 $\text{depend}_P = \text{depend}; \text{indep}_P = \text{indep}$
 if $P = \text{return } (N); Q$ then $\text{depAnal}(Q, \text{indep})$
 if $P = x[\tilde{i}] \stackrel{R}{\leftarrow} T; P'$ then
 if T is a large type then
 $\text{depend}'(x) = \emptyset; \text{indep}'(x) = \bigcup_{\text{defined}(M) \in \mathcal{F}_P} \text{subterms}(M)$
 $\forall y \neq x, \text{depend}'(y) = \text{depend}(y), \text{indep}'(y) = \text{indep}(y) \cup \{x[\tilde{i}]\}$
 $\text{depAnal}(P', \text{depend}', \text{indep}')$
 if $P = \text{let } x[\tilde{i}] : T = M \text{ in } P'$ then
 $\forall y$, if M does not depend on y at P then
 $\text{depend}'(y) = \text{depend}(y); \text{indep}'(y) = \{x[\tilde{i}]\} \cup \text{indep}(y)$
 else
 if $\text{depend}(y) \neq \top$ then
 $\text{depend}'(y) = \text{depend}(y) \cup \{(x, M\text{depend}(y))\}$
 else
 $\text{depend}'(y) = \top$
 $\text{indep}'(y) = \text{indep}(y)$
 $\text{depAnal}(P', \text{depend}', \text{indep}')$
 if $P = \text{find } (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j}$
 suchthat $\text{defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j$ then P_j else P'
 then
 for each $j \leq m, M'_j = \text{simplifyTerm}(M_j, P)$
 replace M_j with M'_j
 if $M'_j = \text{false}$ then remove the j -th branch
 if $M'_j = \text{true}$ and $l_j = 0$ then replace P' with yield
 if $m = 0$ then
 replace P with P' ; $\text{depAnal}(P', \text{depend}, \text{indep})$
 else if $m = 1, m_1 = l_1 = 0$, and $M_1 = \text{true}$ then
 replace P with P_1 ; $\text{depAnal}(P_1, \text{depend}, \text{indep})$
 else
 $\forall y$, if $\forall j, k, M_{jk}$ and M'_j do not depend on y at P then
 $\text{depend}'(y) = \text{depend}(y)$
 for each $j \leq m, \text{indep}_j(y) = \text{indep}(y) \cup \{M' \mid$
 $M' \in \text{subterms}(M) \text{ for some } \text{defined}(M) \in \mathcal{F}_{P_j}, M' \text{ does not depend on } y \text{ at } P\}$
 else
 $\text{depend}'(y) = \top; \text{for each } j \leq m, \text{indep}_j(y) = \text{indep}(y)$
 for each $j \leq m, \text{depAnal}(P_j, \text{depend}', \text{indep}_j)$
 $\text{depAnal}(P', \text{depend}', \text{indep})$

Figure 21: Local dependency analysis (2)

- For oracle definitions, `depAnal` sets $\text{depend}_Q(y)$ to \top , so that depend_Q gives no information, and propagates `indep`. Indeed, when $y[\tilde{i}']$ is set in some oracle body P_0 , the value of $y[\tilde{i}']$ may be passed to another oracle by P_0 or read by `find` in other oracle bodies executed after P_0 , so as soon as P_0 passes control to another oracle by the first return after the definition of y , we lose track of exactly which variables depend on $y[\tilde{i}']$. However, variables already defined before P_0 passes control to another oracle and proved to be independent of $y[\tilde{i}']$ remain independent of $y[\tilde{i}']$, so we can propagate `indep` in all subprocesses of P_0 .
- In the case of a return, `depAnal` forgets the information in depend_P as mentioned above.
- In the case of a random choice $x[\tilde{i}] \stackrel{R}{\leftarrow} T$, if T is a large type, we create the dependency information for the newly defined variable x : no variable depends on $x[\tilde{i}]$, and all terms already defined before the random choice are independent of $x[\tilde{i}]$. We also note that $x[\tilde{i}]$ is independent of $y[\tilde{i}']$ for other variables y by adding $x[\tilde{i}]$ to $\text{indep}(y)$.
- In the case of an assignment `let` $x[\tilde{i}] : T = M$, if M depends on $y[\tilde{i}']$ for some variable y , then $x[\tilde{i}]$ depends on $y[\tilde{i}']$, so x is added to $\text{depend}(y)$ (if it is not \top); otherwise, $x[\tilde{i}]$ does not depend on $y[\tilde{i}']$ so it is added to $\text{indep}(y)$.
- In the case of a `find`, we first simplify each condition of the `find`, remove branches when we can prove that they are taken with negligible probability, and remove the `find` itself when we know which branch is taken and this branch of the `find` does not define variables. Furthermore, if some condition of `find` depends on $y[\tilde{i}']$ for some variable y , $\text{depend}'(y)$ is set to \top : the control flow depends on $y[\tilde{i}']$ so future assignments in fact depend on $y[\tilde{i}']$ even if the assigned expression itself does not, so we can no longer keep track precisely of which variables depend on $y[\tilde{i}']$. Otherwise, we add all terms that are guaranteed to be defined and independent of $y[\tilde{i}']$ to $\text{indep}(y)$.

3.4 Equational Prover

We use an algorithm inspired by the Knuth-Bendix completion algorithm [51], with differences detailed below.

The prover manipulates pairs \mathcal{F}, \mathcal{R} where \mathcal{F} is a set of facts (M or $\text{defined}(M)$) and \mathcal{R} is a set of rewrite rules $M_1 \rightarrow M_2$. We say that M reduces into M' by $M_1 \rightarrow M_2$ when $M = C[M_1]$ and $M' = C[M_2]$ for some term context C . (That is, all variables in rewrite rules of \mathcal{R} are considered as constants.) The prover starts with a certain set of facts \mathcal{F} and $\mathcal{R} = \emptyset$. Then the prover transforms the pairs $(\mathcal{F}, \mathcal{R})$ by the following rules (the rule $\frac{\mathcal{F}, \mathcal{R}}{\mathcal{F}', \mathcal{R}'}$ means that \mathcal{F}, \mathcal{R} is transformed into $\mathcal{F}', \mathcal{R}'$):

$$\frac{\mathcal{F} \cup \{F\}, \mathcal{R}}{\mathcal{F} \cup \{F'\}, \mathcal{R}} \quad \begin{array}{l} \text{if } F \text{ reduces into } F' \text{ by a rule of } \mathcal{R} \text{ or} \\ \text{a user-defined rewrite rule knowing } \mathcal{F}, \mathcal{R} \end{array} \quad (17)$$

$$\frac{\mathcal{F} \cup \{M_1 \wedge M_2\}, \mathcal{R}}{\mathcal{F} \cup \{M_1, M_2\}, \mathcal{R}} \quad (18)$$

$$\frac{\mathcal{F} \cup \{x[M_1, \dots, M_m] = x[M'_1, \dots, M'_m]\}, \mathcal{R}}{\mathcal{F} \cup \{M_1 = M'_1, \dots, M_m = M'_m\}, \mathcal{R}} \quad \begin{array}{l} \text{when } x \text{ is defined only by random choices} \\ x \stackrel{R}{\leftarrow} T \text{ and } T \text{ is a large type} \end{array} \quad (19)$$

$\frac{\mathcal{F} \cup \{M_1 = M_2\}, \mathcal{R}}{\{\text{false}\}, \mathcal{R}}$ when one of the following conditions holds:

- denoting by M'_1 the term obtained from M_1 by replacing all array indices that are not replication indices with fresh replication indices, we have the following properties: x occurs in M'_1 , x is defined only by random choices $x \stackrel{R}{\leftarrow} T$, T is a large type, M'_1 characterizes a part of x , and M_2 is obtained by optionally applying function symbols to terms of the form $y[\widetilde{M}]$ where y is defined only by random choices and $y \neq x$;
- $\text{simplifyTerm}(M_1 = M_2, P) = \text{false}$, where P is the current program point.

(20)

$$\frac{\mathcal{F} \cup \{M = M'\}, \mathcal{R}}{\mathcal{F}, \mathcal{R} \cup \{M \rightarrow M'\}} \text{ if } M > M' \quad (21)$$

$$\frac{\mathcal{F}, \mathcal{R} \cup \{M_1 \rightarrow M_2\} \text{ if } M_2 \text{ reduces into } M'_2 \text{ by a rule of } \mathcal{R}}{\mathcal{F} \cup \{M_1 = M'_2\}, \mathcal{R}} \text{ or a user-defined rewrite rule knowing } \mathcal{F}, \mathcal{R} \quad (22)$$

$$\frac{\mathcal{F}, \mathcal{R} \cup \{M_1 \rightarrow M_2\}}{\mathcal{F} \cup \{M'_1 = M_2\}, \mathcal{R}} \text{ if } M_1 \text{ reduces into } M'_1 \text{ by a rule of } \mathcal{R} \quad (23)$$

We also use the symmetric of Rules (20) and (21) obtained by swapping the two sides of the equality.

Rule (17) simplifies facts using rewrite rules. Let us define when M reduces into M' by a user-defined rewrite rule knowing \mathcal{F}, \mathcal{R} . For the first kind of rewrite rules of Section 3.1, this is simply when M reduces into M' by the considered rewrite rule. For the second kind of rewrite rules of Section 3.1, consider a rewrite rule $y_1 \stackrel{R}{\leftarrow} T'_1, \dots, y_l \stackrel{R}{\leftarrow} T'_l, \forall x_1 : T_1, \dots, \forall x_m : T_m, M_1 \rightarrow M_2$. Suppose that $M = C[\sigma M_1]$, where C is a term context and σ is a substitution that maps x_j to any term of type T_j for all $j \leq m$ and y_j to terms to the form $z_j[\widetilde{M}_j]$ where z_j is defined only by random choices $z_j \stackrel{R}{\leftarrow} T'_j$ for all $j \leq l$. Let $\text{Cond} = \{\widetilde{M}_j = \widetilde{M}_{j'} \mid j \neq j' \wedge z_j = z_{j'}\}$.

- If $\text{Cond} = \emptyset$, the random choices $z_j[\widetilde{M}_j]$ are independent, so $M = C[\sigma M_1]$ reduces into $M' = C[\sigma M_2]$.
- If $\text{Cond} \neq \emptyset$, let $\text{Cond} = \{\text{cond}_1, \dots, \text{cond}_k\}$. When $\text{cond}_1 \vee \dots \vee \text{cond}_k$ is true, the random choices $z_j[\widetilde{M}_j]$ are not independent, so we must leave M as it is. When $\text{cond}_1 \vee \dots \vee \text{cond}_k$ is false, M reduces into $M' = C[\sigma M_2]$. Suppose that $M' = \text{false}$. Hence, we could rewrite M into $\text{if } \text{cond}_1 \vee \dots \vee \text{cond}_k \text{ then } M \text{ else false}$, that is, $(\text{cond}_1 \vee \dots \vee \text{cond}_k) \wedge M$, that is, $(\text{cond}_1 \wedge M) \vee \dots \vee (\text{cond}_k \wedge M)$. However, since this term itself contains M , this transformation could lead to a loop. We avoid it as follows. If for all $k' = 1, \dots, k$, M is transformed into $M'_{k'}$ by rewrite rules generated by our equational prover from $\mathcal{F} \cup \{\text{cond}_{k'}\}, \mathcal{R}$ and user-defined rewrite rules, using only the first kind of user-defined rewrite rules of Section 3.1, and $M'_{k'} \neq M_k$, then we rewrite M into $(\text{cond}_1 \wedge M'_1) \vee \dots \vee (\text{cond}_k \wedge M'_k)$. Otherwise, M does not rewrite by the considered user-defined rewrite rule knowing \mathcal{F}, \mathcal{R} .

Rule (18) decomposes conjunctions of facts. Rules (19) and (20) exploit the elimination of collisions between random values. Rule (19) takes into account that, when x is defined by a random choice of a large type, two different cells of x have a negligible probability of containing the same value. So when two cells of x contain the same value, we can conclude up to negligible probability that they are the same cell. Rule (20) expresses that M_1 and M_2 have a negligible probability of being equal when x is defined by a random choice of a large type, M_1 characterizes a part of x , and M_2 does not depend of x . The first item of (20) establishes these properties without further dependency analysis and the second item exploits the local dependency analysis.

Additionally, in simplification (**simplify**), if \mathcal{F} contains $M_1 = M_2$ such that x occurs in M_1 , x is defined only by random choices $x \stackrel{R}{\leftarrow} T$, T is a large type, and M_1 characterizes a part of x , we trigger the global dependency analysis (Section 5.1.20). If the global dependency analysis succeeds in transforming the game, the simplification is restarted from the game obtained after global dependency analysis.

Rule (21) is applied only when Rules (17) to (20) cannot be applied. Rule (21) transforms equations into rewrite rules by orienting them. We say that $M > M'$ when either M is the form $x[\widetilde{M}]$, x does not occur in M' , and x is not defined only by random choices, or $M = x[M_1, \dots, M_m]$, $M' = x[M'_1, \dots, M'_m]$, and for all $j \leq m$, $M_j > M'_j$. Intuitively, our goal is to replace M with M' when M' defines the content of the variable M . (Notice that this is not an ordering; the Knuth-Bendix algorithm normally uses a reduction ordering to orient equations. However, we tried some reduction orderings, namely the lexicographic path ordering and the Knuth-Bendix ordering, and obtained disappointing results: the prover fails to prove many equalities because too many equations are left unoriented. The simple heuristic given above succeeds more often, at the expense of a greater risk of non-termination, but that does not cause problems in practice on our examples. We believe that this comes from the particular structure of equations, which come from let definitions and from conditions of find or if, and tend to define variables from other variables without creating dependency cycles.)

Rules (22) and (23) are systematically applied to simplify all rewrite rules of \mathcal{R} after a new rewrite rule has been added by Rule (21). Since all terms in rewrite rules of \mathcal{R} are considered as constants, Rule (23) in fact includes the deduction of equations from critical pairs done by the standard Knuth-Bendix completion algorithm.

We say that \mathcal{F} yields a contradiction when the prover, starting from (\mathcal{F}, \emptyset) , derives false.

Extension: *That depends on the program point for local dependency analysis. We ignore that point for now. For the proof of correspondences, local dependency analysis is not used. The current game Q_0 is used for the basic dependency analysis (rule (19) and first item of (20)).*

Lemma 34 *If for all $j \in J$, \mathcal{F}_j yields a contradiction in a game Q_0 , then CryptoVerif returns a probability p such that for all attackers $\mathcal{A} \in \mathcal{BB}$ and all evaluation contexts C acceptable for Q_0 with any public variables, $\Pr^{\mathcal{A}}[C[Q_0] \preceq \bigvee_{j \in J} \exists \tilde{x}_j \in \widetilde{T}_j, \bigwedge \mathcal{F}_j] \leq p(\mathcal{A}, C)$, where \tilde{x}_j are the replication indices and non-process variables that occur in \mathcal{F}_j and \widetilde{T}_j are their types.*

In particular, the lemma states that, when several sets of facts \mathcal{F}_j yield a contradiction in the same game, CryptoVerif counts only once in the probability p the collisions that are eliminated in proofs that \mathcal{F}_j yields a contradiction for several j .

More generally, let us consider an algorithm φ built from the following grammar:

| | |
|---|----------------------|
| $\varphi ::=$ | algorithm |
| \mathcal{F} yields a contradiction | equational proof |
| $\varphi_1 \wedge \varphi_2$ | conjunction |
| $\varphi_1 \vee \varphi_2$ | disjunction |
| ψ | mathematical formula |
| if ψ then φ_1 else φ_2 | test |

The mathematical formulas ψ in such algorithms must not depend on the executed trace. (They may depend on the syntax of the game Q_0 or on the set of public variables V , for instance.)

We translate such algorithms into logical formulas on traces:

$$\{\{\mathcal{F} \text{ yields a contradiction}\}\} = \neg \exists \tilde{x} \in \widetilde{T}, \bigwedge \mathcal{F} \quad \text{where } \tilde{x} \text{ are the replication indices and non-process variables that occur in } \mathcal{F} \text{ and } \widetilde{T} \text{ are their types.}$$

$$\begin{aligned} \{\{\varphi_1 \wedge \varphi_2\}\} &= \{\{\varphi_1\}\} \wedge \{\{\varphi_2\}\} \\ \{\{\varphi_1 \vee \varphi_2\}\} &= \begin{cases} \{\{\varphi_1\}\} & \text{if } \varphi_1 \\ \{\{\varphi_2\}\} & \text{otherwise} \end{cases} \\ \{\{\psi\}\} &= \psi \\ \{\{\text{if } \psi \text{ then } \varphi_1 \text{ else } \varphi_2\}\} &= \text{if } \psi \text{ then } \{\{\varphi_1\}\} \text{ else } \{\{\varphi_2\}\} \end{aligned}$$

Intuitively, when algorithm φ returns true, CryptoVerif shows that the formula $\{\{\varphi\}\}$ holds for most traces. It bounds the probability of the traces for which this formula does not hold, as shown by the following lemma.

Lemma 35 *If algorithm φ returns true in a game Q_0 , then CryptoVerif returns a probability p such that for all attackers $\mathcal{A} \in \mathcal{BB}$ and all evaluation contexts C acceptable for Q_0 with any public variables, $\Pr^{\mathcal{A}}[C[Q_0] \preceq \neg\{\{\varphi\}\}] \leq p(\mathcal{A}, C)$.*

Proof We show by induction on the definition of φ that, if φ returns true and $Tr \vdash \neg\{\{\varphi\}\}$, then there exists \mathcal{F} such that “ \mathcal{F} yields a contradiction” has been called in the evaluation of φ and returned true, and $Tr \vdash \exists \tilde{x} \in \tilde{T}, \wedge \mathcal{F}$ where \tilde{x} are the replication indices and non-process variables that occur in \mathcal{F} and \tilde{T} are their types.

- Case $\varphi = (\mathcal{F}$ yields a contradiction): obvious.
- Case $\varphi = \varphi_1 \wedge \varphi_2$: Since φ returns true, φ_1 and φ_2 both return true. Since $Tr \vdash \neg(\{\{\varphi_1\}\} \wedge \{\{\varphi_2\}\})$, we have either $Tr \vdash \neg\{\{\varphi_1\}\}$ or $Tr \vdash \neg\{\{\varphi_2\}\}$. In the first case, by induction hypothesis on φ_1 , there exists \mathcal{F} such that “ \mathcal{F} yields a contradiction” has been called in the evaluation of φ_1 and returned true, and $Tr \vdash \exists \tilde{x} \in \tilde{T}, \wedge \mathcal{F}$ where \tilde{x} are the replication indices and non-process variables that occur in \mathcal{F} and \tilde{T} are their types. Moreover, “ \mathcal{F} yields a contradiction” has been called in the evaluation of φ . The second case is symmetric.
- Case $\varphi = \varphi_1 \vee \varphi_2$: If φ_1 returns true, then $\{\{\varphi\}\} = \{\{\varphi_1\}\}$, so $Tr \vdash \neg\{\{\varphi_1\}\}$. We conclude by induction hypothesis on φ_1 , as above. If φ_1 returns false, then φ_2 returns true, since φ returns true. Hence $\{\{\varphi\}\} = \{\{\varphi_2\}\}$, so $Tr \vdash \neg\{\{\varphi_2\}\}$. We conclude by induction hypothesis on φ_2 .
- Case $\varphi = \psi$: since ψ evaluates to true, there is no trace Tr such that $Tr \vdash \neg\psi$, so the property holds trivially.
- Case $\varphi = \text{if } \psi \text{ then } \varphi_1 \text{ else } \varphi_2$: if ψ evaluates to true, then we conclude by induction hypothesis on φ_1 . Indeed, since φ returns true, φ_1 returns true. Since $Tr \vdash \neg\{\{\varphi\}\}$, we have $Tr \vdash \neg\{\{\varphi_1\}\}$. By induction hypothesis, there exists \mathcal{F} such that “ \mathcal{F} yields a contradiction” has been called in the evaluation of φ_1 and returned true, and $Tr \vdash \exists \tilde{x} \in \tilde{T}, \wedge \mathcal{F}$ where \tilde{x} are the replication indices and non-process variables that occur in \mathcal{F} and \tilde{T} are their types. Then “ \mathcal{F} yields a contradiction” has also been called in the evaluation of φ . Similarly, if ψ evaluates to false, then we conclude by induction hypothesis on φ_2 .

We conclude by Lemma 34. □

4 success: Criteria for Proving Security Properties

The command **success** tries to prove the active queries, as explained below. We consider an oracle definition Q_0 that satisfies Properties 3 and 4, and prove secrecy and correspondence properties for Q_0 .

$$\begin{aligned}
C ::= & [] \\
& y[M_1, \dots, M_{k-1}, C, M_{k+1}, \dots, M_m] \\
& f(M_1, \dots, M_{k-1}, C, M_{k+1}, \dots, M_m) \\
& y[\tilde{i}] \stackrel{R}{\leftarrow} T; C \\
& \text{let } y[\tilde{i}] = M \text{ in } C \\
& \text{if } M \text{ then } C \text{ else } N' \\
& \text{if } M \text{ then } N \text{ else } C \\
& \text{find}[\text{unique?}] \left(\bigoplus_{j=1, \dots, m; j \neq k} \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(\widetilde{M}_j) \wedge M_j \text{ then } N_j \right) \\
& \quad \oplus \tilde{u}_k[\tilde{i}] = \tilde{i}_k \leq \tilde{n}_k \text{ suchthat defined}(\widetilde{M}_k) \wedge M_k \text{ then } C \text{ else } N \\
& \text{find}[\text{unique?}] \left(\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(\widetilde{M}_j) \wedge M_j \text{ then } N_j \right) \text{ else } C \\
& \text{event } e(\widetilde{M}); C
\end{aligned}$$

Figure 22: Event contexts

4.1 Secrecy

Let us now define syntactic criteria that allow us to prove secrecy properties of protocols. We first define the function `noleak`, explained below. This function implicitly depends on the current game Q_0 and the public variables V .

$$\text{noleak}(x[\widetilde{M}'], \mathcal{I}, \mathcal{F}) = (\mathcal{F} \text{ yields a contradiction}) \vee ((x \notin V) \wedge \bigwedge_{\mu x[\widetilde{M}] \text{ in } Q_0}$$

– if $\mu x[\widetilde{M}]$ is in M in an assignment `let $y[\tilde{i}] = M$ in Q_0` , M is built from replication indices, variables, function applications, and conditionals, and the current call is not inside a call to `noleak($y[-], -, -$)`, then

$$\text{noleak}(y[\theta\tilde{i}], \mathcal{I} \cup \{\theta\tilde{i}\}, \mathcal{F} \cup \theta\mathcal{F}_\mu \cup \{\theta\widetilde{M} = \widetilde{M}'\})$$

where θ is a renaming of \tilde{i} to fresh replication indices

– if $\mu x[\widetilde{M}]$ is in `event $e(M_1, \dots, M_{k-1}, C[\mu x[\widetilde{M}]], M_{k+1}, \dots, M_m)$ in Q_0` for C defined in Figure 22, then true

– otherwise, $\mathcal{F} \cup \theta\mathcal{F}_\mu \cup \{\theta\widetilde{M} = \widetilde{M}'\}$ yields a contradiction
where θ is a renaming of I_μ to fresh replication indices)

The function call `noleak($x[\widetilde{M}'], \mathcal{I}, \mathcal{F}$)` shows that $x[\widetilde{M}']$ does not leak to the adversary, assuming \mathcal{F} holds. The set \mathcal{I} contains all replication indices that appear in \mathcal{F} . If \mathcal{F} yields a contradiction, `noleak` is true, since it shows the absence of leak *assuming \mathcal{F} holds*. Otherwise, $x[\widetilde{M}']$ may leak either because $x \in V$ so x is a public variable, or because of an occurrence of a term $x[\widetilde{M}']$ in the game that reads $x[\widetilde{M}']$ (so $\widetilde{M} = \widetilde{M}'$ holds) and such that the result of $x[\widetilde{M}']$ leaks.

- In case $x[\widetilde{M}']$ occurs in the term M in an assignment `let $y[\tilde{i}] = M$` , the function `noleak` recursively tries to prove that $y[\tilde{i}]$ does not leak, when $y[\tilde{i}]$ may use $x[\widetilde{M}']$, that is, when $\widetilde{M}' = \widetilde{M}$. The fact $\widetilde{M}' = \widetilde{M}$ and the facts \mathcal{F}_μ that hold at the program point μ of $x[\widetilde{M}']$

are added to the known facts \mathcal{F} in the recursive call. Indeed, these facts are known to hold in this case. The replication indices are renamed to fresh indices in order to avoid using the same index variable for indices that can actually take different values.

- In case $x[\widetilde{M}]$ occurs in the arguments of an event, the arguments of the event do not leak to the adversary, so this occurrence of $x[\widetilde{M}]$ does not make $x[\widetilde{M}']$ leak.
- In all other cases, we consider that the result of $x[\widetilde{M}]$ may leak, so, in order to prove that $x[\widetilde{M}']$ does not leak, we show that the occurrence of $x[\widetilde{M}]$ at μ cannot read $x[\widetilde{M}']$, by showing that $\widetilde{M} = \widetilde{M}'$, \mathcal{F}_μ , and \mathcal{F} together yield a contradiction.

Definition 15 (μ follows a definition of x) We say that μ follows a definition of x when $x[\widetilde{i}] \stackrel{R}{\leftarrow} T; {}^\mu \dots$, let $x[\widetilde{i}] = M$ in ${}^\mu \dots$, find[unique?] ($\bigoplus_{j=1}^m \widetilde{u}_j[\widetilde{i}] = \widetilde{i}_j \leq \widetilde{n}_j$ such that defined(\widetilde{M}_j) \wedge M_j then ${}^{\mu_j} \dots$) else \dots with $\mu_j = \mu$ and x in \widetilde{u}_j for some $j \leq m$, or $O(x[\widetilde{i}] : T) := {}^\mu P$ occurs in Q_0 .

We do not mention `get` in the previous definition, because it is excluded by Property 3. For each μ that follows a definition of x in Q_0 , we define $\text{defRand}_\mu(x)$ as follows:

$$\text{defRand}_\mu(x) = \begin{cases} x[\widetilde{i}] & \text{if } x[\widetilde{i}] \stackrel{R}{\leftarrow} T; {}^\mu \dots \text{ occurs in } Q_0 \\ y[\widetilde{M}] & \text{if let } x[\widetilde{i}] : T = y[\widetilde{M}] \text{ in } {}^\mu \dots \text{ occurs in } Q_0 \text{ and} \\ & y \text{ is defined only by random choices in } Q_0 \end{cases}$$

In all other cases, $\text{defRand}_\mu(x)$ is not defined. The variable $\text{defRand}_\mu(x)$ is the random variable that defines x just before program point μ . When x itself is chosen randomly at that point, $\text{defRand}_\mu(x)$ is simply $x[\widetilde{i}]$, where \widetilde{i} are the current replication indices. When x is defined by an assignment of a variable $y[\widetilde{M}]$ that is random, $\text{defRand}_\mu(x)$ is that variable. Otherwise, we give up and do not define $\text{defRand}_\mu(x)$.

$$\begin{aligned} \text{prove}^{\text{1-ses.secr.}(x)}(\mu) &= \text{defRand}_\mu(x) \text{ is defined and} \\ &\quad \text{noleak}(\theta \text{defRand}_\mu(x), \{\theta I_\mu\}, \theta \mathcal{F}_\mu) \\ &\quad \text{where } \theta \text{ is a renaming of } I_\mu \text{ to fresh replication indices} \\ \text{prove}^{\text{1-ses.secr.}(x)}(\mathcal{S}) &= \bigwedge_{\mu \in \mathcal{S}} \text{prove}^{\text{1-ses.secr.}(x)}(\mu) \end{aligned}$$

The function call $\text{prove}^{\text{1-ses.secr.}(x)}(\mu)$ proves one-session secrecy for the definition of x just before program point μ . It considers only the cases in which x is defined either by a random choice or by an assignment from a random choice. In other cases, the proof fails. (These other cases can typically be handled by first removing assignments as needed.) Intuitively, $\text{prove}^{\text{1-ses.secr.}(x)}(\mu)$ guarantees that, when $x[\widetilde{i}]$ is defined just before program point μ , the random variable $\text{defRand}_\mu(x)$ that defines $x[\widetilde{i}]$ does not leak, knowing that the facts \mathcal{F}_μ hold. Only events and variables $y[\widetilde{i}']$ that do not leak depend on the random choice that defines $x[\widetilde{i}]$; the sent messages and the control flow of the process are independent of $x[\widetilde{i}]$, so the adversary obtains no information on $x[\widetilde{i}]$. That guarantees the one-session secrecy of $x[\widetilde{i}]$ when it is defined just before μ . This is verified for all program points in \mathcal{S} by $\text{prove}^{\text{1-ses.secr.}(x)}(\mathcal{S})$. When x is defined by assignment of $z[\widetilde{M}]$, this proof of one-session secrecy allows some array cells of z to leak, provided the array cells $z[\widetilde{M}]$ used to define x do not leak.

In order to prove secrecy, we also define $\text{prove}^{\text{distinct}(x)}(\mu_1, \mu_2) = (z_1 \neq z_2) \vee (\theta_1 \mathcal{F}_{\mu_1} \cup \theta_2 \mathcal{F}_{\mu_2} \cup \{\theta_1 \widetilde{M}_1 = \theta_2 \widetilde{M}_2, \widetilde{i}_1 \neq \widetilde{i}_2\})$ yields a contradiction), where $\text{defRand}_{\mu_1}(x) = z_1[\widetilde{M}_1]$, $\text{defRand}_{\mu_2}(x) = z_2[\widetilde{M}_2]$, \widetilde{i} are the current replication indices at the definition of x , θ_1 and θ_2 are two distinct renamings of \widetilde{i} to fresh replication indices, $\widetilde{i}_1 = \theta_1 \widetilde{i}$, and $\widetilde{i}_2 = \theta_2 \widetilde{i}$. Intuitively, $\text{prove}^{\text{distinct}(x)}(\mu_1, \mu_2)$ guarantees that, if $x[\widetilde{i}_1]$ is defined at μ_1 , so $x[\widetilde{i}_1] = z_1[\theta_1 \widetilde{M}_1]$, and $x[\widetilde{i}_2]$ is defined at μ_2 , so $x[\widetilde{i}_2] = z_2[\theta_2 \widetilde{M}_2]$, with $\widetilde{i}_1 \neq \widetilde{i}_2$, then the random variables that define x in these two cases, $z_1[\theta_1 \widetilde{M}_1]$ and $z_2[\theta_2 \widetilde{M}_2]$, are different, that is, $z_1 \neq z_2$ or $\theta_1 \widetilde{M}_1 \neq \theta_2 \widetilde{M}_2$. Therefore, $z_1[\theta_1 \widetilde{M}_1]$ is independent of $z_2[\theta_2 \widetilde{M}_2]$, so $x[\widetilde{i}_1]$ is independent of $x[\widetilde{i}_2]$. Combining this information with the proof of one-session secrecy, we can prove secrecy of x : we define

$$\text{prove}^{\text{Secrecy}(x)}(\mathcal{S}) = \text{prove}^{\text{1-ses.secr.}(x)}(\mathcal{S}) \wedge \bigwedge_{\mu_1, \mu_2 \in \mathcal{S}} \text{prove}^{\text{distinct}(x)}(\mu_1, \mu_2)$$

The proof of bit secrecy is the same as for one-session secrecy:

$$\text{prove}^{\text{bit secr.}(x)}(\mathcal{S}) = \text{prove}^{\text{1-ses.secr.}(x)}(\mathcal{S})$$

The proof of (one-session or bit) secrecy is justified by the following proposition.

Proposition 1 ((One-session or bit) secrecy) *Consider an oracle definition Q_0 that satisfies Properties 3 and 4. Let sp be $\text{1-ses.secr.}(x)$, $\text{Secrecy}(x)$, or $\text{bit secr.}(x)$. Let $\mathcal{S} = \{\mu \mid \mu \text{ follows a definition of } x\}$. If $\text{prove}^{sp}(\mathcal{S})$ and for all attackers $\mathcal{A} \in \mathcal{BB}$ and all evaluation contexts C acceptable for Q_0 , $\Pr^{\mathcal{A}}[C[Q_0] \preceq \neg\{\text{prove}^{sp}(\mathcal{S})\}] \leq p(\mathcal{A}, C)$, then Q_0 satisfies sp with public variables V ($x \notin V$) up to probability p' such that $p'(\mathcal{A}, C) = p(\mathcal{A}, C[C_{sp}[]])$ and $\text{Bound}_{Q_0}(V \cup \{x\}, sp, D_{\text{false}}, p)$.*

The proof of Proposition 1 relies on the following definitions and lemma. We have

$$\{\{\text{noleak}(x[\widetilde{M}'], \mathcal{I}, \mathcal{F})\}\} = (\forall \mathcal{I}, \neg \bigwedge \mathcal{F}) \vee ((x \notin V) \wedge \bigwedge_{\mu x[\widetilde{M}] \text{ in } Q_0}$$

– if $\mu x[\widetilde{M}]$ is in M in an assignment let $y[\widetilde{i}] = M$, M is built from replication indices, variables, function applications, and conditionals, and the current call is not inside a call to $\text{noleak}(y, -, -)$, then

$$\{\{\text{noleak}(y[\theta \widetilde{i}], \mathcal{I} \cup \{\theta \widetilde{i}\}, \mathcal{F} \cup \theta \mathcal{F}_{\mu} \cup \{\theta \widetilde{M} = \widetilde{M}'\})\}\}$$

where θ is a renaming of \widetilde{i} to fresh replication indices

– if $\mu x[\widetilde{M}]$ is in event $e(M_1, \dots, M_{k-1}, C[\mu x[\widetilde{M}]], M_{k+1}, \dots, M_m)$ for C defined in Figure 22, then true

– otherwise, $\forall (\mathcal{I} \cup \theta I_{\mu}), \neg \bigwedge (\mathcal{F} \cup \theta \mathcal{F}_{\mu} \cup \{\theta \widetilde{M} = \widetilde{M}'\})$

where θ is a renaming of I_{μ} to fresh replication indices)

$\{\{\text{noleak}(x[\widetilde{M}'], \mathcal{I}, \mathcal{F})\}\}$ is the logical formula that is guaranteed when $\text{noleak}(x[\widetilde{M}'], \mathcal{I}, \mathcal{F})$ succeeds, up to a small probability computed by the equational prover and that bounds the probability $\Pr^{\mathcal{A}}[C[Q_0] \preceq \neg\{\text{noleak}(x[\widetilde{M}'], \mathcal{I}, \mathcal{F})\}]$. It is obtained by collecting formulas guaranteed by each call to “ \mathcal{F} yields a contradiction”: such a call guarantees $\forall \widetilde{z} \in \widetilde{T}'', \neg \bigwedge \mathcal{F}$ up to a small probability that it evaluates, where \widetilde{z} are the non-process variables in \mathcal{F} and \widetilde{T}'' are their types. In the definition of $\{\{\text{noleak}(x[\widetilde{M}'], \mathcal{I}, \mathcal{F})\}\}$, the notation $\forall \mathcal{I}$ means that all variables in \mathcal{I} are

universally quantified in their respective types. The notation $\forall(\mathcal{I} \cup \theta I_\mu)$ is similar. We have similarly

$$\begin{aligned} \{\{\text{prove}^{1\text{-ses.secr.}(x)}(\mu)\}\} &= \begin{cases} \{\{\text{noleak}(\theta\text{defRand}_\mu(x), \{\theta I_\mu\}, \theta\mathcal{F}_\mu)\}\} \\ \quad \text{if } \text{defRand}_\mu(x) \text{ is defined,} \\ \quad \text{where } \theta \text{ is a renaming of } I_\mu \text{ to fresh replication indices} \\ \text{false} \quad \text{otherwise} \end{cases} \\ \{\{\text{prove}^{1\text{-ses.secr.}(x)}(\mathcal{S})\}\} &= \bigwedge_{\mu \in \mathcal{S}} \{\{\text{prove}^{1\text{-ses.secr.}(x)}(\mu)\}\} \\ \{\{\text{prove}^{\text{distinct}(x)}(\mu_1, \mu_2)\}\} &= (z_1 \neq z_2) \vee (\forall \tilde{i}_1, \forall \tilde{i}_2, \neg \bigwedge \theta_1 \mathcal{F}_{\mu_1} \cup \theta_2 \mathcal{F}_{\mu_2} \cup \{\theta_1 \tilde{M}_1 = \theta_2 \tilde{M}_2, \tilde{i}_1 \neq \tilde{i}_2\}) \\ \{\{\text{prove}^{\text{Secrecy}(x)}(\mathcal{S})\}\} &= \{\{\text{prove}^{1\text{-ses.secr.}(x)}(\mathcal{S})\}\} \wedge \bigwedge_{\mu_1, \mu_2 \in \mathcal{S}} \{\{\text{prove}^{\text{distinct}(x)}(\mu_1, \mu_2)\}\} \\ \{\{\text{prove}^{\text{bit secr.}(x)}(\mathcal{S})\}\} &= \{\{\text{prove}^{1\text{-ses.secr.}(x)}(\mathcal{S})\}\} \end{aligned}$$

The only semantic rules that can add $x[\tilde{a}]$ to the environment E are (NewT), (LetT), (FindT1), (New), (Let), (Find1), (OracleCall), (Return), (AttIO), and (AttYield). (`get` is excluded by Property 3.) By Corollary 1, the target term or process of these rules is a subterm or subprocess of Q_0 up to renaming of oracles. Hence, the target configuration Conf of these rules is at some program point μ in Tr . In this case, we say that $x[\tilde{a}]$ is defined just before μ in a trace Tr . Furthermore, given $x[\tilde{a}]$ and Tr , there is at most one program point μ such that $x[\tilde{a}]$ is defined just before μ in Tr , by Lemma 29.

Let sp be $1\text{-ses.secr.}(x)$, $\text{Secrecy}(x)$, or $\text{bit secr.}(x)$. Let Tr be a trace of $C[C_{sp}[Q_0]]$. Let $E = E_{Tr}$. We define the set $\text{Tidx}(Tr)$ of indices of successful test queries as follows:

- When sp is $1\text{-ses.secr.}(x)$: Let μ_t be the program point of the oracle that performs the test query in $Q_{1\text{-ses.secr.}(x)}$: ${}^{\mu_t}O_s(u_1 : [1, n_1], \dots, u_m : [1, n_m])$. If there is a (AttIO) or (OracleCall) reduction in Tr with $\mu' = \mu_t$, we define E' to be the environment after that reduction. If $x[E(u_1), \dots, E(u_m)] \in \text{Dom}(E')$, we let $\text{Tidx}(Tr) = \{\epsilon\}$ (the empty sequence of indices). Otherwise, $\text{Tidx}(Tr) = \emptyset$.
- When sp is $\text{Secrecy}(x)$: Let μ_t be the program point of the oracle definition that performs the test query in $Q_{\text{Secrecy}(x)}$: ${}^{\mu_t}O_s[i_s](u_1 : [1, n_1], \dots, u_m : [1, n_m])$. Let $\text{Tidx}(Tr) = \{a \in [1, n_s] \mid \text{there is a (AttIO) or (OracleCall) reduction in } Tr \text{ with } \mu' = \mu_t, \sigma'(\tilde{i}) = a, \text{ and } x[E(u_1[a]), \dots, E(u_m[a])] \in \text{Dom}(E') \text{ where } E' \text{ is the environment after that reduction, and for all a (AttIO) or (OracleCall) reductions in } Tr \text{ before the latter reduction, with } \mu' = \mu_t, \sigma'(\tilde{i}) = a', E(u_1[a']) = E(u_1[a]), \dots, \text{ and } E(u_m[a']) = E(u_m[a]), \text{ we have } x[E(u_1[a]), \dots, E(u_m[a])] \notin \text{Dom}(E'') \text{ where } E'' \text{ is the environment after that reduction}\}$.

The test query with index a is the first successful test query for $x[E(u_1[a]), \dots, E(u_m[a])]$, so $x[E(u_1[a]), \dots, E(u_m[a])]$ is defined at that test query, that is, $x[E(u_1[a]), \dots, E(u_m[a])] \in \text{Dom}(E')$. For all previous test queries on the same indices, $x[E(u_1[a]), \dots, E(u_m[a])]$ was not defined, that is, $x[E(u_1[a]), \dots, E(u_m[a])] \notin \text{Dom}(E'')$. The bound n_s and the variables u_1, \dots, u_m come from $Q_{\text{Secrecy}(x)}$.

- When sp is $\text{bit secr.}(x)$: Let μ_t be the program point of the oracle definition in $Q_{\text{bit secr.}(x)}$: ${}^{\mu_t}O'_s(b' : \text{bool})$. If there is a (AttIO) or (OracleCall) reduction in Tr with $\mu' = \mu_t$, we define E' to be the environment after that reduction. If $x \in \text{Dom}(E')$, we let $\text{Tidx}(Tr) = \{\epsilon\}$. Otherwise, $\text{Tidx}(Tr) = \emptyset$.

Let $\text{Tpp}(Tr) = \{\mu \mid \exists a \in \text{Tidx}(Tr), x[E(u_1[a]), \dots, E(u_m[a])]\}$ is defined just before μ in Tr . When sp is $\text{bit secr.}(x)$, $m = 0$, so this definition reduces to $\text{Tpp}(Tr) = \emptyset$ if $\text{Tidx}(Tr) = \emptyset$ and $\text{Tpp}(Tr) = \{\mu_\epsilon\}$ where x is defined just before μ_ϵ in Tr otherwise. We write $Tr \vdash sp$ when $Tr \vdash \{\text{prove}^{sp}(\text{Tpp}(Tr))\}$.

Lemma 36 *Consider a process Q_0 that satisfies Properties 3 and 4. Let sp be $1\text{-ses.secr.}(x)$, $\text{Secrecy}(x)$, or $\text{bit secr.}(x)$. Let C be an evaluation context acceptable for $C_{sp}[Q_0]$ with any public variables V ($x \notin V$) that does not contain \mathbb{S} nor $\bar{\mathbb{S}}$ and \mathcal{A} be an attacker in \mathcal{BB} . We have*

$$\Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \mathbb{S} \wedge sp] = \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \bar{\mathbb{S}} \wedge sp].$$

Proof of the cases $sp = 1\text{-ses.secr.}(x)$ and $sp = \text{Secrecy}(x)$ Let Tr be a full trace of $C[C_{sp}[Q_0]]$ for attacker \mathcal{A} such that $Tr \vdash sp$ and b is defined in Tr , that is, $b \in \text{Dom}(E_{Tr})$.

Let $E = E_{Tr}$. Let \tilde{i} be the current replication indices at the definition of x in Q_0 . For $j \in \text{Tidx}(Tr)$, let $\tilde{a}_j = E(u_1[j]), \dots, E(u_m[j])$, so that the test query at index j tests $x[\tilde{a}_j]$. Let Conf_j be the target configuration of the semantic rule that adds $x[\tilde{a}_j]$ to E , and μ_j be such that $x[\tilde{a}_j]$ is defined just before μ_j in Tr . So Conf_j is at program point μ_j in Tr . Let $z_j[\tilde{M}_j] = \text{defRand}_{\mu_j}(x)$ (which is always defined since $\mu_j \in \text{Tpp}(Tr)$ and $Tr \vdash sp$, so $Tr \vdash \{\text{prove}^{1\text{-ses.secr.}(x)}(\mu_j)\}$). Let \tilde{b}_j be such that $E, \{\tilde{i} \mapsto \tilde{a}_j\}, \tilde{M}_j \Downarrow \tilde{b}_j$. Then $x[\tilde{a}_j]$ is added to the environment E by (NewT), (LetT), (New), or (Let), we have $E(x[\tilde{a}_j]) = E(z_j[\tilde{b}_j])$, and $z_j[\tilde{b}_j]$ is chosen at random by (NewT) or (New) in Tr by definition of $\text{defRand}_{\mu_j}(x)$. Let us prove that, for all $j_1 \neq j_2$ in $\text{Tidx}(Tr)$, we have $z_{j_1} \neq z_{j_2}$ or $\tilde{b}_{j_1} \neq \tilde{b}_{j_2}$.

- When sp is $1\text{-ses.secr.}(x)$, this is trivially true since $\text{Tidx}(Tr)$ contains at most one element.
- When sp is $\text{Secrecy}(x)$, we have $\mu_{j_1}, \mu_{j_2} \in \text{Tpp}(Tr)$ and $Tr \vdash \text{Secrecy}(x)$, so we have $Tr \vdash \{\text{prove}^{\text{distinct}(x)}(\mu_{j_1}, \mu_{j_2})\}$, so $z_{j_1} \neq z_{j_2}$ or $Tr \vdash \forall \tilde{i}_1, \forall \tilde{i}_2, \neg \bigwedge \theta_1 \mathcal{F}_{\mu_{j_1}} \cup \theta_2 \mathcal{F}_{\mu_{j_2}} \cup \{\theta_1 \tilde{M}_{j_1} = \theta_2 \tilde{M}_{j_2}, \tilde{i}_1 \neq \tilde{i}_2\}$ where θ_1 and θ_2 are two distinct renamings of \tilde{i} to fresh replication indices, $\tilde{i}_1 = \theta_1 \tilde{i}$, and $\tilde{i}_2 = \theta_2 \tilde{i}$. In the latter case, let $\rho = \{\tilde{i}_1 \mapsto \tilde{a}_{j_1}, \tilde{i}_2 \mapsto \tilde{a}_{j_2}\}$. We have $\sigma_{\text{Conf}_{j_1}} = [\tilde{i} \mapsto \tilde{a}_{j_1}]$. By Corollary 3, $Tr, \rho \vdash \theta_1 \mathcal{F}_{\mu_{j_1}}$. Similarly, $Tr, \rho \vdash \theta_2 \mathcal{F}_{\mu_{j_2}}$. Since $j_1 \neq j_2$, $\tilde{a}_{j_1} \neq \tilde{a}_{j_2}$. (By construction of $\text{Tidx}(Tr)$, we consider only the first successful test query for a certain $x[\tilde{a}_j]$.) So $Tr, \rho \vdash \tilde{i}_1 \neq \tilde{i}_2$. Therefore, $Tr, \rho \vdash \theta_1 \tilde{M}_{j_1} \neq \theta_2 \tilde{M}_{j_2}$, so $\tilde{b}_{j_1} \neq \tilde{b}_{j_2}$.

For $j \in \text{Tidx}(Tr)$, let us choose elements v_j in T , where T is the type of x . Let us consider the following two sets of traces:

1. Tr modified by choosing $b = \text{true}$ and $z_j[\tilde{b}_j] = v_j$ for all $j \in \text{Tidx}(Tr)$. (The variable b is chosen and used in Q_{sp} . Note that the variable y of Q_{sp} is not defined when $b = \text{true}$. This set contains a single trace.)
2. Tr modified by choosing $b = \text{false}$ and $y[j] = v_j$ and $z_j[\tilde{b}_j] = v'_j$ for any $v'_j \in T$, for all $j \in \text{Tidx}(Tr)$. (This set contains $|T|^{|\text{Tidx}(Tr)|}$ traces.)

The trace Tr is one of the traces in these two sets: just choose the values of b , $z_j[\tilde{b}_j]$, and $y[j]$ if b is false that are used in Tr .

We show by induction on the derivation of these traces Tr_s (Tr_1 is in set 1 and Tr_2 is in set 2) that they have matching configurations $\text{Conf}'_s = E_s, \sigma, M_s, \mathcal{T}, \mu \mathcal{E} v_s$, $\text{Conf}'_s = Q_s, \mathcal{O} r$, or $\text{Conf}'_s = E_s, (\sigma, P_s) :: St_s, Q_s, \mathcal{O} r, \mathcal{T}, \mu \mathcal{E} v_s$ for $s \in \{1, 2\}$ that differ as follows:

- $E_1(b) = \text{true}$ while $E_2(b) = \text{false}$.
- $E_1(u'_s[j])$ when sp is $\text{Secrecy}(x)$ and $E_1(y[j])$ are undefined even when $E_2(u'_s[j])$ and $E_2(y[j])$ are defined for some indices $j \in \text{Tidx}(Tr)$.
- $E_1(z[\tilde{a}])$ differs from $E_2(z[\tilde{a}])$ for some z and \tilde{a} such that for some $\tilde{M}, \mathcal{I}, \mathcal{F}$, we have $Tr \vdash \{\{\text{noleak}(z[\tilde{M}], \mathcal{I}, \mathcal{F})\}\}$ and $Tr \vdash \exists \mathcal{I}, (\tilde{M} = \tilde{a}) \wedge \bigwedge \mathcal{F}$.
- Values inside M_1 and M_2 may differ when Conf'_s (for $s \in \{1, 2\}$) occurs in the derivation of

$$E_s, \sigma, \text{let } y[\tilde{i}] = M \text{ in } M', \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1}^* E_s[y[\sigma\tilde{i}] \mapsto a_s], \sigma, M', \mathcal{T}, \mu\mathcal{E}v_s$$

or of $E_s, (\sigma, \text{let } y[\tilde{i}] = M \text{ in } P) :: \mathcal{S}t_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1}^*$

$$E_s[y[\sigma\tilde{i}] \mapsto a_s], (\sigma, P) :: \mathcal{S}t_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s$$

where M is built from replication indices, variables, function applications, and conditionals and for some $\tilde{M}, \mathcal{I}, \mathcal{F}$, we have $Tr \vdash \{\{\text{noleak}(y[\tilde{M}], \mathcal{I}, \mathcal{F})\}\}$ and $Tr \vdash \exists \mathcal{I}, (\tilde{M} = \sigma\tilde{i}) \wedge \bigwedge \mathcal{F}$.

- Terms M_1 and M_2 may differ when Conf'_s (for $s \in \{1, 2\}$) occurs in the derivation of

$$E_s, \sigma, \text{event } e(\tilde{M}_s); M', \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1} E_s, \sigma, \text{event } e(\tilde{M}'_s); M', \mathcal{T}, \mu\mathcal{E}v_s$$

or of $E_s, (\sigma, \text{event } e(\tilde{M}_s); P) :: \mathcal{S}t_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1}$

$$E_s, (\sigma, \text{event } e(\tilde{M}'_s); P) :: \mathcal{S}t_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s$$

where the terms M_s match and the only rules above this reduction and under Conf'_s are (CtxT) with matching simple contexts any number of times followed by (CtxT) with context $\text{event } e(a_{s,1}, \dots, a_{s,k-1}, [], N_{k+1}, \dots, N_l); N$ or (Ctx) with context $\text{event } e(a_{s,1}, \dots, a_{s,k-1}, [], N_{k+1}, \dots, N_l); P$ once, where

- a *simple context* is a context of the form $x[a_1, \dots, a_{k-1}, [], N_{k+1}, \dots, N_m]$ or $f(a_1, \dots, a_{k-1}, [], N_{k+1}, \dots, N_m)$ for some $x, f, k, m, a_1, \dots, a_{k-1}, N_{k+1}, \dots, N_m$,
- simple contexts C_s (for $s \in \{1, 2\}$) *match* when $C_s = x[a_{s,1}, \dots, a_{s,k-1}, [], N_{k+1}, \dots, N_m]$ or $C_s = f(a_{s,1}, \dots, a_{s,k-1}, [], N_{k+1}, \dots, N_m)$ for some $x, f, k, m, a_{s,1}, \dots, a_{s,k-1}, N_{k+1}, \dots, N_m$, and
- terms M_s (for $s \in \{1, 2\}$) *match* when $M_s = a_{s,0}$, or $M_s = x[a_{s,1}, \dots, a_{s,k-1}, M'_s, N_{k+1}, \dots, N_m]$ or $M_s = f(a_{s,1}, \dots, a_{s,k-1}, M'_s, N_{k+1}, \dots, N_m)$ for some $x, f, k, m, a_{s,0}, a_{s,1}, \dots, a_{s,k-1}, N_{k+1}, \dots, N_m$ and matching M'_s .

- Terms M_1 and M_2 (resp. processes P_1 and P_2) may differ when

$$M_s = C_1[\dots C_k[\text{event } e(\tilde{M}_s); M'] \dots],$$

$$P_s = C_0[C_1[\dots C_k[\text{event } e(\tilde{M}_s); M'] \dots]],$$

or $P_s = \text{event } e(\tilde{M}_s); P$

where $k \in \mathbb{N}$, C_1, \dots, C_k are term contexts defined in Figure 7, C_0 is a process context defined in Figure 12, and the terms \tilde{M}_s match, for $s \in \{1, 2\}$.

- Arguments of events in $\mu\mathcal{E}v_1$ and $\mu\mathcal{E}v_2$ may differ.

- The events \mathbb{S} and $\bar{\mathbb{S}}$ are swapped: when $\mu\mathcal{E}v_1$ contains \mathbb{S} , $\mu\mathcal{E}v_2$ contains $\bar{\mathbb{S}}$, and conversely.
- Some additional configurations corresponding to the execution Q_{sp} differ.

The proof can be sketched as follows. The different choice of b leads to $E_1(b) = \text{true}$ and $E_2(b) = \text{false}$. The only semantic rule that reads the environment is (Var), when it evaluates an occurrence of the variable in question. By Definition 7, $b \notin \text{var}(Q_0) \cup V$, so the only occurrences of b are in Q_{sp} .

If the adversary calls oracle $O_s[\dots](\tilde{a})$ and $x[\tilde{a}]$ is not defined, then Q_{sp} simply yields. If the adversary calls oracle $O_s[\dots](\tilde{a})$ and $x[\tilde{a}]$ is defined, then $\tilde{a} = \tilde{a}_j$ for some $j \in \text{Tidx}(Tr)$. In set 1, $E_1(b) = \text{true}$, so Q_{sp} returns $E_1(x[\tilde{a}_j]) = E_1(z_j[\tilde{b}_j]) = v_j$. In set 2, $E_2(b) = \text{false}$, so when it is the first time that the adversary calls oracle $O_s[\dots](\tilde{a})$ and $x[\tilde{a}]$ is defined, Q_{sp} chooses a fresh $y[j]$ equal to v_j , and returns $E_2(y[j]) = v_j$; when the adversary calls again $O_s[\dots](\tilde{a})$, Q_{sp} finds $u'_s = j$ (by construction of $\text{Tidx}(Tr)$), and returns $E_2(y[j]) = v_j$. So in both sets, Q_{sp} returns the same value.

If the adversary calls $O'_s(b')$, then the result of the test $b' = b$ differs between set 1 and set 2, since $E_1(b) \neq E_2(b)$, so if set 1 executes \mathbb{S} , then set 2 executes $\bar{\mathbb{S}}$ and conversely. That is why the events \mathbb{S} and $\bar{\mathbb{S}}$ are swapped.

By Definition 7, $u'_s, y \notin \text{var}(Q_0) \cup V$, so the only occurrences of u'_s and y are in Q_{sp} . Therefore, the changes that come from differences in the definition of u'_s and y are already taken into account above.

The value of $E_1(z_j[\tilde{b}_j])$ is different from the one of $E_2(z_j[\tilde{b}_j])$. Since $Tr \vdash sp$, we have $Tr \vdash \{\{\text{prove}^{1\text{-ses.secr.}(x)}(\mu_j)\}\}$, so $Tr \vdash \{\{\text{noleak}(\theta z_j[\tilde{M}_j], \{\theta I_{\mu_j}\}, \theta \mathcal{F}_{\mu_j})\}\}$ where θ is a renaming of I_{μ_j} to fresh replication indices. We have $\sigma_{\text{Conf}_j} = [I_{\mu_j} \mapsto \tilde{a}_j]$. Let $\rho = \{\theta I_{\mu_j} \mapsto \tilde{a}_j\}$. By Corollary 3, $Tr, \rho \vdash \theta \mathcal{F}_{\mu_j}$. Moreover, $Tr, \rho \vdash \theta \tilde{M}_j = \tilde{b}_j$. So $Tr \vdash \exists \theta I_{\mu_j}, (\theta \tilde{M}_j = \tilde{b}_j) \wedge \bigwedge \theta \mathcal{F}_{\mu_j}$. Hence, for $\tilde{M} = \theta \tilde{M}_j$, $\mathcal{I} = \{\theta I_{\mu_j}\}$, and $\mathcal{F} = \theta \mathcal{F}_{\mu_j}$, we have $Tr \vdash \{\{\text{noleak}(z_j[\tilde{M}], \mathcal{I}, \mathcal{F})\}\}$ and $Tr \vdash \exists \mathcal{I}, (\tilde{M} = \tilde{b}_j) \wedge \bigwedge \mathcal{F}$.

The difference between $E_1(z[\tilde{a}])$ and $E_2(z[\tilde{a}])$ for z and \tilde{a} such that for some $\tilde{M}, \mathcal{I}, \mathcal{F}$, we have $Tr \vdash \{\{\text{noleak}(z[\tilde{M}], \mathcal{I}, \mathcal{F})\}\}$ and $Tr \vdash \exists \mathcal{I}, (\tilde{M} = \tilde{a}) \wedge \bigwedge \mathcal{F}$ has consequences when (Var) evaluates $z[\tilde{a}]$:

$$E_s, \sigma, {}^\mu z[\tilde{a}], \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1} E_s, \sigma, E_s(z[\tilde{a}]), \mathcal{T}, \mu\mathcal{E}v_s.$$

By Lemma 8, Property 4 applied to the configuration $\text{Conf} = E_s, \sigma, {}^\mu z[\tilde{a}], \mathcal{T}, \mu\mathcal{E}v_s$ with $l = 0$, we have

$$E'_s, \sigma, {}^\mu z[\tilde{M}'], \mathcal{T}', \mu\mathcal{E}v'_s \xrightarrow{1}^* E_s, \sigma, {}^\mu z[\tilde{a}], \mathcal{T}, \mu\mathcal{E}v_s$$

by any number of applications of (CtxT), where ${}^\mu z[\tilde{M}']$ is a subterm of $C[C_{sp}[Q_0]]$. Furthermore, if the evaluation of \tilde{M}' itself uses (Var) that evaluates a $z[\tilde{a}]$ that differs, we replace $z[\tilde{M}']$ by the smallest subterm of \tilde{M}' that evaluates a $z[\tilde{a}]$ that differs. By this replacement, we guarantee that the evaluation of \tilde{M}' proceeds in the same way in set 1 and set 2 and yields the same \tilde{a} . Recall that \tilde{M}' are simple terms by Invariants 2 and 5, so the evaluation of \tilde{M}' does not change $E_s, \mathcal{T}, \mu\mathcal{E}v_s$. So we have

$$E_s, \sigma, {}^\mu z[\tilde{M}'], \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1}^* E_s, \sigma, {}^\mu z[\tilde{a}], \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1} E_s, \sigma, E_s(z[\tilde{a}]), \mathcal{T}, \mu\mathcal{E}v_s$$

for $s \in \{1, 2\}$, where ${}^\mu z[\tilde{M}']$ is a subterm of $C[C_{sp}[Q_0]]$, by any number of applications of (CtxT) followed by one application of (Var). Let θ be a renaming of I_μ to fresh replication indices and $\rho = \{\theta I_\mu \mapsto \sigma I_\mu\}$. By Corollary 3, $Tr_s, \rho \vdash \theta \mathcal{F}_\mu$. Moreover, $E_s, \sigma, \tilde{M}' \Downarrow \tilde{a}$, so $E_s, \rho, \theta \tilde{M}' \Downarrow \tilde{a}$. Since E_{Tr_s} extends E_s , we have $Tr_s, \rho \vdash \theta \tilde{M}' = \tilde{a}$. Since Tr is among the traces Tr_s , we have

$Tr, \rho \vdash \theta\mathcal{F}_\mu$ and $Tr, \rho \vdash \theta\widetilde{M}' = \widetilde{a}$. There exists ρ' with domain \mathcal{I} such that $Tr, \rho' \vdash (\widetilde{M} = \widetilde{a}) \wedge \bigwedge \mathcal{F}$. So $Tr, \rho \cup \rho' \vdash \bigwedge (\mathcal{F} \cup \theta\mathcal{F}_\mu \cup \{\theta\widetilde{M}' = \widetilde{a}, \widetilde{M} = \widetilde{a}\})$. Since $Tr \vdash \{\{\text{noleak}(z[\widetilde{M}], \mathcal{I}, \mathcal{F})\}\}$ and $Tr \vdash \exists \mathcal{I}, (\widetilde{M} = \widetilde{a}) \wedge \bigwedge \mathcal{F}$, we have $Tr \vdash \neg(\forall \mathcal{I}, \neg \bigwedge \mathcal{F})$, so Tr satisfies the second disjunct of $\{\{\text{noleak}(z[\widetilde{M}], \mathcal{I}, \mathcal{F})\}\}$. Therefore, $z \notin V$, so the occurrence of ${}^\mu z[\widetilde{M}']$ evaluated above is either in Q_{sp} , and in this case z is actually x and this case has already been studied above, or in Q_0 . In the latter situation, we are in one of the following three cases:

- ${}^\mu z[\widetilde{M}']$ is in M in an assignment $\mu' \text{ let } y[\widetilde{i}] = M$ in Q_0 , M is built from replication indices, variables, function applications, and conditionals and $Tr \vdash \{\{\text{noleak}(y[\theta\widetilde{i}], \mathcal{I} \cup \{\theta\widetilde{i}\}, \mathcal{F} \cup \theta\mathcal{F}_\mu \cup \{\theta\widetilde{M}' = \widetilde{M}\})\}\}$ where θ is a renaming of \widetilde{i} to fresh replication indices. Let $\widetilde{M}'' = \theta\widetilde{i}$, $\mathcal{I}' = \mathcal{I} \cup \{\theta\widetilde{i}\}$, and $\mathcal{F}' = \mathcal{F} \cup \theta\mathcal{F}_\mu \cup \{\theta\widetilde{M}' = \widetilde{M}\}$. Since μ' is above μ , by Lemma 6, there is a configuration inside μ' before *Conf* in Tr . By Lemma 8 applied to that configuration (Property 1 when the assignment is a process, Property 4 with $l = 0$ when it is a term), the assignment $\mu' \text{ let } y[\widetilde{i}] = M$ is evaluated by

$$E_s, \sigma, \mu' \text{ let } y[\widetilde{i}] = M \text{ in } M', \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1}^* E_s[y[\sigma\widetilde{i}] \mapsto a_s], \sigma, M', \mathcal{T}, \mu\mathcal{E}v_s$$

or $E_s, (\sigma, \mu' \text{ let } y[\widetilde{i}] = M \text{ in } P) :: St_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1}^*$

$$E_s[y[\sigma\widetilde{i}] \mapsto a_s], (\sigma, P) :: St_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s$$

for $s \in \{1, 2\}$. We have $Tr \vdash \{\{\text{noleak}(y[\widetilde{M}''], \mathcal{I}', \mathcal{F}')\}\}$. Moreover $\widetilde{i} = I_\mu$ and $Tr, \rho \cup \rho' \vdash \theta\widetilde{i} = \sigma\widetilde{i}$ by definition of ρ . Hence $Tr, \rho \cup \rho' \vdash \widetilde{M}'' = \sigma\widetilde{i}$ and $Tr, \rho \cup \rho' \vdash \mathcal{F}'$, so $Tr \vdash \exists \mathcal{I}', (\widetilde{M}'' = \sigma\widetilde{i}) \wedge \bigwedge \mathcal{F}'$. Hence we are in a case in which different values inside terms M_1 and M_2 are allowed. Furthermore, the added values $E_s(y[\sigma\widetilde{i}]) = a_s$ may differ. Let $\widetilde{a}' = \sigma\widetilde{i}$. We have $Tr \vdash \{\{\text{noleak}(y[\widetilde{M}''], \mathcal{I}', \mathcal{F}')\}\}$ and $Tr \vdash \exists \mathcal{I}', (\widetilde{M}'' = \widetilde{a}') \wedge \bigwedge \mathcal{F}'$, so $E_1(y[\widetilde{a}'])$ is indeed allowed to differ from $E_2(y[\widetilde{a}'])$.

- ${}^\mu z[\widetilde{M}']$ is in μ' event $e(M_1, \dots, M_{k-1}, C[{}^\mu z[\widetilde{M}']], M_{k+1}, \dots, M_m)$ in Q_0 , for C defined in Figure 22. Since μ' is above μ , by Lemma 6, there is a configuration inside μ' before *Conf* in Tr . By Lemma 8 applied to that configuration (Property 1 when the event is a process, Property 4 with $l = 0$ when it is a term), the evaluation of the event starts from a configuration at μ' in Tr . The evaluation of event $e(M_1, \dots, M_{k-1}, C[{}^\mu z[\widetilde{M}']], M_{k+1}, \dots, M_m)$ first evaluates M_1, \dots, M_{k-1} to values using (CtxT) or (Ctx) with an event context. (If they evaluated to abort event values, $C[{}^\mu z[\widetilde{M}']]$ would not be evaluated.) Then it evaluates the context C : $y[\widetilde{i}] \stackrel{R}{\leftarrow} T$; C is evaluated by (NewT), let $y[\widetilde{i}] = M$ in C is evaluated by (LetT), if M then C else N' is evaluated by (IfT1) (M must evaluate to true because otherwise, ${}^\mu z[\widetilde{M}']$ would not be evaluated), if M then N else C is evaluated by (IfT2) (M must not evaluate to true because otherwise, ${}^\mu z[\widetilde{M}']$ would not be evaluated), event $e(\widetilde{M})$; C is evaluated by (EventT), and find contexts are evaluated by rules for find, until we reach

$$\text{event } e(a_{s,1}, \dots, a_{s,k-1}, C_{s,1}[\dots C_{s,l}[{}^\mu z[\widetilde{M}']] \dots], M_{k+1}, \dots, M_m)$$

for $s \in \{1, 2\}$, where $C_{s,1}, \dots, C_{s,l}$ are matching simple contexts. (Values may differ in case M_1, \dots, M_{k-1} , or terms in C contain other occurrences of variables whose value differs.) At this point, the reduction proceeds as follows:

$$E_s, \sigma, \text{event } e(\widetilde{M}_s); M', \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1} E_s, \sigma, \text{event } e(\widetilde{M}'_s); M', \mathcal{T}, \mu\mathcal{E}v_s$$

or $E_s, (\sigma, \text{event } e(\widetilde{M}_s); P), \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1} E_s, (\sigma, \text{event } e(\widetilde{M}'_s); P), \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s$

by (Var), (CtxT) with matching simple contexts any number of times followed by (CtxT) or (Ctx) with context $\text{event } e(a_{s,1}, \dots, a_{s,k-1}, [], M_{k+1}, \dots, M_m); \dots$ once, where

$$\begin{aligned} \widetilde{M}_s &= a_{s,1}, \dots, a_{s,k-1}, C_{s,1}[\dots C_{s,l}[\mu z[\widetilde{M}']] \dots], M_{k+1}, \dots, M_m \\ \text{and } \widetilde{M}'_s &= a_{s,1}, \dots, a_{s,k-1}, C_{s,1}[\dots C_{s,l}[a_s] \dots], M_{k+1}, \dots, M_m \end{aligned}$$

for $s \in \{1, 2\}$. Further reductions still manipulate configurations of the same form until the event itself is executed by (EventT) or (Event), which adds the event e with possibly different arguments to $\mu\mathcal{E}v_s$.

- $Tr \vdash \forall(\mathcal{I} \cup \theta I_\mu), \neg \wedge(\mathcal{F} \cup \theta \mathcal{F}_\mu \cup \{\theta \widetilde{M}' = \widetilde{M}\})$ where θ is a renaming of I_μ to fresh replication indices. We have $Tr, \rho \cup \rho' \vdash \neg \wedge(\mathcal{F} \cup \theta \mathcal{F}_\mu \cup \{\theta \widetilde{M}' = \widetilde{M}\})$. That yields a contradiction, so this case does not happen.

The sequence of events $\mu\mathcal{E}v_1$ (resp. $\mu\mathcal{E}v_2$) is never read by the semantic rules. It is only read by the distinguisher. Therefore, changes in this sequence of events do not modify the rest of the trace.

That concludes the proof that traces in set 1 and set 2 match.

Furthermore, the trace in set 1 and the traces in set 2 have the same probability. All full traces of $C[C_{sp}[Q_0]]$ for attacker \mathcal{A} that define b and that satisfy sp belong to set 1 or to set 2 for some Tr (for instance using the trace in question as Tr). Therefore, these sets form a partition of the full traces of $C[C_{sp}[Q_0]]$ for attacker \mathcal{A} that define b and that satisfy sp , and the sets that execute S have the same probability as the sets that execute \bar{S} . Moreover, the traces of $C[C_{sp}[Q_0]]$ that do not define b execute neither S nor \bar{S} . So $\Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : S \wedge sp] = \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \bar{S} \wedge sp]$. \square

Proof of the case $sp = \text{bit secr.}(x)$ Let Tr be a full trace of $C[C_{sp}[Q_0]]$ for attacker \mathcal{A} such that $Tr \vdash sp$. Let $E = E_{Tr}$.

If $\text{Tidx}(Tr) = \emptyset$, then Tr executes neither S nor \bar{S} .

Otherwise, $\text{Tidx}(Tr) = \{\epsilon\}$ and x is defined in Tr , that is, $x \in \text{Dom}(E)$. Let Conf_ϵ be the target configuration of the semantic rule that adds x to E , and μ_ϵ be such that x is defined just before μ_ϵ in Tr . So Conf_ϵ is at program point μ_ϵ in Tr . Let $z_\epsilon[\widetilde{M}_\epsilon] = \text{defRand}_{\mu_\epsilon}(x)$ (which is always defined since $\mu_\epsilon \in \text{Tpp}(Tr)$ and $Tr \vdash sp$, so $Tr \vdash \{\{\text{prove}^{1\text{-ses.secr.}(x)}(\mu_\epsilon)\}\}$). Let \widetilde{b}_ϵ be such that $E, \emptyset, \widetilde{M}_\epsilon \Downarrow \widetilde{b}_\epsilon$. Then x is added to the environment E by (NewT), (LetT), (New), or (Let), we have $E(x) = E(z_\epsilon[\widetilde{b}_\epsilon])$, and $z_\epsilon[\widetilde{b}_\epsilon]$ is chosen at random by (NewT) or (New) in Tr by definition of $\text{defRand}_{\mu_\epsilon}(x)$.

Let us consider the following two traces:

1. Tr_1 is Tr modified by choosing $z_\epsilon[\widetilde{b}_\epsilon] = \text{true}$.
2. Tr_2 is Tr modified by choosing $z_\epsilon[\widetilde{b}_\epsilon] = \text{false}$.

The trace Tr is one of these two traces: just choose the value $z_\epsilon[\widetilde{b}_\epsilon]$ that is used in Tr .

We show by induction on the derivation of these traces Tr_s that they have matching configurations $\text{Conf}'_s = E_s, \sigma, M_s, \mathcal{T}, \mu\mathcal{E}v_s$, $\text{Conf}'_s = Q_s, \mathcal{O}r$, or $\text{Conf}'_s = E_s, (\sigma, P_s) :: St_s, Q_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s$ for $s \in \{1, 2\}$ that differ as follows:

- $E_1(z[\widetilde{a}])$ differs from $E_2(z[\widetilde{a}])$ for some z and \widetilde{a} such that for some $\widetilde{M}, \mathcal{I}, \mathcal{F}$, we have $Tr \vdash \{\{\text{inleak}(z[\widetilde{M}], \mathcal{I}, \mathcal{F})\}\}$ and $Tr \vdash \exists \mathcal{I}, (\widetilde{M} = \widetilde{a}) \wedge \wedge \mathcal{F}$.

- Values inside M_1 and M_2 may differ when $Conf'_s$ (for $s \in \{1, 2\}$) occurs in the derivation of

$$E_s, \sigma, \text{let } y[\tilde{i}] = M \text{ in } M', \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1^*} E_s[y[\sigma\tilde{i}] \mapsto a_s], \sigma, M', \mathcal{T}, \mu\mathcal{E}v_s$$

or of $E_s, (\sigma, \text{let } y[\tilde{i}] = M \text{ in } P) :: \mathcal{S}t_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1^*}$

$$E_s[y[\sigma\tilde{i}] \mapsto a_s], (\sigma, P) :: \mathcal{S}t_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s$$

where M is built from replication indices, variables, function applications, and conditionals and for some $\widetilde{M}, \mathcal{I}, \mathcal{F}$, we have $Tr \vdash \{\text{noleak}(y[\widetilde{M}], \mathcal{I}, \mathcal{F})\}$ and $Tr \vdash \exists \mathcal{I}, (\widetilde{M} = \sigma\tilde{i}) \wedge \bigwedge \mathcal{F}$.

- Terms M_1 and M_2 may differ when $Conf'_s$ (for $s \in \{1, 2\}$) occurs in the derivation of

$$E_s, \sigma, \text{event } e(\widetilde{M}_s); M', \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1} E_s, \sigma, \text{event } e(\widetilde{M}'_s); M', \mathcal{T}, \mu\mathcal{E}v_s$$

or of $E_s, (\sigma, \text{event } e(\widetilde{M}_s); P) :: \mathcal{S}t_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s \xrightarrow{1}$

$$E_s, (\sigma, \text{event } e(\widetilde{M}'_s); P) :: \mathcal{S}t_s, \mathcal{Q}_s, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v_s$$

where the terms M_s match and the only rules above this reduction and under $Conf'_s$ are (CtxT) with matching simple contexts any number of times followed by (CtxT) with context $\text{event } e(a_{s,1}, \dots, a_{s,k-1}, [], N_{k+1}, \dots, N_l); N$ or (Ctx) with context $\text{event } e(a_{s,1}, \dots, a_{s,k-1}, [], N_{k+1}, \dots, N_l); P$ once, where simple contexts and matching are defined as in the cases $sp = \text{1-ses.secr.}(x)$ and $sp = \text{Secrecy}(x)$.

- Terms M_1 and M_2 (resp. processes P_1 and P_2) may differ when

$$M_s = C_1[\dots C_k[\text{event } e(\widetilde{M}_s); M'] \dots],$$

$$P_s = C_0[C_1[\dots C_k[\text{event } e(\widetilde{M}_s); M'] \dots]],$$

or $P_s = \text{event } e(\widetilde{M}_s); P$

where $k \in \mathbb{N}$, C_1, \dots, C_k are term contexts defined in Figure 7, C_0 is a process context defined in Figure 12, and the terms \widetilde{M}_s match, for $s \in \{1, 2\}$.

- Arguments of events in $\mu\mathcal{E}v_1$ and $\mu\mathcal{E}v_2$ may differ.
- The events S and \bar{S} are swapped: when $\mu\mathcal{E}v_1$ contains S , $\mu\mathcal{E}v_2$ contains \bar{S} , and conversely.
- Some additional configurations corresponding to the execution Q_{sp} differ.

The proof can be sketched as follows.

If the adversary calls $O''_s(b')$, either directly with rule (AttIO) or through the context with rule (OracleCall), then x is defined (since $\text{Tid}_x(Tr) = \{\epsilon\}$) and the result of the test $x = b'$ differs between Tr_1 and Tr_2 , since $E_1(x) = E_1(z_\epsilon[\tilde{b}_\epsilon]) = \text{true} \neq E_2(x) = E_2(z_\epsilon[\tilde{b}_\epsilon]) = \text{false}$, so if Tr_1 executes S , then Tr_2 executes \bar{S} and conversely. That is why the events S and \bar{S} are swapped.

The value of $E_1(z_\epsilon[\tilde{b}_\epsilon])$ is different from the one of $E_2(z_\epsilon[\tilde{b}_\epsilon])$. Since $Tr \vdash sp$, we have $Tr \vdash \{\text{prove}^{\text{1-ses.secr.}}(x)(\mu_\epsilon)\}$, so $Tr \vdash \{\text{noleak}(\theta z_\epsilon[\widetilde{M}_\epsilon], \{\theta I_{\mu_\epsilon}\}, \theta \mathcal{F}_{\mu_\epsilon})\}$ where θ is a renaming of I_{μ_ϵ} to fresh replication indices. Here, I_{μ_ϵ} is empty since x is defined under no replication. We have $\sigma_{Conf_\epsilon} = []$. Let $\rho = \emptyset$. By Corollary 3, $Tr, \rho \vdash \theta \mathcal{F}_{\mu_\epsilon}$. Moreover, $Tr, \rho \vdash \theta \widetilde{M}_\epsilon = \tilde{b}_\epsilon$. So $Tr \vdash \exists \theta I_{\mu_\epsilon}, (\theta \widetilde{M}_\epsilon = \tilde{b}_\epsilon) \wedge \bigwedge \theta \mathcal{F}_{\mu_\epsilon}$. Hence, for $\widetilde{M} = \theta \widetilde{M}_\epsilon$, $\mathcal{I} = \{\theta I_{\mu_\epsilon}\}$, and $\mathcal{F} = \theta \mathcal{F}_{\mu_\epsilon}$, we have $Tr \vdash \{\text{noleak}(z_\epsilon[\widetilde{M}], \mathcal{I}, \mathcal{F})\}$ and $Tr \vdash \exists \mathcal{I}, (\widetilde{M} = \tilde{b}_\epsilon) \wedge \bigwedge \mathcal{F}$.

The difference between $E_1(z[\tilde{a}])$ and $E_2(z[\tilde{a}])$ for z and \tilde{a} such that for some $\widetilde{M}, \mathcal{I}, \mathcal{F}$, we have $Tr \vdash \{\{\text{noleak}(z[\widetilde{M}], \mathcal{I}, \mathcal{F})\}\}$ and $Tr \vdash \exists \mathcal{I}, (\widetilde{M} = \tilde{a}) \wedge \bigwedge \mathcal{F}$ has consequences when (Var) evaluates $z[\tilde{a}]$, as in the cases $sp = 1\text{-ses.secr.}(x)$ and $sp = \text{Secrecy}(x)$. That concludes the proof that traces Tr_1 and Tr_2 match.

Furthermore, the traces Tr_1 and Tr_2 have the same probability. All full traces of $C[C_{sp}[Q_0]]$ for attacker \mathcal{A} such that Tidx is non-empty and that satisfy sp are Tr_1 or Tr_2 for some Tr (for instance using the trace in question as Tr). Therefore, Tr_1 and Tr_2 form a partition of the full traces of $C[C_{sp}[Q_0]]$ for attacker \mathcal{A} such that Tidx is non-empty and that satisfy sp , and half of these traces execute \mathbb{S} , the other half execute $\bar{\mathbb{S}}$. Moreover, the traces of $C[C_{sp}[Q_0]]$ such that Tidx is empty execute neither \mathbb{S} nor $\bar{\mathbb{S}}$. So $\Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \mathbb{S} \wedge sp] = \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \bar{\mathbb{S}} \wedge sp]$. \square

Proof of Proposition 1 Let sp be $1\text{-ses.secr.}(x)$, $\text{Secrecy}(x)$, or $\text{bit secr.}(x)$. Let C be an evaluation context acceptable for $C_{sp}[Q_0]$ with public variables V ($x \notin V$) that does not contain \mathbb{S} nor $\bar{\mathbb{S}}$ and \mathcal{A} be an attacker in \mathcal{BB} . We have

$$\begin{aligned}
\text{Adv}_{Q_0}^{sp}(\mathcal{A}, C) &= \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \mathbb{S}] - \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \bar{\mathbb{S}}] \\
&= \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \mathbb{S} \wedge sp] + \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \mathbb{S} \wedge \neg sp] \\
&\quad - \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \bar{\mathbb{S}} \wedge sp] - \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \bar{\mathbb{S}} \wedge \neg sp] \\
&= \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \mathbb{S} \wedge \neg sp] - \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \bar{\mathbb{S}} \wedge \neg sp] && \text{by Lemma 36} \\
&\leq \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \neg sp] \\
&\leq \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \neg\{\{\text{prove}^{sp}(\{\mu \mid \mu \text{ follows a definition of } x\})\}\}] \\
&\quad \text{because, for all } Tr, \text{Tpp}(Tr) \subseteq \{\mu \mid \mu \text{ follows a definition of } x\} \\
&\leq \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]]] \leq \neg\{\{\text{prove}^{sp}(\{\mu \mid \mu \text{ follows a definition of } x\})\}\}] && \text{by Lemma 1} \\
&\leq p(\mathcal{A}, C[C_{sp}[[]]]) = p'(\mathcal{A}, C)
\end{aligned}$$

So Q_0 satisfies sp with public variables V up to probability p' . Moreover,

$$\begin{aligned}
\text{Adv}_{Q_0}(\mathcal{A}, C[C_{sp}[[]], sp, D_{\text{false}}]) &= \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \mathbb{S}] - \Pr^{\mathcal{A}}[C[C_{sp}[Q_0]] : \bar{\mathbb{S}} \vee \text{NonUnique}_{Q_0}] \\
&\leq \text{Adv}_{Q_0}^{sp}(\mathcal{A}, C) \leq p(\mathcal{A}, C[C_{sp}[[]]])
\end{aligned}$$

so $\text{Bound}_{Q_0}(V \cup \{x\}, sp, D_{\text{false}}, p)$. \square

Example 5 As we shall see in Example 10, the oracle definition Q_0 of Example 1 can be transformed into the following oracle definition Q''_0 :

$$\begin{aligned}
Q''_0 &= \text{Start}() := x_k \stackrel{R}{\leftarrow} T_k; x_{mk} \stackrel{R}{\leftarrow} T_{mk}; \text{return } (); (Q''_A \mid Q''_B) \\
Q''_A &= \text{foreach } i \leq n \text{ do } O_A[i]() := x'_k \stackrel{R}{\leftarrow} T_k; x_r \stackrel{R}{\leftarrow} T_r; \\
&\quad \text{let } x_m : \text{bitstring} = \text{enc}'(Z_k, x_k, x_r) \text{ in} \\
&\quad \text{return } (x_m, \text{mac}'(x_m, x_{mk})) \\
Q''_B &= \text{foreach } i' \leq n \text{ do } O_B[i'](x'_m, x_{ma}) := \\
&\quad \text{find } u \leq n \text{ suchthat defined}(x_m[u], x'_k[u]) \wedge \\
&\quad \quad x'_m = x_m[u] \wedge \text{verify}'(x'_m, x_{mk}, x_{ma}) \text{ then} \\
&\quad \text{let } x''_k : T_k = x'_k[u] \text{ in return } ()
\end{aligned}$$

and $Q_0 \approx_p^{x''_k} Q'_0$. In order to prove the one-session secrecy of x''_k , we notice that x''_k is defined by let $x''_k : T_k = x'_k[u]$, the only variable access to x'_k in Q'_0 is let $x''_k : T_k = x'_k[u]$, and x''_k is not used in Q'_0 . So by Proposition 1, Q'_0 satisfies the one-session secrecy of x''_k without public variables up to probability 0. (We have $\text{defRand}_\mu(x''_k) = x'_k[u]$ and $\{\{\text{noleak}(x'_k[u], \mathcal{I}, \mathcal{F})\}\} = (\forall \mathcal{I}, \neg \wedge \mathcal{F}) \vee ((x'_k \notin V) \wedge \{\{\text{noleak}(x''_k[\theta i], \mathcal{I}', \mathcal{F}')\}\}) = \text{true}$ since $x'_k \notin V$, $x''_k \notin V$, and $\{\{\text{noleak}(x''_k[\theta i], \mathcal{I}', \mathcal{F}')\}\} = (\forall \mathcal{I}', \neg \wedge \mathcal{F}') \vee ((x''_k \notin V) \wedge \text{true}) = \text{true}$. So $\Pr^A[C[Q'_0]] \leq \neg\{\{\text{prove}^{1\text{-ses.secr.}(x''_k)}(\mathcal{S})\}\} = 0$.) By Lemma 21, the oracle definition Q_0 of Example 1 also satisfies the one-session secrecy of x''_k without public variables up to probability $p'(\mathcal{A}, C) = 2p(\mathcal{A}, C[C_{1\text{-ses.secr.}(x''_k)}[\]], t_S)$. However, this process does not preserve the secrecy of x''_k , because the adversary can force several sessions of B to use the same key x''_k , by replaying the message sent by A . (Accordingly, $\text{prove}^{\text{Secrecy}(x)}(\mathcal{S})$ is not satisfied.)

The criteria given in this section might seem restrictive, but in fact, they should be sufficient for all protocols, provided the previous transformation steps are powerful enough to transform the protocol into a simpler protocol, on which these criteria can then be applied.

4.2 Correspondences

4.2.1 Example

We illustrate the proof of correspondences on the following example, inspired by the corrected Woo-Lam public key protocol [73]:

$$\begin{aligned} B \rightarrow A &: (N, B) \\ A \rightarrow B &: \{pk_A, B, N\}_{sk_A} \end{aligned}$$

This protocol is a simple nonce challenge: B sends to A a fresh nonce N and its identity. A replies by signing the nonce N , B 's identity, and A 's public key (which we use here instead of A 's identity for simplicity: this avoids having to relate identities and keys; CryptoVerif can obviously also handle the version with A 's identity). The signatures are assumed to be (existentially) unforgeable under chosen message attacks (UF-CMA) [49], so, when B receives the signature, B is convinced that A is present. The signature cannot be a replay because the nonce N is signed.

In our calculus, this protocol is encoded by the following oracle definition G_0 , explained below:

$$\begin{aligned} G_0 &= O_0() := rk_A \stackrel{R}{\leftarrow} \text{keyseed}; \text{let } pk_A = \text{pkgen}(rk_A) \text{ in} \\ &\quad \text{let } sk_A = \text{skgen}(rk_A) \text{ in return } (pk_A); (Q_A \mid Q_B) \\ Q_A &= \text{foreach } i_A \leq n \text{ do } O_A[i_A](x_N : \text{nonce}, x_B : \text{host}) := \\ &\quad \text{event } e_A(pk_A, x_B, x_N); r \stackrel{R}{\leftarrow} \text{seed}; \\ &\quad \text{return } (\text{sign}(\text{concat}(pk_A, x_B, x_N), sk_A, r)) \\ Q_B &= \text{foreach } i_B \leq n \text{ do } O_{B1}[i_B](x_{pk_A} : \text{pkey}) := N \stackrel{R}{\leftarrow} \text{nonce}; \\ &\quad \text{return } (N, B); O_{B2}[i_B](s : \text{signature}) :=; \\ &\quad \text{if } \text{verify}(\text{concat}(x_{pk_A}, B, N), x_{pk_A}, s) \text{ then} \\ &\quad \text{if } x_{pk_A} = pk_A \text{ then event } e_B(x_{pk_A}, B, N) \end{aligned}$$

The oracle definition G_0 is assumed to run in interaction with an adversary, which also models the network. G_0 first receives a call to oracle O_0 without argument, from the adversary. Then, it chooses randomly with uniform probability a bitstring rk_A in the type *keyseed*, by the construct

$rk_A \stackrel{R}{\leftarrow} \text{keyseed}$. Then, G_0 generates the public key pk_A corresponding to the coins rk_A , by calling the public-key generation algorithm pkgen . Similarly, G_0 generates the secret key sk_A by calling skgen . It then returns the public key pk_A , so that the adversary has this public key.

After this return the control passes to the adversary. Several oracles are then made available, which represent the roles of A and B in the protocol: the oracle definition $Q_A \mid Q_B$ is the parallel composition of Q_A and Q_B ; it makes simultaneously available the oracles defined in Q_A and Q_B . Let Q'_A and Q'_B be such that $Q_A = \text{foreach } i_A \leq n \text{ do } Q'_A$ and $Q_B = \text{foreach } i_B \leq n \text{ do } Q'_B$. The replication $\text{foreach } i_A \leq n \text{ do } Q'_A$ represents n copies of the oracle definition Q'_A , indexed by the replication index i_A . The oracle definition Q'_A begins with oracle $O_A[i_A]$; the oracle is indexed with i_A so that the adversary can choose which copy of oracle O_A it calls by calling $O_A[i_A]$ for the appropriate value of i_A . The situation is similar for Q'_B , which expects a call to oracle $O_{B1}[i_B]$. The adversary can then run each copy of Q'_A or Q'_B simply by calling the appropriate oracles $O_A[i_A]$ or $O_{B1}[i_B]$.

The oracle definition Q'_B first expects a call to oracle $O_{B1}[i_B]$ with a message x_{pk_A} in the type pkey of public keys. This message is not really part of the protocol. It serves for starting a new session of the protocol, in which B interacts with the participant of public key x_{pk_A} . For starting a session between A and B , this message should be pk_A . Then, Q'_B chooses randomly with uniform probability a nonce N in the type nonce . The type nonce is *large*: collisions between independent random numbers chosen uniformly in a large type are eliminated by CryptoVerif. Q'_B returns the message (N, B) to the adversary and gives it back control. The attacker is expected to call oracle $O_A[i_A]$ with this message (N, B) , but may proceed differently in order to mount an attack against the protocol.

Upon a call to oracle $O_A[i_A]$ with message (x_N, x_B) , where the bitstring x_N is in the type nonce and x_B in the type host , the process Q'_A executes the event $e_A(pk_A, x_B, x_N)$. This event does not change the state of the system. Events just record that a certain program point has been reached, with certain values of the arguments of the event. Then, Q'_A chooses randomly with uniform probability a bitstring r in the type seed ; this random bitstring is next used as coins for the signature algorithm. Finally, Q'_A returns the signed message $\{pk_A, x_B, x_N\}_{sk_A}$. (The function concat concatenates its arguments, with information on the length of these arguments, so that the arguments can be recovered from the concatenation.) The control then passes to the attacker, which should call oracle $O_{B2}[i_B]$ with this message if it wishes to run the protocol correctly.

Upon a call to oracle $O_{B2}[i_B]$ with message s , Q'_B verifies that the signature s is correct and, if $x_{pk_A} = pk_A$, that is, if B runs a session with A , it executes the event $e_B(x_{pk_A}, B, N)$. Our goal is to prove that, if event e_B is executed, then event e_A has also been executed. However, when B runs a session with a participant other than A , it is perfectly correct that B terminates without event e_A being executed; that is why event e_B is executed only when B runs a session with A .

By the unforgeability of signatures, the signature verification with pk_A succeeds only for signatures generated with sk_A . So, when we verify that the signature is correct, we can furthermore check that it has been generated using sk_A . So, after game transformations explained below, we obtain the following final game:

$$G_1 = O_0() := rk_A \stackrel{R}{\leftarrow} \text{keyseed}; \\ \text{let } pk_A = \text{pkgen}'(rk_A) \text{ in return } (pk_A); (Q_{1A} \mid Q_{1B})$$

$$\begin{aligned}
Q_{1A} &= \text{foreach } i_A \leq n \text{ do } O_A[i_A](x_N : \text{nonce}, x_B : \text{host}) := \\
&\quad \text{event } e_A(pk_A, x_B, x_N); \\
&\quad \text{let } m = \text{concat}(pk_A, x_B, x_N) \text{ in} \\
&\quad r \stackrel{R}{\leftarrow} \text{seed}; \text{return } (\text{sign}'(m, \text{skgen}'(rk_A), r)) \\
Q_{1B} &= \text{foreach } i_B \leq n \text{ do } O_{B1}[i_B](x_{pk_A} : \text{pkey}) := N \stackrel{R}{\leftarrow} \text{nonce}; \\
&\quad \text{return } (N, B); O_{B2}[i_B](s : \text{signature}) := \\
&\quad \text{find } u \leq n \text{ suchthat defined}(m[u], x_B[u], x_N[u]) \\
&\quad \quad \wedge (x_{pk_A} = pk_A) \wedge (B = x_B[u]) \wedge (N = x_N[u]) \\
&\quad \quad \wedge \text{verify}'(\text{concat}(x_{pk_A}, B, N), x_{pk_A}, s) \text{ then} \\
&\quad \text{event } e_B(x_{pk_A}, B, N)
\end{aligned}$$

The assignment $sk_A = \text{skgen}(rk_A)$ has been removed and $\text{skgen}(rk_A)$ has been substituted for sk_A , in order to make the term $\text{sign}(m, \text{skgen}(rk_A), r)$ appear. This term is needed for the security of the signature scheme to apply.

In Q_{1A} , the signed message is stored in variable m , and this variable is used when computing the signature.

Finally, using the unforgeability of signatures, the signature verification has been replaced with an array lookup: the signature verification can succeed only when $\text{concat}(x_{pk_A}, B, N)$ has been signed with sk_A , so we look for the message $\text{concat}(x_{pk_A}, B, N)$ in the array m and the event e_B is executed only when this message is found. In other words, we look for an index $u \leq n$ such that $m[u]$ is defined and $m[u] = \text{concat}(x_{pk_A}, B, N)$. By definition of m , $m[u] = \text{concat}(pk_A, x_B[u], x_N[u])$, so the equality $m[u] = \text{concat}(x_{pk_A}, B, N)$ can be replaced with $(x_{pk_A} = pk_A) \wedge (B = x_B[u]) \wedge (N = x_N[u])$. (Recall that the result of the concat function contains enough information to recover its arguments.) This transformation replaces the function symbols pkgen , skgen , sign , and verify with primed function symbols pkgen' , skgen' , sign' , and verify' respectively, to avoid repeated applications of the unforgeability of signatures with the same key. (The unforgeability of signatures is applied only to unprimed symbols.)

The soundness of the game transformations shows that $G_0 \approx G_1$. We will prove that G_1 satisfies the correspondences (4) and (6) with any public variables V , in particular with $V = \emptyset$. By Lemma 27, G_0 also satisfies these correspondences with public variables $V = \emptyset$. Let us sketch how the proof of correspondence (4) for the game G_1 will proceed. Let Q'_{1A} and Q'_{1B} such that $Q_{1A} = \text{foreach } i_A \leq n \text{ do } Q'_{1A}$ and $Q_{1B} = \text{foreach } i_B \leq n \text{ do } Q'_{1B}$. Assume that event e_B is executed in the copy of Q'_{1B} of index i_B , that is, $e_B(x_{pk_A}[i_B], B, N[i_B])$ is executed. (Recall that the variables x_{pk_A}, N, u, \dots are implicitly arrays.) Then the condition of the find above e_B holds, that is, $m[u[i_B]], x_B[u[i_B]]$, and $x_N[u[i_B]]$ are defined, $x_{pk_A}[i_B] = pk_A$, $B = x_B[u[i_B]]$, and $N[i_B] = x_N[u[i_B]]$. Moreover, since $m[u[i_B]]$ is defined, the assignment that defines m has been executed in the copy of Q'_{1A} of index $i_A = u[i_B]$. Then the event $e_A(pk_A, x_B, x_N)$, located above the definition of m , must have been executed in that copy of Q'_{1A} , that is, $e_A(pk_A, x_B[u[i_B]], x_N[u[i_B]])$ has been executed. The equalities in the condition of the find imply that this event is also $e_A(x_{pk_A}[i_B], B, N[i_B])$. To sum up, if $e_B(x_{pk_A}[i_B], B, N[i_B])$ has been executed, then $e_A(x_{pk_A}[i_B], B, N[i_B])$ has been executed, so we have the correspondence (4). This reasoning is typical of the way the prover shows correspondences. In particular, the conditions of array lookups are key in these proofs, because they allow us to relate values in processes that run in parallel (here, the processes that represent A and B), and interesting correspondences relate events that occur in such processes. Next, we detail and formalize this reasoning, both for non-injective and injective correspondences.

4.2.2 Non-unique Events

The only correspondence that involves a non-unique event e is $\text{event}(e) \Rightarrow \text{false}$, and it is simply proved by noticing that the event e no longer occurs in the game after the transformation **prove_unique** (Section 5.1.4). Therefore, non-unique events are not concerned by the proofs of Sections 4.2.3 and 4.2.4.

4.2.3 Non-injective Correspondences

Intuitively, in order to prove that Q_0 satisfies a non-injective correspondence $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$, with $\tilde{x} = \text{var}(\psi)$ and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$, we collect all facts that hold at events in ψ and show that these facts imply ϕ using the equational prover.

When **casesInCorresp** = **false**, CryptoVerif uses \mathcal{F}_μ to collect these facts. When **casesInCorresp** = **true** (the default), it uses $\mathcal{F}_{\mu,c}$ for more precision. In this section, we detail the proof with $\mathcal{F}_{\mu,c}$. The usage of \mathcal{F}_μ can be considered as using a single case c , relying on Lemma 31 instead of Lemma 32. Formally, we collect facts that hold when the event F in ψ has been executed, as follows.

Definition 16 (μ executes F , $\mathcal{F}_{F,\mu,c}$) When $F = \text{event}(e(M_1, \dots, M_m))$ and ${}^\mu\text{event } e(M'_1, \dots, M'_m); \dots$ occurs in Q_0 or, for $m = 0$, ${}^\mu\text{event_abort } e$ or ${}^\mu\text{find}[\text{unique}_e] \dots$ occurs in Q_0 , we say that μ executes F .

If μ executes F and for all ${}^\mu\text{event } e(M'_1, \dots, M'_m); \dots$ in Q_0 , M'_1, \dots, M'_m are simple terms, then we define $\mathcal{F}_{F,\mu,c}^0 = \mathcal{F}_{\mu,c} \cup \{M'_j = M_j \mid j \leq m\} \cup \{\text{lastdefprogrampoint}(\mu, I_\mu) \text{ if } {}^\mu\text{event_abort } e \text{ or } {}^\mu\text{find}[\text{unique}_e] \dots \text{ occurs in } Q_0\}$. If additionally F is not a non-unique event, then we define $\mathcal{F}_{F,\mu,c} = \mathcal{F}_{F,\mu,c}^0 \cup \mathcal{F}_\mu^{\text{Fut}}$.

Intuitively, when the event F in ψ has been executed, it has been executed by some subterm or subprocess of Q_0 , so there exists a subterm or subprocess ${}^\mu\text{event } e(M'_1, \dots, M'_m); \dots$ or, for $m = 0$, ${}^\mu\text{event_abort } e$ or ${}^\mu\text{find}[\text{unique}_e] \dots$ in Q_0 such that, the event $e(M'_1, \dots, M'_m)$ has been executed and it is equal to the event F , hence $M'_j = M_j$ holds for $j \leq m$. Moreover, since the program point μ , which executes F , has been reached, $\mathcal{F}_{\mu,c}$ holds for some case c (Lemma 32). Furthermore, when the event aborts, it is the last step of the trace, so $\text{lastdefprogrampoint}(\mu, I_\mu)$ also holds. Hence $\mathcal{F}_{F,\mu,c}^0$ holds for some case c . Additionally, assuming we consider traces that do not execute non-unique events, since the adversary cannot stop execution of the process until the next return, yield, or $\text{event_abort } e$, $\mathcal{F}_\mu^{\text{Fut}}$ also holds (Lemma 33), so $\mathcal{F}_{F,\mu,c}$ holds for some case c . This is proved more formally in Lemma 38 below. (The case of $\text{get}[\text{unique}_e]$ is not mentioned in Definition 16 because it is excluded by Property 3.)

We restrict ourselves to the case in which M'_1, \dots, M'_m are simple terms because only simple terms allowed in sets of facts.

Let θ be a substitution equal to the identity on the variables \tilde{x} of ψ . This substitution gives values to existentially quantified variables \tilde{y} of ϕ . We say that $\mathcal{F} \Vdash_\theta \phi$ when we can show that \mathcal{F} implies $\theta\phi$. Formally, we define:

$\mathcal{F} \Vdash_\theta M$ if and only if $\mathcal{F} \cup \{-\theta M\}$ yields a contradiction

$\mathcal{F} \Vdash_\theta \text{event}(e(M_1, \dots, M_m))$ if and only if there exist M'_0, \dots, M'_m such that $M'_0 : \text{event}(e(M'_1, \dots, M'_m)) \in \mathcal{F}$ and $\mathcal{F} \cup \{\bigvee_{j=1}^m \theta M_j \neq M'_j\}$ yields a contradiction

$\mathcal{F} \Vdash_\theta \phi_1 \wedge \phi_2$ if and only if $\mathcal{F} \Vdash_\theta \phi_1$ and $\mathcal{F} \Vdash_\theta \phi_2$

$\mathcal{F} \Vdash_\theta \phi_1 \vee \phi_2$ if and only if $\mathcal{F} \Vdash_\theta \phi_1$ or $\mathcal{F} \Vdash_\theta \phi_2$

Terms θM are proved by contradiction, using the equational prover. Events θF are proved by looking for some event F' in \mathcal{F} and showing by contradiction that $\theta F = F'$, using the equational prover.

Let $\varphi = \llbracket \forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket$ be a non-injective correspondence that does not use non-unique events, with $\psi = F_1 \wedge \dots \wedge F_m$, $\tilde{x} = \text{var}(\psi)$, and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$. Suppose that, in Q_0 , the arguments of the events that occur in ψ are always simple terms. Suppose that, for all $j \leq m$, μ_j that executes F_j and c_j is a case for \mathcal{F}_{μ_j, c_j} . For $j \leq m$, let θ_j be a renaming of I_{μ_j} to fresh replication indices. (The renamings θ_j have pairwise disjoint images.) Let θ be a family parameterized by $\mu_1, c_1, \dots, \mu_m, c_m$ of substitutions equal to the identity on \tilde{x} . We define $\text{prove}^\varphi(\theta, \mu_1, c_1, \dots, \mu_m, c_m) = (\theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m} \xrightarrow{\theta_{(\mu_1, c_1, \dots, \mu_m, c_m)}} \phi)$. This function defines the algorithm that we use to prove the correspondence φ assuming for all $j \leq m$, F_j is executed in μ_j and we are in case c_j . We also define $\text{prove}^\varphi(\theta, \mathcal{S}) = \bigwedge_{(\mu_1, c_1, \dots, \mu_m, c_m) \in \mathcal{S}} \text{prove}^\varphi(\theta, \mu_1, c_1, \dots, \mu_m, c_m)$.

Non-injective correspondences are proved as follows.

Proposition 2 *Let $\varphi = \llbracket \forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket$ be a non-injective correspondence that does not use non-unique events, with $\psi = F_1 \wedge \dots \wedge F_m$, $\tilde{x} = \text{var}(\psi)$, and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$. Let Q_0 be a process that satisfies Properties 3 and 4. Suppose that, in Q_0 , the arguments of the events that occur in ψ are always simple terms. Let $\mathcal{S} = \{(\mu_1, c_1, \dots, \mu_m, c_m) \mid \forall j \leq m, \mu_j \text{ executes } F_j \text{ and } c_j \text{ is a case for } \mathcal{F}_{\mu_j, c_j}\}$. If there exists a family of substitutions θ equal to the identity on \tilde{x} such that $\text{prove}^\varphi(\theta, \mathcal{S})$ and for all $\mathcal{A} \in \mathcal{BB}$ and all evaluation contexts C acceptable for Q_0 , $\text{Pr}^{\mathcal{A}}[C[Q_0] \preceq \neg \llbracket \text{prove}^\varphi(\theta, \mathcal{S}) \rrbracket] \leq p(\mathcal{A}, C)$, then $\text{Bound}_{Q_0}(V, \varphi, D_{\text{false}}, p)$ for any V .*

Intuitively, when $\psi = F_1 \wedge \dots \wedge F_m$ holds, $\theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m}$ hold for some $\mu_1, c_1, \dots, \mu_m, c_m$. For some θ equal to the identity on ψ , $\theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m}$ implies $\theta\phi$, so $\theta\phi$ holds. Hence the correspondence is satisfied. The proof of Proposition 2 relies on the following properties and lemmas. We have

$$\begin{aligned} \{\{\mathcal{F} \mapsto_\theta M\}\} &= \forall \tilde{z} \in \tilde{T}'', \neg \left(\bigwedge \mathcal{F} \wedge \neg \theta M \right) \\ \{\{\mathcal{F} \mapsto_\theta \text{event}(e(M_1, \dots, M_m))\}\} &= \forall \tilde{z} \in \tilde{T}'', \neg \left(\bigwedge \mathcal{F} \wedge \bigvee_{j=1}^m \theta M_j \neq M'_j \right) \\ &\quad \text{for some } M'_0 : \text{event}(e(M'_1, \dots, M'_m)) \in \mathcal{F} \\ \{\{\mathcal{F} \mapsto_\theta \phi_1 \wedge \phi_2\}\} &= \{\{\mathcal{F} \mapsto_\theta \phi_1\}\} \wedge \{\{\mathcal{F} \mapsto_\theta \phi_2\}\} \\ \{\{\mathcal{F} \mapsto_\theta \phi_1 \vee \phi_2\}\} &= \begin{cases} \{\{\mathcal{F} \mapsto_\theta \phi_1\}\} & \text{if } \mathcal{F} \mapsto_\theta \phi_1 \\ \{\{\mathcal{F} \mapsto_\theta \phi_2\}\} & \text{otherwise} \end{cases} \end{aligned}$$

where \tilde{z} are the non-process variables in \mathcal{F} , in the image of θ , and in \tilde{x} , and \tilde{T}'' are their types.

Lemma 37 $\{\{\mathcal{F} \mapsto_\theta \phi\}\} \Rightarrow \forall \tilde{z} \in \tilde{T}'', \neg \left(\bigwedge \mathcal{F} \wedge \neg \theta\phi \right)$, where \tilde{z} are the non-process variables in \mathcal{F} , in the image of θ , and in \tilde{x} , and \tilde{T}'' are their types.

Proof By induction on ϕ . □

Lemma 38 *Let Q_0 be a process that satisfies Properties 3 and 4. Let Tr be a full trace of $C[Q_0]$ for some attacker $\mathcal{A} \in \mathcal{BB}$. Let $\mu\mathcal{E}v$ be the sequence of events in the last configuration of Tr . Let $F = \text{event}(e(\tilde{M}))$ where \tilde{M} is a tuple of terms and e is an event that does not occur in C . Suppose that the arguments of e in Q_0 are always simple terms. Let ρ be a mapping of the variables of \tilde{M} and τ to their values. Suppose that $\text{Tr}, \rho \vdash F @ \tau$.*

Then there exist a program point μ (in Q_0) that executes F and a case c such that, for any θ' renaming of I_μ to fresh replication indices, there exists a mapping σ with domain $\theta' I_\mu$ such that $\mu\mathcal{E}v(\rho(\tau)) = (\mu, \sigma(\theta' I_\mu)) : e(\dots)$ and $Tr, \sigma \cup \rho \vdash \theta' \mathcal{F}_{F, \mu, c}^0$. If additionally, Tr does not execute a non-unique event of Q_0 , then $Tr, \sigma \cup \rho \vdash \theta' \mathcal{F}_{F, \mu, c}$.

Proof Let $E = E_{Tr}$. Since $Tr, \rho \vdash F@_\tau$ and the variables of \widetilde{M} and τ are defined in ρ , there exists \widetilde{a} such that $\rho, \widetilde{M} \Downarrow \widetilde{a}$ and $\mu\mathcal{E}v(\rho(\tau)) = (\mu, \widetilde{a}_0) : e(\widetilde{a})$ for some μ and \widetilde{a}_0 . The rule of the semantics that may have added this element to $\mu\mathcal{E}v$ is (Event), (EventAbort), (CtxEvent), (FindE), (Find3), (GetE), (Get3), or (EventT).

- In case (Event), the initial configuration of the rule (Event) is of the form

$$E_1, (\sigma_1, {}^\mu\text{event } e(\widetilde{a}); P) :: \mathcal{S}t_0, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_1, \mu\mathcal{E}v_1.$$

By Lemma 8, Property 1 applied to this configuration, we have reductions

$$\begin{aligned} \text{Conf} &= E_0, (\sigma_0, {}^\mu\text{event } e(\widetilde{M}'); P) :: \mathcal{S}t_0, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{p_0}_{t_0} \dots \xrightarrow{p_1}_{t_1} E_1, (\sigma_1, {}^\mu\text{event } e(\widetilde{a}); P) :: \mathcal{S}t_0, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_1, \mu\mathcal{E}v_1 \\ &\xrightarrow{1} E_1, (\sigma_1, P) :: \mathcal{S}t_0, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_1, (\mu\mathcal{E}v_1, (\mu, \text{Im}(\sigma_1)) : e(\widetilde{a})) \end{aligned}$$

where ${}^\mu\text{event } e(\widetilde{M}'); P$ is a subprocess of $C[Q_0]$ up to renaming of oracles, by any number of applications of (Ctx) and a final application of (Event).

- In case (EventAbort), (Find3), or (Get3), the rules that can conclude with a process ${}^\mu P$ with $P = \text{event_abort } e$, $P = \text{find}[\text{unique}_e] \dots$, or $P = \text{get}[\text{unique}_e] \dots$ are (New), (Let), (If1), (If2), (Find1), (Find2), (Insert), (Get1), (Get2), (OracleCall), (Return), (Yield), (AttIO), (AttYield), (Event) and by Corollary 1, their target process is a subprocess of $C[Q_0]$ up to renaming of oracles. So we have a reduction

$$\begin{aligned} \text{Conf} &= E_0, (\sigma_0, {}^\mu P) :: \mathcal{S}t_0, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{p}_t E_0, (\sigma_0, \text{abort}) :: \mathcal{S}t_0, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, (\mu\mathcal{E}v_0, (\mu, \text{Im}(\sigma_0)) : e) \end{aligned}$$

where ${}^\mu P$ is a subprocess of $C[Q_0]$ up to renaming of oracles, by (EventAbort), (Find3), or (Get3).

- In case (EventT), the initial configuration of the rule (EventT) is of the form $E_1, \sigma_1, {}^\mu\text{event } e(\widetilde{a}); N, \mathcal{T}_1, \mu\mathcal{E}v_1$. By Lemma 8, Property 4 applied to this configuration with $l = 0$, we have reductions

$$\begin{aligned} \text{Conf} &= E_0, \sigma_0, {}^\mu\text{event } e(\widetilde{M}'); N, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{p_0}_{t_0} \dots \xrightarrow{p_1}_{t_1} E_1, \sigma_1, {}^\mu\text{event } e(\widetilde{a}); N, \mathcal{T}_1, \mu\mathcal{E}v_1 \\ &\xrightarrow{1} E_1, \sigma_1, N, \mathcal{T}_1, (\mu\mathcal{E}v_1, (\mu, \text{Im}(\sigma_1)) : e(\widetilde{a})) \end{aligned}$$

where ${}^\mu\text{event } e(\widetilde{M}'); N$ is a subterm of $C[Q_0]$, by any number of applications of (CtxT) and a final application of (EventT).

- In case (CtxEvent), the only rule that can conclude with a process $C[\text{event_abort } (\mu, \widetilde{a}_0) : e]$ is (Ctx). (The rules (New), (Let), (If1), (If2), (Find1), (Find2), (Insert), (Get1), (Get2), (OracleCall), (Return), (Yield), (AttIO), (AttYield), (Event) cannot conclude with

$C[\text{event_abort } (\mu, \tilde{a}_0) : e]$ because, by Corollary 1, their target process is a subprocess of $C[Q_0]$ up to renaming of oracles, and the initial process $C[Q_0]$ does not contain the abort event value $\text{event_abort } (\mu, \tilde{a}_0) : e$. Hence, there is a rule that concludes with a term $\text{event_abort } (\mu, \tilde{a}_0) : e$.

In cases (FindE) and (GetE), there is also a rule that concludes with a term $\text{event_abort } (\mu, \tilde{a}_0) : e$.

The only rules that conclude with a term $\text{event_abort } (\mu, \tilde{a}_0) : e$ are (FindTE), (FindT3), (GetTE), (GetT3), (EventAbortT), and (CtxEventT). (It cannot be (NewT), (LetT), (IfT1), (IfT2), (FindT1), (FindT2), (InsertT), (GetT1), (GetT2), (EventT), (DefinedYes) because, by Corollary 1, their target term is a subterm of $C[Q_0]$, and the initial process $C[Q_0]$ does not contain the abort event value $\text{event_abort } (\mu, \tilde{a}_0) : e$.) In cases (FindTE) and (GetTE), there is recursively another rule that concludes with $\text{event_abort } (\mu, \tilde{a}_0) : e$. In case (CtxEventT), the only rule that can conclude with $C[\text{event_abort } (\mu, \tilde{a}_0) : e]$ is (CtxT), so there is recursively another rule that concludes with $\text{event_abort } (\mu, \tilde{a}_0) : e$. Therefore, $\text{event_abort } (\mu, \tilde{a}_0) : e$ ultimately comes from an application of (EventAbortT), (FindT3), or (GetT3):

$$\text{Conf} = E_0, \sigma_0, {}^\mu N, \mathcal{T}_0, \mu \mathcal{E}v_0 \xrightarrow{t} E_0, \sigma_0, \text{event_abort } (\mu, \text{Im}(\sigma_0)) : e, \mathcal{T}_0, \mu \mathcal{E}v_0$$

where $N = \text{event_abort } e$, $N = \text{find}[\text{unique}_e] \dots$, or $N = \text{get}[\text{unique}_e] \dots$. Furthermore, the only rules that can conclude with such a term ${}^\mu N$ are (NewT), (IfT1), (IfT2), (LetT), (FindT1), (FindT2), (InsertT), (GetT1), (GetT2), (EventT), (DefinedYes) and, by Corollary 1, their target term is a subterm of $C[Q_0]$.

In all cases, since e does not occur in C , μ is in fact a program point of Q_0 . Therefore, the process or term at program point μ in Q_0 is of the form ${}^\mu \text{event } e(\tilde{M}'); \dots, {}^\mu \text{event_abort } e, {}^\mu \text{find}[\text{unique}_e] \dots, \text{ or } {}^\mu \text{get}[\text{unique}_e] \dots$. In cases (Event) and (EventT), we have $\sigma_0 = \sigma_1$ by Lemma 3. In all cases, $\text{Dom}(\sigma_0) = I_\mu$ are the current replication indices at program point μ by Lemma 2, and $\text{Im}(\sigma_0) = \tilde{a}_0$, so $\sigma_0 = [I_\mu \mapsto \tilde{a}_0]$. Moreover, \tilde{M}' are simple terms, so their evaluation can be written $E_0, \{I_\mu \mapsto \tilde{a}_0\}, \tilde{M}' \Downarrow \tilde{a}$. Let $\sigma = \{\theta' I_\mu \mapsto \tilde{a}_0\}$. We have $\mu \mathcal{E}v(\rho(\tau)) = (\mu, \sigma(\theta' I_\mu)) : e(\tilde{a})$.

We have $E_0, \sigma_0, \tilde{M}' \Downarrow \tilde{a}$ so $E_0, \sigma, \theta' \tilde{M}' \Downarrow \tilde{a}$. The environment E_{Tr} extends E_0 , so $E_{Tr}, \sigma, \theta' \tilde{M}' \Downarrow \tilde{a}$, so $Tr, \sigma \cup \rho \vdash \theta' \tilde{M}' = M$.

The configuration Conf is at program point μ in Tr , so by Corollary 3, we have $Tr, \sigma \vdash \theta' \mathcal{F}_{\mu, c}$.

When the process or term at μ is ${}^\mu \text{event_abort } e, {}^\mu \text{find}[\text{unique}_e] \dots, \text{ or } {}^\mu \text{get}[\text{unique}_e] \dots$, the environment and replication indices in Conf are the same as at the end of the trace, since the execution after Conf applies (EventAbort), (Find3), or (Get3), which terminate the trace keeping the same environment and replication indices as in Conf or (EventAbortT), (FindT3), or (GetT3) which build an abort event value keeping the same environment and replication indices as in Conf , followed by some rules among (FindTE), (GetTE), (CtxT), (CtxEventT), (FindE), (GetE), (Ctx), and (CtxEvent), which preserve the environment and replication indices that come with the abort event value. So $Tr, \sigma \vdash \text{lastdefprogrampoint}(\mu, \theta' I_\mu)$.

Therefore, μ executes F and $Tr, \sigma \cup \rho \vdash \theta' \mathcal{F}_{F, \mu, c}^0$ for some c .

Suppose additionally that Tr does not execute any non-unique event of Q_0 . Let Tr'' be the prefix of Tr that stops at the first evaluated return $\text{return } (b); Q$ or yield that follows Conf , or $Tr'' = Tr$ if Tr contains no return nor yield after Conf . By Lemma 33, we have $Tr'' \vdash \mathcal{F}_{\mu}^{\text{Fut}}$, so $Tr'', \sigma \vdash \theta' \mathcal{F}_{\mu}^{\text{Fut}}$. Moreover, E extends $E_{Tr''}$ and by Lemma 3, $\mu \mathcal{E}v_{Tr}$ extends $\mu \mathcal{E}v_{Tr''}$, so $Tr, \sigma \vdash \theta' \mathcal{F}_{\mu}^{\text{Fut}}$. Therefore, $Tr, \sigma \cup \rho \vdash \theta' \mathcal{F}_{F, \mu, c}$ for some c . \square

Lemma 39 Let $\varphi = \llbracket \forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket$ be a non-injective correspondence that does not use non-unique events, with $\psi = F_1 \wedge \dots \wedge F_m$, $\tilde{x} = \text{var}(\psi)$, and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$. Let Q_0 be a process that satisfies Properties 3 and 4. Suppose that, in Q_0 , the arguments of the events that occur in ψ are always simple terms.

Let $\mathcal{S} = \{(\mu_1, c_1, \dots, \mu_m, c_m) \mid \forall j \leq m, \mu_j \text{ executes } F_j \text{ and } c_j \text{ is a case for } \mathcal{F}_{\mu_j, c_j}\}$. Let C be an evaluation context acceptable for Q_0 with public variables V that does not contain events used by φ and \mathcal{A} be an attacker in \mathcal{BB} . Let Tr be a full trace of $C[Q_0]$ for \mathcal{A} that does not execute any non-unique event of Q_0 . If $Tr \vdash \neg\varphi$, then for any θ family of substitutions equal to the identity on \tilde{x} , $Tr \vdash \neg\{\text{prove}^\varphi(\theta, \mathcal{S})\}$.

Proof Since $Tr \vdash \neg\varphi$, we have $Tr \vdash \exists \tilde{x} \in \tilde{T}, F_1 \wedge \dots \wedge F_m \wedge \forall \tilde{y} \in \tilde{T}', \neg\phi$. So there exists ρ that maps \tilde{x} to elements of \tilde{T} such that $Tr, \rho \vdash F_1 \wedge \dots \wedge F_m \wedge \forall \tilde{y} \in \tilde{T}', \neg\phi$. By Lemma 38, for all $j \leq m$, there exists a program point μ_j (in Q_0) that executes F_j and a case c_j such that, for any θ_j renaming of I_{μ_j} to fresh replication indices, there exists a mapping σ_j with domain $\theta_j I_{\mu_j}$ such that $Tr, \sigma_j \cup \rho \vdash \theta_j \mathcal{F}_{F_j, \mu_j, c_j}$. Since $Tr, \rho \vdash \forall \tilde{y} \in \tilde{T}', \neg\phi$, we have $Tr, \rho \vdash \neg\theta(\mu_1, c_1, \dots, \mu_m, c_m)\phi$. Therefore, $Tr, \sigma_1 \cup \dots \cup \sigma_m \cup \rho \vdash \theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m} \cup \{\neg\theta(\mu_1, c_1, \dots, \mu_m, c_m)\phi\}$, so $Tr \vdash \exists \theta_1 I_{\mu_1}, \dots, \exists \theta_m I_{\mu_m}, \exists \tilde{x} \in \tilde{T}, \wedge \mathcal{F}_0(\mu_1, c_1, \dots, \mu_m, c_m) \wedge \neg\theta(\mu_1, c_1, \dots, \mu_m, c_m)\phi$ where $\mathcal{F}_0(\mu_1, c_1, \dots, \mu_m, c_m) = \theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m}$. We have

$$\begin{aligned} & \{\text{prove}^\varphi(\theta, \mu_1, c_1, \dots, \mu_m, c_m)\} \\ & \Rightarrow \forall \theta_1 I_{\mu_1}, \dots, \forall \theta_m I_{\mu_m}, \forall \tilde{x} \in \tilde{T}, \neg(\wedge \mathcal{F}_0(\mu_1, c_1, \dots, \mu_m, c_m) \wedge \neg\theta(\mu_1, c_1, \dots, \mu_m, c_m)\phi) \end{aligned}$$

by Lemma 37. (The non-process variables in $\mathcal{F}_0(\mu_1, c_1, \dots, \mu_m, c_m)$, in the image of $\theta(\mu_1, c_1, \dots, \mu_m, c_m)$, and in \tilde{x} are in $\theta_1 I_{\mu_1}, \dots, \theta_m I_{\mu_m}, \tilde{x}$.) So $Tr \vdash \neg\{\text{prove}^\varphi(\theta, \mu_1, c_1, \dots, \mu_m, c_m)\}$, so $Tr \vdash \neg\{\text{prove}^\varphi(\theta, \mathcal{S})\}$. \square

Proof of Proposition 2 Let C be an evaluation context acceptable for Q_0 with public variables V that does not contain events used by φ and \mathcal{A} be an attacker in \mathcal{BB} . We have

$$\begin{aligned} \text{Adv}_{Q_0}(\mathcal{A}, C, \varphi, D_{\text{false}}) &= \Pr^{\mathcal{A}}[C[Q_0] : \neg\varphi \wedge \neg\text{NonUnique}_{Q_0}] \\ &\leq \Pr^{\mathcal{A}}[C[Q_0] : \neg\{\text{prove}^\varphi(\theta, \mathcal{S})\}] && \text{by Lemma 39} \\ &\leq \Pr^{\mathcal{A}}[C[Q_0] \leq \neg\{\text{prove}^\varphi(\theta, \mathcal{S})\}] && \text{by Lemma 1} \\ &\leq p(\mathcal{A}, C) \end{aligned}$$

So $\text{Bound}_{Q_0}(V, \varphi, D_{\text{false}}, p)$. \square

Example 6 Let us prove that the example G_1 satisfies (4). We first study the facts \mathcal{F}_{μ_B} that hold at the program point μ_B that executes event e_B . \mathcal{F}_{μ_B} contains $\text{defined}(m[u[i_B]])$, $\text{defined}(x_B[u[i_B]])$, $\text{defined}(x_N[u[i_B]])$, $x_{pk_A}[i_B] = pk_A$, $B = x_B[u[i_B]]$, and $N[i_B] = x_N[u[i_B]]$, because the condition of find holds at μ_B . Moreover, we have $\text{defined}(m[u[i_B]]) \in \mathcal{F}_{\mu_B}$, and, when $m[i_A]$ is defined, $(\mu_A, i_A) : \text{event}(e_A(pk_A, x_B[i_A], x_N[i_A]))$ holds, so $(\mu_A, i_A) : \text{event}(e_A(pk_A, x_B[i_A], x_N[i_A]))\{u[i_B]/i_A\} \in \mathcal{F}_{\mu_B}$, that is, $(\mu_A, u[i_B]) : \text{event}(e_A(pk_A, x_B[u[i_B]], x_N[u[i_B]])) \in \mathcal{F}_{\mu_B}$. In other words, since m is defined at index $u[i_B]$, event e_A has been executed in the copy of Q'_{1A} of index $u[i_B]$. (\mathcal{F}_{μ_B} also contains other facts, which are useless for proving the desired correspondences, so we do not list them.)

For $\psi = F = \text{event}(e_B(x, y, z))$, μ_B is the only program point that executes F , so this event has been executed in some copy of Q'_{1B} of index i'_B , with $x_{pk_A}[i'_B] = x$, $B = y$, $N[i'_B] = z$. Then, when ψ holds, the facts $\theta' \mathcal{F}_{F, \mu_B} \supseteq \theta' \mathcal{F}_{\mu_B} \cup \{x_{pk_A}[i'_B] = x, B = y, N[i'_B] = z\}$ hold for some value of i'_B , with $\theta' = \{i'_B/i_B\}$. (We consider a single case c here, so we can simply omit the case c .)

Furthermore, the substitution $\theta(\mu_B)$ is the identity since all variables of ϕ also occur in ψ . Then we just have to show that $\theta' \mathcal{F}_{F, \mu_B}$ implies $\phi = \text{event}(e_A(x, y, z))$, that is, $\theta' \mathcal{F}_{F, \mu_B} \vdash_{\theta(\mu_B)} \text{event}(e_A(x, y, z))$. Since $(\mu_A, u[i_B]) : \text{event}(e_A(pk_A, x_B[u[i_B]], x_N[u[i_B]])) \in \mathcal{F}_{\mu_B}$, we have $(\mu_A, u[i'_B]) : \text{event}(e_A(pk_A, x_B[u[i'_B]], x_N[u[i'_B]])) \in \theta' \mathcal{F}_{F, \mu_B}$, so the equational prover just has to prove by contradiction that $e_A(pk_A, x_B[u[i'_B]], x_N[u[i'_B]]) = e_A(x, y, z)$, that is, $pk_A = x$, $x_B[u[i'_B]] = y$, and $x_N[u[i'_B]] = z$. The proof succeeds using the following equalities of $\theta' \mathcal{F}_{F, \mu_B}$: $x_{pk_A}[i'_B] = x$, $B = y$, $N[i'_B] = z$, $x_{pk_A}[i'_B] = pk_A$, $B = x_B[u[i'_B]]$, and $N[i'_B] = x_N[u[i'_B]]$.

Hence, G_1 satisfies (4) with any public variables V : if $\psi = \text{event}(e_B(x, y, z))$ has been executed, then $\phi = \text{event}(e_A(x, y, z))$ has been executed.

In the implementation, the substitution θ is initially defined as the identity on $\tilde{x} = \text{var}(\psi)$. It is defined on other variables when checking $\mathcal{F} \vdash_{\theta} M$ by trying to find θ such that $\theta M \in \mathcal{F}$, and when checking $\mathcal{F} \vdash_{\theta} \text{event}(e(M_1, \dots, M_m))$ by trying to find θ such that $\theta \text{event}(e(M_1, \dots, M_m)) \in \mathcal{F}$. When we do not manage to find the image by θ of all variables of M , resp. M_1, \dots, M_m , the check fails. When there are several suitable facts $\theta M \in \mathcal{F}$ or $\theta \text{event}(e(M_1, \dots, M_m)) \in \mathcal{F}$, the system tries all possibilities.

4.2.4 Injective Correspondences

Injective correspondences are more difficult to check than non-injective ones, because they require distinguishing between several executions of the same event. We achieve that by relying on the pair (program points, replication indices) that is recorded in the sequence $\mu \mathcal{E} \nu$ together with each event: distinct executions of events either occur at different program points or have different values of replication indices.

We extend Definition 16 to injective events, with exactly the same definition as for non-injective events.

The proof of injective correspondences extends that for non-injective correspondences: for a correspondence $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$, we additionally prove that distinct executions of the injective events of ψ correspond to distinct executions of each injective event of ϕ , that is, if the injective events of ψ have different pairs (program point, replication indices), then each injective event of ϕ has a different pair (program point, replication indices). In order to achieve this proof, we collect the following information for each injective event of ϕ :

- the set of facts \mathcal{F} that are known to hold, which will be used to reason on replication indices of events;
- the program point and replication indices of the considered injective event of ϕ , stored in a pair M_0 ; these program point and indices are computed when we prove that this event is executed;
- the program point and replication indices of the injective events of ψ , stored as a mapping $\mathcal{I} = \{j \mapsto (\mu_j, \theta_j I_{\mu_j}) \mid F_j \text{ is an injective event}\}$, where $\psi = F_1 \wedge \dots \wedge F_m$, μ_j is the program point that executes F_j , and θ_j is a renaming of I_{μ_j} to fresh replication indices, for $j \leq m$;
- the set \mathcal{V} containing the replication indices in \mathcal{F} and the variables \tilde{x} of ψ ; these variables will be renamed to fresh variables in order to avoid conflicts of variable names between different events.

This information is stored in a set \mathcal{S} , which contains quadruples $(\mathcal{F}, M_0, \mathcal{I}, \mathcal{V})$. We will show that, if the pair (program point, replication indices) of two executions of the injective events of ψ are different, then the pair (program, replication indices) of the corresponding executions of

the considered injective event of ϕ are also different. The equality between pairs (program point, replication indices) is obviously defined as the equality between program points and between replication indices. Formally, we consider $(\mathcal{F}, M_0, \mathcal{I}, \mathcal{V})$ and $(\mathcal{F}', M'_0, \mathcal{I}', \mathcal{V}')$ in \mathcal{S} . We rename the variables \mathcal{V}' of the second element to fresh variables by a substitution θ'' and show that, if $\mathcal{I} \neq \theta''\mathcal{I}'$, then $M_0 \neq \theta''M'_0$ (knowing \mathcal{F} and $\theta''\mathcal{F}'$). This property implies injectivity.

Since this reasoning is done for each injective event in ϕ , we collect the associated sets \mathcal{S} in a pseudo-formula \mathcal{C} , obtained by replacing each injective event of ϕ with a set \mathcal{S} and all other leaves of ϕ with \perp .

We say that $\vdash \mathcal{C}$ when for all non-bottom leaves \mathcal{S} of \mathcal{C} , for all $(\mathcal{F}, M_0, \mathcal{I}, \mathcal{V})$, $(\mathcal{F}', M'_0, \mathcal{I}', \mathcal{V}')$ in \mathcal{S} , $\mathcal{F} \cup \theta''\mathcal{F}' \cup \{\bigvee_{j \in \text{Dom}(\mathcal{I})} \mathcal{I}(j) \neq \theta''\mathcal{I}'(j), M_0 = \theta''M'_0\}$ yields a contradiction, where the substitution θ'' is a renaming of variables in \mathcal{V}' to distinct fresh variables. As explained above, the condition $\vdash \mathcal{C}$ guarantees injectivity.

We extend the definition of $\mathcal{F} \Vdash_{\theta} \phi$ used for non-injective correspondences to $\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi$, which means that \mathcal{F} implies $\theta\phi$ and \mathcal{C} correctly collects the tuples $(\mathcal{F}, M_0, \mathcal{I}, \mathcal{V})$ associated to this proof. Formally, we define:

$$\begin{aligned} \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \perp} M & \text{ if and only if } \mathcal{F} \cup \{\neg\theta M\} \text{ yields a contradiction} \\ \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \perp} \text{event}(e(M_1, \dots, M_m)) & \text{ if and only if} \\ & \text{there exist } M'_0, \dots, M'_m \text{ such that } M'_0 : \text{event}(e(M'_1, \dots, M'_m)) \in \mathcal{F} \text{ and} \\ & \mathcal{F} \cup \{\bigvee_{j=1}^m \theta M_j \neq M'_j\} \text{ yields a contradiction} \\ \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{S}} \text{inj-event}(e(M_1, \dots, M_m)) & \text{ if and only if} \\ & \text{there exist } M'_0, \dots, M'_m \text{ such that } M'_0 : \text{event}(e(M'_1, \dots, M'_m)) \in \mathcal{F}, \\ & \mathcal{F} \cup \{\bigvee_{j=1}^m \theta M_j \neq M'_j\} \text{ yields a contradiction, and } (\mathcal{F}, M'_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}. \\ \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1 \wedge \mathcal{C}_2} \phi_1 \wedge \phi_2 & \text{ if and only if } \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1 \text{ and } \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_2} \phi_2 \\ \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1 \vee \mathcal{C}_2} \phi_1 \vee \phi_2 & \text{ if and only if } \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1 \text{ or } \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_2} \phi_2 \end{aligned}$$

These formulas differ from the non-injective case in that we propagate \mathcal{I} , \mathcal{V} , \mathcal{C} and, in the case of injective events, we make sure that quadruples $(\mathcal{F}, M'_0, \mathcal{I}, \mathcal{V})$ are collected correctly by requiring that $(\mathcal{F}, M'_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}$.

Let $\varphi = \llbracket \forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket$ be a correspondence that does not use non-unique events, with $\psi = F_1 \wedge \dots \wedge F_m$, $\tilde{x} = \text{var}(\psi)$, and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$. Suppose that, in Q_0 , the arguments of the events that occur in ψ are always simple terms. Suppose that, for all $j \leq m$, μ_j executes F_j and c_j is a case for \mathcal{F}_{μ_j, c_j} . For $j \leq m$, let θ_j be a renaming of I_{μ_j} to fresh replication indices. (The renamings θ_j have pairwise disjoint images.) Let \mathcal{C} be a pseudo formula and θ be a family parameterized by $\mu_1, c_1, \dots, c_m, \mu_m$ of substitutions equal to the identity on \tilde{x} . We define $\text{prove}^{\varphi}(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m) = (\mathcal{F} \Vdash_{\theta(\mu_1, c_1, \dots, \mu_m, c_m)}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi)$ where $\mathcal{F} = \theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m}$, $\mathcal{I} = \{j \mapsto (\mu_j, \theta_j I_{\mu_j}) \mid F_j \text{ is an injective event}\}$, and $\mathcal{V} = \text{var}(\theta_1 I_{\mu_1}) \cup \dots \cup \text{var}(\theta_m I_{\mu_m}) \cup \{\tilde{x}\}$. The algorithm $\text{prove}^{\varphi}(\mathcal{C}, \theta, \mu_1, \dots, \mu_m)$ shows that the non-injective version of the correspondence φ holds assuming the events in $\psi = F_1 \wedge \dots \wedge F_m$ are executed at program points μ_1, \dots, μ_m respectively. Indeed, in this case, the facts $\mathcal{F} = \theta_1 \mathcal{F}_{F_1, \mu_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m}$ hold and the formula $\mathcal{F} \Vdash_{\theta(\mu_1, \dots, \mu_m)}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi$ shows that this implies $\theta(\mu_1, \dots, \mu_m)\phi$. (The substitution $\theta(\mu_1, \dots, \mu_m)\phi$ determines the values of \tilde{y} .) Additionally, $\text{prove}^{\varphi}(\mathcal{C}, \theta, \mu_1, \dots, \mu_m)$ makes sure that \mathcal{C} correctly collects the information needed to prove injectivity. We also define $\text{prove}^{\varphi}(\mathcal{C}, \theta, \mathcal{S}) = (\vdash \mathcal{C}) \wedge \bigwedge_{(\mu_1, c_1, \dots, \mu_m, c_m) \in \mathcal{S}} \text{prove}^{\varphi}(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m)$. This algorithm proves the correspondence φ assuming the events in ψ are executed at program points in \mathcal{S} . It verifies injectivity via $\vdash \mathcal{C}$.

The following proposition shows the soundness of the proof of injective correspondences based on this algorithm.

Proposition 3 Let $\varphi = \llbracket \forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket$ be a correspondence that does not use non-unique events, with $\psi = F_1 \wedge \dots \wedge F_m$, $\tilde{x} = \text{var}(\psi)$, and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$. Let Q_0 be a process that satisfies Properties 3 and 4. Suppose that, in Q_0 , the arguments of the events that occur in ψ are always simple terms.

Let $\mathcal{S} = \{(\mu_1, c_1, \dots, \mu_m, c_m) \mid \forall j \leq m, \mu_j \text{ executes } F_j \text{ and } c_j \text{ is a case for } \mathcal{F}_{\mu_j, c_j}\}$. Assume that there exist a pseudo-formula \mathcal{C} and a family of substitutions θ equal to the identity on \tilde{x} such that $\text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S})$. Assume that for all $\mathcal{A} \in \mathcal{BB}$ and all evaluation contexts C acceptable for Q_0 , $\text{Pr}^{\mathcal{A}}[C[Q_0] \preceq \neg\{\text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S})\}] \leq p(\mathcal{A}, C)$. Then $\text{Bound}_{Q_0}(V, \varphi, D_{\text{false}}, p)$ for any V .

In the implementation, the value of \mathcal{C} is computed by adding $(\mathcal{F}, M'_0, \mathcal{I}, \mathcal{V})$ to \mathcal{S} when handling injective events during the checking of $\text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S})$. We check $\vdash \mathcal{C}$ incrementally, after each addition of an element to \mathcal{C} . The proof of Proposition 3 relies on the following definitions and lemmas. We have

$$\begin{aligned} \{\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \perp} M\}\} &= \forall \tilde{z} \in \tilde{T}'', \neg \left(\bigwedge \mathcal{F} \wedge \neg \theta M \right) \\ \{\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \perp} \text{event}(e(M_1, \dots, M_m))\}\} &= \forall \tilde{z} \in \tilde{T}'', \neg \left(\bigwedge \mathcal{F} \wedge \bigvee_{j=1}^m \theta M_j \neq M'_j \right) \\ &\quad \text{for some } M'_0 : \text{event}(e(M'_1, \dots, M'_m)) \in \mathcal{F} \\ \{\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{S}} \text{inj-event}(e(M_1, \dots, M_m))\}\} &= \forall \tilde{z} \in \tilde{T}'', \neg \left(\bigwedge \mathcal{F} \wedge \bigvee_{j=1}^m \theta M_j \neq M'_j \right) \\ &\quad \text{for some } M'_0 : \text{event}(e(M'_1, \dots, M'_m)) \in \mathcal{F} \text{ and } (\mathcal{F}, M'_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S} \\ \{\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1 \wedge \mathcal{C}_2} \phi_1 \wedge \phi_2\}\} &= \{\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1\}\} \wedge \{\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_2} \phi_2\}\} \\ \{\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1 \vee \mathcal{C}_2} \phi_1 \vee \phi_2\}\} &= \begin{cases} \{\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1\}\} & \text{if } \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1 \\ \{\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_2} \phi_2\}\} & \text{otherwise} \end{cases} \end{aligned}$$

where $\tilde{z} = \mathcal{V}$ and \tilde{T}'' are the types of these variables. We also have

$$\begin{aligned} \{\{\vdash \mathcal{C}\}\} &= \bigwedge_{\mathcal{S} \neq \perp \text{ leaf of } \mathcal{C}} \bigwedge_{(\mathcal{F}, M_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}} \bigwedge_{(\mathcal{F}', M'_0, \mathcal{I}', \mathcal{V}') \in \mathcal{S}} \forall \tilde{z} \in \tilde{T}'', \\ &\quad \neg \left(\bigwedge \mathcal{F} \wedge \bigwedge \theta'' \mathcal{F}' \wedge \left(\bigvee_{j \in \text{Dom}(\mathcal{I})} \mathcal{I}(j) \neq \theta'' \mathcal{I}'(j) \right) \wedge M_0 = \theta'' M'_0 \right) \end{aligned}$$

where the substitution θ'' is a renaming of variables in \mathcal{V}' to distinct fresh variables, $\tilde{z} = \mathcal{V} \cup \theta'' \mathcal{V}'$, and \tilde{T}'' are the types of these variables.

We define formula $(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi)$ as follows:

$$\begin{aligned} \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \perp} M) &= \theta M \\ \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \perp} \text{event}(e(M_1, \dots, M_m))) &= \theta \text{event}(e(M_1, \dots, M_m)) \\ \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{S}} \text{inj-event}(e(M_1, \dots, M_m))) &= \\ &\quad \bigvee_{(\mathcal{F}, N_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}} \exists \tau \in \mathbb{N}, N_0 : \theta \text{event}(e(M_1, \dots, M_m)) @ \tau \\ \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1 \wedge \mathcal{C}_2} \phi_1 \wedge \phi_2) &= \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1) \wedge \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_2} \phi_2) \\ \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1 \vee \mathcal{C}_2} \phi_1 \vee \phi_2) &= \begin{cases} \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1) & \text{if } \mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1 \\ \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_2} \phi_2) & \text{otherwise} \end{cases} \end{aligned}$$

The formula $\text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi)$ generalizes $\theta \phi$ to the case of injective events.

Lemma 40 $\{\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi\} \Rightarrow \forall \tilde{z} \in \tilde{T}'', \neg \left(\bigwedge \mathcal{F} \wedge \neg \text{formula}(\mathcal{F} \Vdash_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi) \right)$ where $\tilde{z} = \mathcal{V}$ and \tilde{T}'' are the types of these variables.

Proof By induction on ϕ . This result is similar to Lemma 37. The case of injective events is new. The case of disjunction differs, but is straightforward by induction hypothesis. \square

The next lemma shows that, for events e in the considered correspondence, two distinct executions of event e have distinct pairs (program point, replication indices). When the term M contains no array accesses, we define $\sigma(M)$ by $\sigma, M \Downarrow \sigma(M)$.

Lemma 41 Assume that the event e is used in the correspondence φ . Let Q_0 be a process that satisfies Property 3. Let C be an evaluation context acceptable for Q_0 with public variables V that does not contain events used by φ and \mathcal{A} be an attacker in \mathcal{BB} . If $\text{initConfig}(C[Q_0], \mathcal{A}) \xrightarrow{p_1}_{t_1} \dots \xrightarrow{p_{m-1}}_{t_{m-1}} E, St, Q, Or, T, \mu Ev, \mu Ev(\tau) = (\mu, \tilde{a}) : e(a_1, \dots, a_m), \mu Ev(\tau') = (\mu', \tilde{a}') : e(a'_1, \dots, a'_m)$, and $\tau \neq \tau'$, then $(\mu, \tilde{a}) \neq (\mu', \tilde{a}')$.

Proof Let us fix the event symbol e . We define the multisets $Events(\tilde{a}, M)$, $Events(\tilde{a}, P)$, and $Events(\tilde{a}, Q)$ by

$$Events(\tilde{a}, {}^{\mu}i) = \emptyset$$

$$Events(\tilde{a}, {}^{\mu}x[M_1, \dots, M_m]) = \biguplus_{j \in \{1, \dots, m\}} Events(\tilde{a}, M_j)$$

$$Events(\tilde{a}, {}^{\mu}f(M_1, \dots, M_m)) = \biguplus_{j \in \{1, \dots, m\}} Events(\tilde{a}, M_j)$$

$$Events(\tilde{a}, {}^{\mu}x[\tilde{i}] \stackrel{R}{\leftarrow} T; N) = Events(\tilde{a}, N)$$

$$Events(\tilde{a}, {}^{\mu}\text{let } x[\tilde{i}] : T = M \text{ in } N) = Events(\tilde{a}, M) \uplus Events(\tilde{a}, N)$$

$$Events(\tilde{a}, {}^{\mu}\text{if } M \text{ then } N \text{ else } N') = Events(\tilde{a}, M) \uplus \max(Events(\tilde{a}, N), Events(\tilde{a}, N'))$$

$$Events(\tilde{a}, {}^{\mu}\text{find}[\text{unique?}] \left(\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ such that } \text{defined}(\tilde{M}_j) \wedge M'_j \text{ then } N_j \right) \text{ else } N') =$$

$$\biguplus_{j=1}^m \biguplus_{\tilde{a}_j \leq \tilde{n}_j} Events((\tilde{a}, \tilde{a}_j), M'_j) \uplus \max(\max_{j=1}^m Events(\tilde{a}, N_j), Events(\tilde{a}, N'))$$

$$Events(\tilde{a}, {}^{\mu}\text{event } e'(M_1, \dots, M_m); M) = \biguplus_{j \in \{1, \dots, m\}} Events(\tilde{a}, M_j) \uplus Events(\tilde{a}, M) \text{ if } e' \neq e$$

$$Events(\tilde{a}, {}^{\mu}\text{event } e(M_1, \dots, M_m); M) = \{(\mu, \tilde{a})\} \uplus \biguplus_{j \in \{1, \dots, m\}} Events(\tilde{a}, M_j) \uplus Events(\tilde{a}, M)$$

$$Events(\tilde{a}, {}^{\mu}\text{event_abort } e') = \emptyset \text{ if } e' \neq e$$

$$Events(\tilde{a}, {}^{\mu}\text{event_abort } e) = \{(\mu, \tilde{a})\}$$

$$Events(\tilde{a}, {}^{\mu}0) = \emptyset$$

$$Events(\tilde{a}, {}^{\mu}(Q_1 \mid Q_2)) = Events(\tilde{a}, Q_1) \uplus Events(\tilde{a}, Q_2)$$

$$Events(\tilde{a}, {}^{\mu}\text{foreach } i \leq n \text{ do } Q) = \biguplus_{a \in [1, n]} Events((\tilde{a}, a), Q)$$

$$Events(\tilde{a}, {}^{\mu}\text{newOracle } O; Q) = Events(\tilde{a}, Q)$$

$$Events(\tilde{a}, {}^{\mu}O[\tilde{i}](x[\tilde{i}] : T) := P) = Events(\tilde{a}, P)$$

$$\begin{aligned}
Events(\tilde{a}, \text{ }^{\mu}\text{return } (N); Q) &= Events(\tilde{a}, N) \uplus Events(\tilde{a}, Q) \\
Events(\tilde{a}, \text{ }^{\mu}\text{x}[\tilde{i}] \stackrel{R}{\leftarrow} T; P) &= Events(\tilde{a}, P) \\
Events(\tilde{a}, \text{ }^{\mu}\text{let } x[\tilde{i}] : T = M \text{ in } P) &= Events(\tilde{a}, M) \uplus Events(\tilde{a}, P) \\
Events(\tilde{a}, \text{ }^{\mu}\text{let } x[\tilde{i}] : T = O[M_1, \dots, M_k, ?u_1[\tilde{i}] \leq n_1, \dots, ?u_m[\tilde{i}] \leq n_m](M_{k+1}) \text{ in } P_1 \text{ else } P_2) &= \\
&\bigoplus_{i=1}^{k+1} Events(\tilde{a}, M_i) \uplus \max(Events(\tilde{a}, P_1), Events(\tilde{a}, P_2)) \\
Events(\tilde{a}, \text{ }^{\mu}\text{if } M \text{ then } P \text{ else } P') &= Events(\tilde{a}, M) \uplus \max(Events(\tilde{a}, P), Events(\tilde{a}, P')) \\
Events(\tilde{a}, \text{ }^{\mu}\text{find}[unique?](\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(\tilde{M}_j) \wedge M'_j \text{ then } P_j) \text{ else } P) &= \\
&\bigoplus_{j=1}^m \bigoplus_{\tilde{a}_j \leq \tilde{n}_j} Events((\tilde{a}, \tilde{a}_j), M'_j) \uplus \max(\max_{j=1}^m Events(\tilde{a}, P_j), Events(\tilde{a}, P)) \\
Events(\tilde{a}, \text{ }^{\mu}\text{event } e'(M_0, \dots, M_m); P) &= \bigoplus_{j \in \{1, \dots, m\}} Events(\tilde{a}, M_j) \uplus Events(\tilde{a}, P) \text{ if } e' \neq e \\
Events(\tilde{a}, \text{ }^{\mu}\text{event } e(M_0, \dots, M_m); P) &= \{(\mu, \tilde{a})\} \uplus \bigoplus_{j \in \{1, \dots, m\}} Events(\tilde{a}, M_j) \uplus Events(\tilde{a}, P) \\
Events(\tilde{a}, \text{ }^{\mu}\text{event_abort } e') &= \emptyset \text{ if } e' \neq e \\
Events(\tilde{a}, \text{ }^{\mu}\text{event_abort } e) &= \{(\mu, \tilde{a})\} \\
Events(\tilde{a}, \text{ }^{\mu}\text{yield}) &= \emptyset
\end{aligned}$$

(get and insert are omitted because they do not occur in the game by Property 3 and as mentioned in the beginning of Section 3, get and insert in the adversary context can be encoded using find.)

We define the multisets

$$\begin{aligned}
Events(\mu\mathcal{E}v) &= \{(\mu, \tilde{a}) \mid (\mu, \tilde{a}) : e(\dots) \in \mu\mathcal{E}v\} \\
Events(E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v) &= Events(\text{Im}(\sigma), M) \uplus Events(\mu\mathcal{E}v) \\
Events(\mathcal{Q}, \mathcal{O}r) &= \bigoplus_{(\sigma', \mathcal{Q}') \in \mathcal{Q}} Events(\text{Im}(\sigma'), \mathcal{Q}') \\
Events(E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v) &= \bigoplus_{(\sigma, P) \in St} Events(\text{Im}(\sigma), P) \uplus \\
&\bigoplus_{(\sigma', \mathcal{Q}') \in \mathcal{Q}} Events(\text{Im}(\sigma'), \mathcal{Q}') \uplus Events(\mu\mathcal{E}v)
\end{aligned}$$

The latter multiset contains all pairs (μ, \tilde{a}) (program point, value of replication indices) for events $e(\dots)$ that may be executed in a trace that contains the configuration $E, St, \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu\mathcal{E}v$.

The multiset $Events(\sigma_0, C[Q_0])$ contains no duplicates. Indeed, we show by induction on M that $Events(\tilde{a}, M)$ is included in the multiset $\{(\mu, \tilde{a}')\}$ where μ is a program point inside M and \tilde{a} is a prefix of \tilde{a}' , and similarly for P and Q . That allows to show that all multiset unions in the computation of $Events$ are disjoint unions, since all recursive calls in the computation of $Events$ are either with disjoint processes or terms, or with different extensions of \tilde{a} .

Moreover, by induction on the derivations, if $E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v \xrightarrow{p} E', \sigma', M', \mathcal{T}', \mu\mathcal{E}v'$, then $Events(E, \sigma, M, \mathcal{T}, \mu\mathcal{E}v) \supseteq Events(E', \sigma', M', \mathcal{T}', \mu\mathcal{E}v')$; if $\mathcal{Q}, \mathcal{O}r \rightsquigarrow \mathcal{Q}', \mathcal{O}r'$, then $Events(\mathcal{Q}, \mathcal{O}r) \supseteq Events(\mathcal{Q}', \mathcal{O}r')$; and if $Conf \xrightarrow{p} Conf'$, then $Events(Conf) \supseteq Events(Conf')$.

Therefore, the multiset $Events(\{(\sigma_0, C[Q_0])\}, (\text{fo}(C[Q_0]), \emptyset))$ contains no duplicates, and neither do the multisets $Events(\text{reduce}(\{(\sigma_0, C[Q_0])\}, (\text{fo}(C[Q_0]), \emptyset)), Events(\text{initConfig}(C[Q_0], \mathcal{A})),$

and $Events(Conf)$, where $Conf = E, St, Q, Or, T, \mu Ev$ is the final configuration of the considered trace. Hence, $Events(\mu Ev)$ contains no duplicates, which implies the desired result. \square

Lemma 42 *Let $\varphi = \llbracket \forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket$ be a correspondence that does not use non-unique events, with $\psi = F_1 \wedge \dots \wedge F_m$, $\tilde{x} = \text{var}(\psi)$, and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$. Let Q_0 be a process that satisfies Properties 3 and 4. Suppose that, in Q_0 , the arguments of the events that occur in ψ are always simple terms.*

Let $\mathcal{S} = \{(\mu_1, c_1, \dots, \mu_m, c_m) \mid \forall j \leq m, \mu_j \text{ executes } F_j \text{ and } c_j \text{ is a case for } \mathcal{F}_{\mu_j, c_j}\}$. Let C be an evaluation context acceptable for Q_0 with public variables V that does not contain events used by φ and \mathcal{A} be an attacker in \mathcal{BB} . Let Tr be a full trace of $C[Q_0]$ for \mathcal{A} that does not execute any non-unique event of Q_0 . If $Tr \vdash \neg \varphi$, then for any family of substitutions θ equal to the identity on \tilde{x} , for any pseudo-formula \mathcal{C} , $Tr \vdash \neg \llbracket \text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S}) \rrbracket$.

Proof By contraposition, we suppose that, $Tr \vdash \llbracket \text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S}) \rrbracket$, so for every μ_1 that executes F_1, \dots , for every μ_m that executes F_m , for every c_1, \dots, c_m , $Tr \vdash \llbracket \text{prove}^\varphi(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m) \rrbracket$ and $Tr \vdash \llbracket \neg \mathcal{C} \rrbracket$, and we show that $Tr \vdash \varphi$.

Let μEv be the sequence of events in the last configuration of Tr .

We use the notations of Definition 12. We construct the functions f_1, \dots, f_k as follows. Let ρ be a mapping of τ_1, \dots, τ_m to elements of \mathbb{N} and of \tilde{x} to elements of \tilde{T} . Suppose that $Tr, \rho \vdash \psi^\tau$. Then, for all $j \leq m$, $Tr, \rho \vdash F_j^\tau$. By Lemma 38, there exists a program point μ_j (in Q_0) that executes F_j and a case c_j such that, for any θ_j renaming of I_{μ_j} to fresh replication indices, there exists a mapping σ_j with domain $\theta_j I_{\mu_j}$ such that $\mu Ev(\rho(\tau_j)) = (\mu_j, \sigma_j(\theta_j I_{pp_j})) : \dots$ and $Tr, \sigma_j \cup \rho \vdash \theta_j \mathcal{F}_{F_j, \mu_j, c_j}$. Let $\rho_1 = \sigma_1 \cup \dots \cup \sigma_m \cup \rho$. We have $Tr, \rho_1 \vdash \theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m}$.

Let $\mathcal{F}(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m) = \theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m} \cup \{\neg \text{formula}(\text{prove}^\varphi(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m))\}$. By Lemma 40, $\llbracket \text{prove}^\varphi(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m) \rrbracket \Rightarrow \forall \theta_1 I_{\mu_1}, \dots, \forall \theta_m I_{\mu_m}, \forall \tilde{x} \in \tilde{T}, \neg \bigwedge \mathcal{F}(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m)$. So $Tr \vdash \forall \theta_1 I_{\mu_1}, \dots, \forall \theta_m I_{\mu_m}, \forall \tilde{x} \in \tilde{T}, \neg \bigwedge \mathcal{F}(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m)$.

Then $Tr, \rho_1 \vdash \neg \bigwedge \mathcal{F}(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m)$, so $Tr, \rho_1 \vdash \text{formula}(\text{prove}^\varphi(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m))$, that is, $Tr, \rho_1 \vdash \text{formula}(\mathcal{F} \stackrel{\mathcal{I}, \mathcal{V}, \mathcal{C}}{\models}_\theta \phi)$, with $\mathcal{F} = \theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m}$, $\mathcal{I} = \{j \mapsto (\mu_j, \theta_j I_{\mu_j}) \mid F_j \text{ is an injective event}\}$, $\mathcal{V} = \text{var}(\theta_1 I_{\mu_1}) \cup \dots \cup \text{var}(\theta_m I_{\mu_m}) \cup \{\tilde{x}\}$, and $\theta = \theta(\mu_1, c_1, \dots, \mu_m, c_m)$.

Consider an injective event in ϕ , associated to function f_l .

- If that injective event corresponds to

$$\bigvee_{(\mathcal{F}, N_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}} \exists \tau \in \mathbb{N}, N_0 : \theta \text{event}(e(M_1, \dots, M_m)) @ \tau$$

in $\text{formula}(\mathcal{F} \stackrel{\mathcal{I}, \mathcal{V}, \mathcal{C}}{\models}_\theta \phi)$, we have

$$Tr, \rho_1 \vdash \bigvee_{(\mathcal{F}, N_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}} \exists \tau \in \mathbb{N}, N_0 : \theta \text{event}(e(M_1, \dots, M_m)) @ \tau$$

since $\text{formula}(\mathcal{F} \stackrel{\mathcal{I}, \mathcal{V}, \mathcal{C}}{\models}_\theta \phi)$ is a conjunction. So $Tr, \rho_1[\tau \mapsto a] \vdash N_0 : \theta \text{event}(e(M_1, \dots, M_m)) @ \tau$ for some $a \in \mathbb{N}$ and some N_0 such that $(\mathcal{F}, N_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}$. We define $f_l(\rho(\tau_1), \dots, \rho(\tau_m), \rho(\tilde{x})) = a$, so that $Tr, \rho_1 \vdash N_0 : \theta \text{event}(e(M_1, \dots, M_m)) @ f_l(\tau_1, \dots, \tau_m, \tilde{x})$.

Moreover, if $j \in I$, then F_j is an injective event, $F_j = \text{inj-event}(e_j(M_{j,1}, \dots, M_{j,m}))$. Moreover, $Tr, \rho \vdash \psi^\tau$, so $Tr, \rho \vdash F_j^\tau$, so $Tr, \rho \vdash \text{event}(e_j(M_{j,1}, \dots, M_{j,m})) @ \tau_j$. Since $\mu Ev(\rho(\tau_j)) = (pp_j, \sigma_j(\theta_j I_{\mu_j})) : \dots = (\mu_j, \rho_1(\theta_j I_{\mu_j})) : \dots$ and $\mathcal{I}(j) = (\mu_j, \theta_j I_{\mu_j})$, we have $Tr, \rho_1 \vdash \mathcal{I}(j) : \text{event}(e_j(M_{j,1}, \dots, M_{j,m})) @ \tau_j$.

- If that injective event is in a removed disjunct in formula ($\mathcal{F} \models_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi$), then we define $f_l(\rho(\tau_1), \dots, \rho(\tau_m), \rho(\tilde{x})) = \perp$.

Then we have $Tr, \rho_1 \vdash \theta \phi^\tau$, so $Tr, \rho_1 \vdash \exists \tilde{y} \in \tilde{T}', \phi^\tau$.

Hence, applying this construction for all ρ , we obtain $Tr \vdash \forall \tau_1, \dots, \tau_m \in \mathbb{N}, \forall \tilde{x} \in \tilde{T}, (\psi^\tau \Rightarrow \exists \tilde{y} \in \tilde{T}', \phi^\tau)$. It remains to show $\text{Inj}(I, f_l)$ for each $l \in \{1, \dots, k\}$.

Suppose $f_l(a_1, \dots, a_m, \tilde{a}) = f_l(a'_1, \dots, a'_m, \tilde{a}') \neq \perp$. Let $\rho = \{\tau_1 \mapsto a_1, \dots, \tau_m \mapsto a_m, \tilde{x} \mapsto \tilde{a}\}$. Let \mathcal{S} be the leaf of \mathcal{C} corresponding to the event associated to f_l . By the construction above, we have $\theta, (\mathcal{F}, N_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}$, and an extension ρ_1 of ρ such that

$$Tr, \rho_1 \vdash N_0 : \theta \text{event}(e(M_1, \dots, M_m)) @ f_l(\tau_1, \dots, \tau_m, \tilde{x}) \quad (24)$$

$$Tr, \rho_1 \vdash \mathcal{F} \quad (25)$$

$$\text{for } j \in I, Tr, \rho_1 \vdash \mathcal{I}(j) : \text{event}(e_j(M_{j,1}, \dots, M_{j,m})) @ \tau_j \quad (26)$$

Let $\rho' = \{\tau_1 \mapsto a'_1, \dots, \tau_m \mapsto a'_m, \tilde{x} \mapsto \tilde{a}'\}$. In the same way, we have $\theta', (\mathcal{F}', N'_0, \mathcal{I}', \mathcal{V}') \in \mathcal{S}$, and an extension ρ'_1 of ρ' such that

$$Tr, \rho'_1 \vdash N'_0 : \theta' \text{event}(e(M_1, \dots, M_m)) @ f_l(\tau_1, \dots, \tau_m, \tilde{x})$$

$$Tr, \rho'_1 \vdash \mathcal{F}'$$

$$\text{for } j \in I, Tr, \rho'_1 \vdash \mathcal{I}'(j) : \text{event}(e_j(M_{j,1}, \dots, M_{j,m})) @ \tau_j.$$

Let θ'' be a renaming of the domain of ρ'_1 to fresh variables. We have

$$Tr, \rho'_1 \theta''^{-1} \vdash \theta'' N'_0 : \theta'' \theta' \text{event}(e(M_1, \dots, M_m)) @ f_l(\tau_1, \dots, \tau_m, \tilde{x}) \quad (27)$$

$$Tr, \rho'_1 \theta''^{-1} \vdash \theta'' \mathcal{F}' \quad (28)$$

$$\text{for } j \in I, Tr, \rho'_1 \theta''^{-1} \vdash \theta'' \text{event}(e_j(\mathcal{I}'(j), M_{j,1}, \dots, M_{j,m})) @ \tau_j. \quad (29)$$

Therefore, by (24) and (27),

$$Tr, \rho_1 \cup \rho'_1 \theta''^{-1} \vdash \theta N_0 : \text{event}(e(M_1, \dots, M_m)) @ f_l(a_1, \dots, a_m, \tilde{a})$$

$$Tr, \rho_1 \cup \rho'_1 \theta''^{-1} \vdash \theta'' N'_0 : \theta'' \theta' \text{event}(e(M_1, \dots, M_m)) @ f_l(a'_1, \dots, a'_m, \tilde{a}').$$

Since $f_l(a_1, \dots, a_m, \tilde{a}) = f_l(a'_1, \dots, a'_m, \tilde{a}')$, the events are the same, so

$$Tr, \rho_1 \cup \rho'_1 \theta''^{-1} \vdash N_0 = \theta'' N'_0. \quad (30)$$

We also have by (25) and (28),

$$Tr, \rho_1 \cup \rho'_1 \theta''^{-1} \vdash \mathcal{F} \cup \theta'' \mathcal{F}'. \quad (31)$$

Since $Tr \vdash \{\vdash \mathcal{C}\}$, we have

$$Tr, \rho_1 \cup \rho'_1 \theta''^{-1} \vdash \neg \left(\bigwedge \mathcal{F} \wedge \bigwedge \theta'' \mathcal{F}' \wedge \left(\bigvee_{j \in \text{Dom}(\mathcal{I})} \mathcal{I}(j) \neq \theta'' \mathcal{I}'(j) \right) \wedge N_0 = \theta'' N'_0 \right)$$

so using (31) and (30), we conclude that

$$Tr, \rho_1 \cup \rho'_1 \theta''^{-1} \vdash \bigwedge_{j \in I} \mathcal{I}(j) = \theta'' \mathcal{I}'(j)$$

since $\text{Dom}(\mathcal{I}) = I$. Let $\mu \mathcal{E}v$ be the sequence of events at the end of Tr . For $j \in I$, let $b_j = \rho_1(\mathcal{I}(j)) = \rho'_1(\mathcal{I}'(j))$. By (26) and (29), we have $\mu \mathcal{E}v(a_j) = b_j : e_j(\dots)$ and $\mu \mathcal{E}v(a'_j) = b_j : e_j(\dots)$. By Lemma 41, we have $a_j = a'_j$. That proves $\text{Inj}(I, f_l)$, and concludes the proof that $Tr \vdash \varphi$. \square

Proof of Proposition 3 Let C be an evaluation context acceptable for Q_0 with public variables V that does not contain events used by φ and \mathcal{A} be an attacker in \mathcal{BB} . We have

$$\begin{aligned} \text{Adv}_{Q_0}(\mathcal{A}, C, \varphi, D_{\text{false}}) &= \Pr^{\mathcal{A}}[C[Q_0] : \neg\varphi \wedge \neg\text{NonUnique}_{Q_0}] \\ &\leq \Pr^{\mathcal{A}}[C[Q_0] : \neg\{\{\text{prove}^\varphi(C, \theta, \mathcal{S})\}\}] && \text{by Lemma 42} \\ &\leq \Pr^{\mathcal{A}}[C[Q_0] \preceq \neg\{\{\text{prove}^\varphi(C, \theta, \mathcal{S})\}\}] && \text{by Lemma 1} \\ &\leq p(\mathcal{A}, C) \end{aligned}$$

So $\text{Bound}_{Q_0}(V, \varphi, D_{\text{false}}, p)$. □

Example 7 Let us prove that the example G_1 satisfies (6). We prove the correspondence $(\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}' ; \phi) = (\forall x : pkey, y : host, z : nonce; \text{inj-event}(e_B(x, y, z)) \Rightarrow \text{inj-event}(e_A(x, y, z)))$. The program point μ_B executes $F = \text{event}(e_B(x, y, z))$ and $\mathcal{F} = \theta' \mathcal{F}_{F, \mu_B} \supseteq \theta' \mathcal{F}_{\mu_B} \{i'_B/i_B\} \cup \{x_{pk_A}[i'_B] = x, B = y, N[i'_B] = z\}$ with $\theta' = \{i'_B/i_B\}$.

As in the proof of $\mathcal{F} \mapsto_{\theta(\mu_B)} \text{event}(e_A(x, y, z))$ in Example 6, we show $\mathcal{F} \mapsto_{\theta(\mu_B)}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \text{event}(e_A(x, y, z))$ where $\mathcal{I} = \{1 \mapsto (\mu_B, i'_B)\}$ encodes the program points and replication indices of the events of ψ , $\mathcal{V} = \{i'_B, x, y, z\}$ contains the replication indices of \mathcal{F} and the variables of ψ , $\mathcal{C} = \mathcal{S} = \{(\mathcal{F}, (\mu_A, u[i'_B]), \mathcal{I}, \mathcal{V})\}$. ($\mathcal{C} = \mathcal{S}$ because the formula ψ is reduced to a single event; $M'_0 = (\mu_A, u[i'_B])$ contains the program point and replication indices of the event e_A contained in \mathcal{F} : $(\mu_A, u[i'_B]) : \text{event}(e_A(pk_A, x_B[u[i'_B]], x_N[u[i'_B]])) \in \mathcal{F}$.)

In order to prove injectivity, it remains to show that $\vdash \mathcal{C}$. Let $\theta'' = \{i''_B/i'_B, x''/x, y''/y, z''/z\}$. We need to show that $\mathcal{F} \cup \theta'' \mathcal{F} \cup \{(\mu_B, i'_B) \neq (\mu_B, i''_B), (\mu_A, u[i'_B]) = (\mu_A, u[i''_B])\}$ yields a contradiction, that is, if the pairs (program point, replication indices) of the event e_B in ψ are distinct $((\mu_B, i'_B) \neq (\mu_B, i''_B))$, then the pairs (program point, replication indices) of the event e_A in ϕ are also distinct $((\mu_A, u[i'_B]) \neq (\mu_A, u[i''_B]))$.

\mathcal{F} contains $N[i'_B] = x_N[u[i'_B]]$, so $\theta'' \mathcal{F}$ contains $N[i''_B] = x_N[u[i''_B]]$. These two equalities combined with $u[i'_B] = u[i''_B]$ imply that $N[i'_B] = x_N[u[i'_B]] = x_N[u[i''_B]] = N[i''_B]$. Since N is defined by random choices of the large type *nonce*, $N[i'_B] = N[i''_B]$ implies $i'_B = i''_B$ up to probability $n^2/2|\text{nonce}|$, by eliminating collisions. This equality contradicts $(\mu_B, i'_B) \neq (\mu_B, i''_B)$, so we obtain the desired injectivity. Therefore, the game G_1 satisfies (6) with any public variables V up to probability $n^2/2|\text{nonce}|$.

5 Game Transformations

5.1 Syntactic Game Transformations

5.1.1 auto_SArename

The transformation **auto_SArename** renames all variables defined in conditions of **find**, so that they have distinct names. This transformation is a particular case of **SArename** (Section 5.1.7) that is particularly simple because these variables do not have array accesses by Invariant 3.

Lemma 43 *The transformation **auto_SArename** requires and preserves Properties 1, 2, and 3. It preserves Property 5. If transformation **auto_SArename** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V} G', D, \text{EvUsed}$ and G' satisfies Property 4.*

5.1.2 expand_tables [33]

The transformation **expand_tables** transforms **insert** and **get** into **find**, since the other transformations do not support tables. It proceeds by storing the inserted list elements in fresh array variables, and looking up in these arrays instead of performing **get**. More precisely, when **insert** $Tbl(M_1, \dots, M_k); P$ is under the replications **foreach** $i_1 \leq n_1$ **do** ... **foreach** $i_l \leq n_l$ **do** , it is transformed into

$$\text{let } y_1[i_1, \dots, i_l] = M_1 \text{ in } \dots \text{let } y_k[i_1, \dots, i_l] = M_k \text{ in } P$$

where y_1, \dots, y_k are fresh array variables, and we add $(y_1, \dots, y_k; i_1 \leq n_1, \dots, i_l \leq n_l)$ in a set S' , to remember them. The construct **get** $[unique?]$ $Tbl(x_1 : T_1, \dots, x_k : T_k)$ **suchthat** M **in** P **else** P' is then transformed into

$$\text{find}[unique?] \left(\bigoplus_{(y_1, \dots, y_k; i_1 \leq n_1, \dots, i_l \leq n_l) \in S'} \begin{array}{l} u_1 = i'_1 \leq n_1, \dots, u_l = i'_l \leq n_l \text{ suchthat} \\ \text{defined}(y_1[\tilde{i}'], \dots, y_k[\tilde{i}']) \wedge M\{y_1[\tilde{i}']/x_1, \dots, y_k[\tilde{i}']/x_k\} \\ \text{then let } x_1 = y_1[\tilde{u}] \text{ in } \dots \text{let } x_k = y_k[\tilde{u}] \text{ in } P \end{array} \right) \\ \text{else } P'$$

where $[unique?]$ is either $[unique_e]$ or empty and has the same value at both occurrences, \tilde{u} stands for u_1, \dots, u_l , and \tilde{i}' stands for i'_1, \dots, i'_l . This construct looks in all arrays used for translating insertion in table Tbl , for indices \tilde{i}' such that $y_1[\tilde{i}'], \dots, y_k[\tilde{i}']$ are defined, that is, an element has been inserted at indices \tilde{i}' , and $M\{y_1[\tilde{i}']/x_1, \dots, y_k[\tilde{i}']/x_k\}$ is true, that is, that element satisfies M . When it finds such an element, it stores it in x_1, \dots, x_k , and runs P . (When it finds several elements, one of them is chosen randomly with uniform probability when $[unique?]$ is empty and the non-unique event e is raised when $[unique?]$ is $[unique_e]$.) When it finds no element, it executes P' . These transformations are described for processes, but similar transformations are performed for **insert** and **get** terms.

After this transformation, **expand_tables** calls **auto_SArename** to guarantee Property 4.

Lemma 44 *The transformation **expand_tables** requires and preserves Properties 1 and 2. It preserves Property 5. If transformation **expand_tables** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, EvUsed \xrightarrow{V}_p G', D, EvUsed$, where p is ϵ_{find} times the maximal number of executions of **get** in G that is not obviously unique (that is, such that the **insert** in that table may be executed several times), and G' satisfies Properties 3 and 4.*

5.1.3 expand

The transformation **expand** transforms terms $\stackrel{R}{\leftarrow}$, **let**, **if**, **find**, **event**, and **event_abort** into processes, so that Property 5 is guaranteed. It simplifies the generated game on the fly, using many of the rules of **simplify** (Section 5.1.21), to avoid generating branches that can actually not be executed.

After this transformation, **expand** calls **auto_SArename** to guarantee Property 4.

Lemma 45 *The transformation **expand** requires and preserves Properties 1, 2, 3, and 4. If transformation **expand** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, EvUsed \xrightarrow{V}_p G', D, EvUsed$, where p is an upper bound of the probability that the simplification steps modify the execution, and G' satisfies Property 5.*

5.1.4 prove_unique

The transformation **prove_unique** tries to prove that each $\text{find}[\text{unique}_e]$ really has a unique possibility at runtime (up to a small probability), so that event e is executed with at most that probability. More precisely, $\text{find}[\text{unique}]$ are already proved; $\text{find}[\text{unique}_e]$ for which no query $\text{event}(e) \Rightarrow \text{false}$ is active are also considered as already proved (they will be proved elsewhere; with the notations of Definition 6, e does not occur in $D_1 \vee D$, so e occurs in $\text{NonUnique}_{Q, D_1 \vee D}$), and they are replaced with $\text{find}[\text{unique}]$. That corresponds to renaming event e to a special non-unique event that is always in $\text{NonUnique}_{Q, D_1 \vee D}$. It remains to prove $\text{find}[\text{unique}_e]$ for which a query $\text{event}(e) \Rightarrow \text{false}$ is active.

Suppose that $P_0 = \text{find}[\text{unique}_e] (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat } \text{defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j \text{ then } P_j) \text{ else } P$ is such a $\text{find}[\text{unique}_e]$. (The same transformation is performed for $\text{find}[\text{unique}_e]$ terms.) CryptoVerif proves uniqueness by proving

- that we obtain a contradiction if the condition of a certain branch holds for two different values of the indices \tilde{i}_j , that is, for all $j \in \{1, \dots, m\}$, $\mathcal{F}_{P_0} \cup \{\text{defined}(M_{j_1}), \dots, \text{defined}(M_{j_{l_j}}), M_j, \text{defined}(\theta M_{j_1}), \dots, \text{defined}(\theta M_{j_{l_j}}), \theta M_j, \tilde{i}_j \neq \theta \tilde{i}_j\}$ yields a contradiction, where the substitution θ maps \tilde{i}_j to fresh replication indices;
- and that we obtain a contradiction if the conditions of two different branches hold simultaneously, that is, for all $j, j' \in \{1, \dots, m\}$ with $j < j'$, $\mathcal{F}_{P_0} \cup \{\text{defined}(M_{j_1}), \dots, \text{defined}(M_{j_{l_j}}), M_j, \text{defined}(\theta M_{j'_1}), \dots, \text{defined}(\theta M_{j'_{l_{j'}}}), \theta M_{j'}\}$ yields a contradiction, where the substitution θ maps $\tilde{i}_{j'}$ to fresh replication indices. (The substitution θ is useful in case the same replication indices are used in both branches j and j' .)

When uniqueness is proved, $\text{find}[\text{unique}_e]$ is replaced with $\text{find}[\text{unique}]$. A subsequent call to **success** (Section 4) will remove the query $\text{event}(e) \Rightarrow \text{false}$ when event e no longer occurs in the game.

Lemma 46 *The transformation **prove_unique** requires and preserves Properties 1, 2, 3, and 4. It preserves Property 5. If transformation **prove_unique** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D, \text{EvUsed}$, where p is an upper bound on the probability that a $\text{find}[\text{unique}_e]$ proved in the transformation actually executes event e .*

5.1.5 remove_assign [26]

The transformation **remove_assign** applied to an assignment $\text{let } x[i_1, \dots, i_l] : T = M \text{ in } P$ replaces x with its value M . (The same transformation is performed for assignment terms.) Precisely, the transformation is performed only when x does not occur in M (non-cyclic assignment) and M contains only variables, function applications, and tests (otherwise, copying the definition of x may break the invariant that each variable is assigned at most once). When x has several distinct definitions, we simply replace $x[i_1, \dots, i_l]$ with M in P . (For accesses to x guarded by find , we do not know which definition of x is actually used.) When x has a single definition or several identical definitions, we replace everywhere in the game $x[M_1, \dots, M_l]$ with $M\{M_1/i_1, \dots, M_l/i_l\}$. We additionally update the defined conditions of find to preserve Invariant 2 and to make sure that, if a condition of find guarantees that $x[M_1, \dots, M_l]$ is defined in the initial game, then so does the corresponding condition of find in the transformed game. (Essentially, when $y[M'_1, \dots, M'_l]$ occurs in M , the transformation typically creates new occurrences of $y[M''_1, \dots, M''_l]$ for some M''_1, \dots, M''_l , so the condition that $y[M''_1, \dots, M''_l]$ is defined must sometimes be explicitly added to conditions of find in order to preserve Invariant 2.) Moreover, we replace as often as possible defined conditions $x[M_1, \dots, M_l]$ with defined conditions

$y[M_1, \dots, M_l]$ where y is defined at the same time as x . When $x \in V$, its definition is kept unchanged. Otherwise, when x is not referred to at all after the transformation, we remove the definition of x . When x is referred to only at the root of **defined** tests, we replace its definition with a constant. (The definition point of x is important, but not its value.)

This removal of assignments is applied to all variables whose value is not used (those are used only at the root of **defined** conditions, or not at all). Depending on the argument of the transformation, it is also applied to other assignments:

- **findcond**: all assignments in conditions of **find**;
- **useless**: assignments that store a variable or a replication index, when the setting **expandAssignXY** is true; otherwise, no other assignment;
- **binder** $x_1 \dots x_n$: the assignments to variables x_1, \dots, x_n .

After this transformation, **remove_assign** calls **auto_SArename** to guarantee Property 4.

With the arguments **findcond** and **useless**, this is repeated as many times as specified by the setting **maxIterRemoveUselessAssign**. Repetition stops if a fixpoint is reached.

Lemma 47 *The transformation **remove_assign** requires and preserves Properties 1, 2, 3, and 4. It preserves Property 5. If transformation **remove_assign** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V} G', D, \text{EvUsed}$.*

Example 8 In the game G_0 of Section 4.2.1, the transformation **remove_assign binder** sk_A substitutes $\text{skgen}(rk_A)$ for sk_A in the whole process and removes the assignment $\text{let } sk_A = \text{skgen}(rk_A)$. After this substitution, $\text{sign}(\text{concat}(pk_A, x_B, x_N), sk_A, r)$ becomes $\text{sign}(\text{concat}(pk_A, x_B, x_N), \text{skgen}(rk_A), r)$ thus exhibiting a term required to apply the security assumption on signatures in the cryptographic transformation of Section 5.2.

5.1.6 use_variable

The transformation **use_variable** $x_1 \dots x_m$ tries to use variables x_1, \dots, x_m instead of recomputing their value. More precisely, at each program point that corresponds to a simple term M and where $x_j[\widetilde{M}]$ is guaranteed to be defined (because x_j is defined above that program point or directly or indirectly because of **defined** conditions above that program point), if all definitions of x_j that can be executed before reaching that program point are $\text{let } x_j[\widetilde{i}] = M_0$ in, then we test whether M is equal to $M_0\{\widetilde{M}/\widetilde{i}\}$ modulo the built-in equations, and if yes, we replace M with $x_j[\widetilde{M}]$.

The **defined** conditions of **find** above the modified program points are updated to make sure that Invariant 2 is preserved. This is needed in particular to make sure that $x_j[\widetilde{M}]$ syntactically occurs in the **defined** conditions when it is used.

Lemma 48 *The transformation **use_variable** requires and preserves Properties 1, 2, 3, and 4. It preserves Property 5. If transformation **use_variable** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V} G', D, \text{EvUsed}$.*

The transformation **use_variable** is a convenient way to perform common subexpression elimination, possibly by first inserting the definition of the desired variable(s) by **insert** (Section 5.1.13). This transformation could also be done by several applications of **replace** (Section 5.1.14). However, **use_variable** is easier to use when it performs the desired replacement.

For performance reasons, the equality tests performed by **use_variable** are considerably less powerful than those performed by **replace**, so if **use_variable** does not replace a term with $x_j[\widetilde{M}]$ at some occurrence, it is worth trying **replace**.

5.1.7 SArename [26]

The transformation **SArename** x (single assignment rename) aims at renaming x so that distinct definitions of x have different names; this is useful for distinguishing cases depending on which definition of x has set $x[\widetilde{i}]$. This transformation can be applied only when $x \notin V$. When x has $m > 1$ definitions, we rename each definition of x to a different variable x_1, \dots, x_m . Terms $x[\widetilde{i}]$ under a definition of $x_j[\widetilde{i}]$ are then replaced with $x_j[\widetilde{i}]$. Each branch of find $FB = \widetilde{u}[\widetilde{i}] = \widetilde{i}' \leq \widetilde{n}$ suchthat defined(M'_1, \dots, M'_l) \wedge M then ... where $x[M_1, \dots, M_l]$ is a subterm of some M'_k for $k \leq l'$ is replaced with m branches $FB\{x_j[M_1, \dots, M_l]/x[M_1, \dots, M_l]\}$ for $1 \leq j \leq m$.

Moreover, the implementation takes into account that some variables cannot be simultaneously defined, to reduce the number of branches of find to generate.

After this transformation, **SArename** calls **auto.SArename** to guarantee Property 4.

As a particular case, **SArename random** performs the following transformation: when y is defined by $y \stackrel{R}{\leftarrow} T$ and has $m > 1$ definitions and all variable accesses to y are of the form $y[i_1, \dots, i_l]$ under a definition of $y[i_1, \dots, i_l]$, where i_1, \dots, i_l are the current replication indices at this definition of y (that is, y has no array access using find), it renames y to y_1, \dots, y_m with a different name for each definition of y by $y \stackrel{R}{\leftarrow} T$.

Lemma 49 *The transformation **SArename** requires and preserves Properties 1, 2, 3, and 4. It preserves Property 5. If transformation **SArename** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D, \text{EvUsed}$, where p is ϵ_{find} times the maximal number of executions of a modified find that is not proved unique.*

Example 9 Consider the following process

$$\begin{aligned} \text{Start}() &:= k_A \stackrel{R}{\leftarrow} T_k; k_B \stackrel{R}{\leftarrow} T_k; \text{return} (); (Q_K \mid Q_S) \\ Q_K &= \text{foreach } i \leq n \text{ do } O[i](h : T_h, k : T_k) := \\ &\quad \text{if } h = A \text{ then let } k' : T_k = k_A \text{ in return } () \text{ else} \\ &\quad \text{if } h = B \text{ then let } k' : T_k = k_B \text{ in return } () \text{ else} \\ &\quad \text{let } k' : T_k = k \text{ in return } () \\ Q_S &= \text{foreach } i' \leq n' \text{ do } O'[i'](h' : T_h) := \\ &\quad \text{find } u = i'' \leq n \text{ suchthat defined}(h[i''], k'[i'']) \wedge h' = h[i''] \text{ then } P_1(k'[u]) \text{ else } P_2 \end{aligned}$$

The process Q_K stores in (h, k') a table of pairs (host name, key): the key for A is k_A , for B , k_B , and for any other h , the adversary can choose the key k . The process Q_S queries this table of keys to find the key $k'[u]$ of host h' , then executes $P_1(k'[u])$. If h' is not found, it executes P_2 .

By the transformation **SArename** k' , we can perform a case analysis, to distinguish the cases in which $k' = k_A$, $k' = k_B$, or $k' = k$, by renaming the three definitions of k' to k'_1 , k'_2 , and k'_3 respectively. After transformation, we obtain the following processes:

$$\begin{aligned} Q'_K &= \text{foreach } i \leq n \text{ do } O[i](h : T_h, k : T_k) := \\ &\quad \text{if } h = A \text{ then let } k'_1 : T_k = k_A \text{ in return } () \text{ else} \\ &\quad \text{if } h = B \text{ then let } k'_2 : T_k = k_B \text{ in return } () \text{ else} \\ &\quad \text{let } k'_3 : T_k = k \text{ in return } () \end{aligned}$$

$$\begin{aligned}
Q'_S &= \text{foreach } i' \leq n' \text{ do } O'[i'](h' : T_h) := \\
&\quad \text{find } u = i'' \leq n \text{ suchthat defined}(h[i''], k'_1[i'']) \wedge h' = h[i''] \text{ then } P_1(k'_1[u]) \\
&\quad \oplus u = i'' \leq n \text{ suchthat defined}(h[i''], k'_2[i'']) \wedge h' = h[i''] \text{ then } P_1(k'_2[u]) \\
&\quad \oplus u = i'' \leq n \text{ suchthat defined}(h[i''], k'_3[i'']) \wedge h' = h[i''] \text{ then } P_1(k'_3[u]) \text{ else } P_2
\end{aligned}$$

The **find** in Q_S , which looks for elements in array k' , is transformed in Q'_S into a **find** with three branches, one for each new name of k' (k'_1 , k'_2 , and k'_3 respectively). After the simplification (Section 5.1.21), Q'_S becomes:

$$\begin{aligned}
Q''_S &= \text{foreach } i' \leq n' \text{ do } O'[i'](h' : T_h) := \\
&\quad \text{find } u = i'' \leq n \text{ suchthat defined}(h[i''], k'_1[i'']) \wedge h' = A \text{ then } P_1(k_A) \\
&\quad \oplus u = i'' \leq n \text{ suchthat defined}(h[i''], k'_2[i'']) \wedge h' = B \text{ then } P_1(k_B) \\
&\quad \oplus u = i'' \leq n \text{ suchthat defined}(h[i''], k'_3[i'']) \wedge h' = h[i''] \text{ then } P_1(k[u]) \text{ else } P_2
\end{aligned}$$

since, when $k'_1[u]$ is defined, $k'_1[u] = k_A$ and $h[u] = A$, and similarly for $k'_2[u]$ and $k'_3[u]$.

5.1.8 move [26]

The transformation **move** moves random choices and assignments downwards in the code as much as possible. A random choice $x[\tilde{i}] \stackrel{R}{\leftarrow} T$ or assignment $\text{let } x[\tilde{i}] : T = M$ cannot be moved under a replication, or under a parallel composition when both sides use x , or a $\text{let } y[\tilde{i}] : T = M$ in \dots or a $\text{return } M$ when x occurs in M , or a **find** (or **if**) when the conditions use x . It can be moved under the other constructs, duplicating it if necessary, when we move it under a **find** (or **if**) that uses x in several branches. Note that when the random choice $x[\tilde{i}] \stackrel{R}{\leftarrow} T$ or assignment $\text{let } x[\tilde{i}] : T = M$ cannot be moved under a parallel composition, or a replication, it must be written above the return that is located above the considered parallel composition or replication, so that the syntax of processes is not violated. When there are array accesses to x , the random choice $x[\tilde{i}] \stackrel{R}{\leftarrow} T$ or assignment $\text{let } x[\tilde{i}] : T = M$ can be moved only inside the same oracle body, without moving it under a return or under a **find** that makes an array access to x .

The conditions above are necessary for the soundness of the move. Furthermore, the move is considered beneficial when it satisfies the following conditions:

- for random choices, when the random choice can be moved under a **find** (or **if**). When this transformation duplicates a $x[\tilde{i}] \stackrel{R}{\leftarrow} T$ by moving it under a **find** that uses x in several branches, a subsequent **SArename**(x) enables us to distinguish several cases depending in which branch x is created, which is useful in some proofs.
- for assignments, when there are no array accesses to x , the assignment to x can be moved under a **find** (or **if**), and x is used in a single branch of that **find** (or **if**). In this case, the assignment can be performed only in the branch that uses x , so it will be computed in fewer cases thanks to the move.

The performed moves are determined by the argument of the transformation:

- **all**: moves all random choices and assignments, provided the move is beneficial.
- **noarrayref**: moves all random choices and assignments that do not have array references, provided the move is beneficial.
- **random**: moves all random choices, provided the move is beneficial.

- **random_noarrayref**: moves all random choices that do not have array references, provided the move is beneficial.
- **assign**: moves all assignments, provided the move is beneficial.
- **binder** $x_1 \dots x_n$: move the variables x_1, \dots, x_n (even when the move is not beneficial).

In all cases, only random choices and assignments at the process level (not inside terms) are moved.

Lemma 50 *The transformation **move** requires and preserves Properties 1, 2, 3, and 4. It preserves Property 5. If transformation **move** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_0 G', D, \text{EvUsed}$.*

5.1.9 move array [27]

The transformation **move array** X delays the generation of a random value X until the point at which it is first used (lazy sampling). This transformation is implemented as a particular case of a cryptographic transformation by the following equivalence:

$$\begin{aligned}
& \text{foreach } i \leq n \text{ do } X \stackrel{R}{\leftarrow} T; \\
& \quad (\text{foreach } iX \leq nX \text{ do } \text{OX}() := \text{return}(X) \mid \\
& \quad \text{foreach } ieq \leq neq \text{ do } \text{Oeq}(X' : T) := \text{return}(X' = X)) \\
& \approx_{\#\text{Oeq}/|T|} [\textit{manual}] \\
& \text{foreach } i \leq n \text{ do} \\
& \quad (\text{foreach } iX \leq nX \text{ do } \text{OX}() := \\
& \quad \quad \text{find}[\text{unique}] j \leq nX \text{ suchthat defined}(Y[j]) \\
& \quad \quad \text{then return}(Y[j]) \text{ else } Y \stackrel{R}{\leftarrow} T; \text{return}(Y) \mid \\
& \quad \text{foreach } ieq \leq neq \text{ do } \text{Oeq}(X' : T) := \\
& \quad \quad \text{find}[\text{unique}] j \leq nX \text{ suchthat defined}(Y[j]) \\
& \quad \quad \text{then return}(X' = Y[j]) \text{ else return}(\text{false}))
\end{aligned}$$

where T is the type of X . Two oracles are defined, OX and Oeq . In the left-hand side, OX returns the random X itself. In the right-hand side, OX uses a lookup to test if the random value was already generated; if yes, it returns the previously generated random value $Y[j]$; if no, it generates a fresh random value Y . Transforming the left-hand side into the right-hand side therefore moves the generation of the random number X to the first call to OX , that is, the first usage of X . The oracle Oeq provides an optimized treatment of equality tests $X' = X$: when the random value X was not already generated, we return false instead of generating a fresh X , so we exclude the case that X' is equal to a fresh X . This case has probability $1/|T|$ for each call to Oeq , so the probability of distinguishing the two games is $\#\text{Oeq}/|T|$. (Notice that there never exist several choices of j that satisfy the conditions of the finds in the right-hand side of this equivalence, so these finds can be marked [unique] without modifying their behavior.)

Extension: A generalization of Oeq with other collisions is now supported by *CryptoVerif*.

5.1.10 move up

The transformation **move up** $x_1 \dots x_n$ **to** μ moves the random number generations or assignments of x_1, \dots, x_n upwards in the syntax tree, to the program point μ . This program point must be inside an oracle body.

The program point μ is an integer, which can be determined using the command **show_game** **occ**: this command displays the current game with the corresponding label $\{\mu\}$ at each program point. The command **show_game occ** also allows one to inspect the game, for instance to know the names of fresh variables created by CryptoVerif during previous transformations. Program points and variable names may depend on the version of CryptoVerif. Since CryptoVerif version 2.01, program points can also be designated by expressions like **before regexp**, which designates the program point at the beginning of the line that matches the regular expression *regexp*; **after regexp**, which designates the program point just after the line that matches *regexp*; **before_nth n regexp**, which designates the program point at the beginning of the n -th line that matches the regular expression *regexp*; **after_nth n regexp**, which designates the program point at the beginning of the first line that has an occurrence number after the n -th line that matches the regular expression *regexp*; **at n' regexp**, which designates the program point at the n' -th occurrence number that occurs inside the string that matches the regular expression *regexp* in the displayed game; **at_nth n n' regexp**, which designates the program point at the n' -th occurrence number that occurs inside the string corresponding to the n -th match of the regular expression *regexp* in the displayed game. This way of designating program points is more stable across versions of CryptoVerif.

After the game transformation, a variable x is defined at program point μ , and all other variables x_k are defined by **let** $x_k = x$ **in**. The variable x is a variable x_k itself when x_k has no array accesses and the current replication indices at the definition of x_k are the same as at μ . Otherwise, the variable x is a fresh variable.

All variables x_1, \dots, x_n must have the same type. They must not be defined syntactically above the program point μ . The definitions of the variables x_1, \dots, x_n must be in distinct branches of **if**, **find**, **let**, so that they cannot be simultaneously defined. Either all variables x_1, \dots, x_n must be defined by random number generations or all of them must be defined by assignments.

- If x_1, \dots, x_n are defined by random number generations, this transformation performs eager sampling of x_i . The random number generation of x_1, \dots, x_n must be executed at most once for each execution of program point μ . This is proved by combining that the definitions of the variables x_1, \dots, x_n are in distinct branches of **if**, **find**, **let** with the fact that each of these definitions (at μ_j) is executed at most once for each value of the current replication indices at μ . To show the latter fact, we notice that, since μ_j is syntactically under μ , the current replication indices at μ are a prefix of the replication indices at μ_j . If the replication indices at μ_j are the same as at μ , then the fact is proved. Otherwise, the replication indices at μ_j are \tilde{i}, \tilde{i}_j while the replication indices at μ are \tilde{i} and we show that $\mathcal{F}_{\mu_j} \cup \mathcal{F}_{\mu_j} \{\tilde{i}'_j / \tilde{i}_j\} \cup \{\tilde{i}_j \neq \tilde{i}'_j\}$ yields a contradiction, where \tilde{i}'_j are fresh replication indices.
- If x_1, \dots, x_n are defined by assignments of terms M_j , then all M_j must consist of variables, function applications, and tests; there must be one M_j defined at program point μ , using all defined variables collected in \mathcal{F}_μ (let M_{j_0} be that M_j , which will be used as the definition of x : **let** $x = M_{j_0}$ **in**); and all terms M_j must be equal: for all $j \neq j_0$, $M_j = M_{j_0}$ knowing the facts \mathcal{F}_{μ_j} that hold at the program point μ_j of M_j .

The defined conditions of **find** above μ are updated to syntactically guarantee the definition of M_{j_0} , as required by Invariant 2.

Lemma 51 *The transformation **move up** requires and preserves Properties 1, 2, 3, and 4. It preserves Property 5. If transformation **move up** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D, \text{EvUsed}$, where p is an upper bound of the probability of collisions eliminated*

in the proof that each μ_j is executed at most once for each execution of μ or that for all $j \neq j_0$, $M_j = M_{j_0}$.

5.1.11 move_if_fun

The transformation **move_if_fun** moves the predefined function `if_fun` or transforms it into a term `if ... then ... else ...`. It supports the following variants:

- **move_if_fun** $loc_1 \dots loc_n$, where each loc_j is either a program point or a function symbol. When loc_j is a program point, it moves occurrences of `if_fun` from inside the term at that program point to the root of that term. (The program point μ is designated as explained in Section 5.1.10.) When loc_j is a function symbol, it moves occurrences of `if_fun` from under that function symbol to just above it. The move corresponds to rewriting $C[\text{if_fun}(M_1, M_2, M_3)]$ into $\text{if_fun}(M_1, C[M_2], C[M_3])$, where C is a term context built from the following grammar:

| | |
|--|----------------------|
| $C ::=$ | simple term context |
| $[\]$ | hole |
| $x[M_1, \dots, M_{k-1}, C, M_{k+1}, \dots, M_m]$ | variable |
| $f(M_1, \dots, M_{k-1}, C, M_{k+1}, \dots, M_m)$ | function application |

and the root of C corresponds to a loc_j . (It is at program point loc_j when loc_j is a program point; its root symbol is loc_j when loc_j is a function symbol.) These moves are possible only when C is a simple term context, since otherwise they might lead to defining several times the same variable or repeating events since the context C is duplicated in the second and third arguments of `if_fun` and `if_fun` evaluates all its arguments. For simplicity, we allow them only when $C[\text{if_fun}(M_1, M_2, M_3)]$ is a simple term.

- **move_if_fun level** n , where n is a positive integer, moves occurrences of `if_fun` n function symbols up in the syntax tree (provided those `if_fun` occur under at least n function symbols). As above, these moves are allowed only when they occur inside a simple term.
- **move_if_fun to_term** $\mu_1 \dots \mu_n$ transforms terms $\text{if_fun}(M_1, M_2, M_3)$ that occur at program points μ_1, \dots, μ_n into terms `if M_1 then M_2 else M_3` . When no program point is given, it performs that transformation everywhere in the game.

When M_2 and M_3 have a visible effect, that is, they define some variable with array accesses (including by their usage in various kinds of secrecy queries) or they execute events, the transformation above would not be correct, because `if_fun(M_1, M_2, M_3)` evaluates both M_2 and M_3 while `if M_1 then M_2 else M_3` evaluates either M_2 or M_3 . In this case, we transform $\text{if_fun}(M_1, M_2, M_3)$ into `let $xcond = M_1$ in let $xthen = M_2$ in let $xelse = M_3$ in if $xcond$ then $xthen$ else $xelse$` to make sure that M_1, M_2 , and M_3 are always evaluated, and in that order.

When **autoExpand** = **true** (the default), a call to **expand** is automatically performed after **move_if_fun**, which transforms the terms `let ... = ... in ...` and `if ... then ... else ...` into processes.

Lemma 52 *The transformation **move_if_fun** requires and preserves Properties 1, 2, 3, and 4. The variants **move_if_fun** $loc_1 \dots loc_n$ and **move_if_fun level** n preserve Property 5. If transformation **move_if_fun** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V} G', D, \text{EvUsed}$.*

5.1.12 insert_event [27]

The transformation **insert_event** $e \mu$ inserts **event_abort** e at program point μ . (The program point μ is designated as explained in Section 5.1.10.)

The transformation **insert_event** $e \mu$ also adds to query **event**(e) \Rightarrow false in order to bound the probability of event e .

Lemma 53 *The transformation **insert_event** requires and preserves Properties 1, 2, 3, and 4. It preserves Property 5 if the event is inserted at the process level. If transformation **insert_event** $e \mu$ transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_0 G', D \vee e, \text{EvUsed} \cup \{e\}$.*

5.1.13 insert [27]

The transformation **insert** μins adds instruction *ins* at the program point μ . The program point μ is designated as explained in Section 5.1.10. The instruction *ins* can for instance be a test, in which case all branches of the test will be copies of the code that follows program point μ (so that the semantics of the game is unchanged). It can also be an assignment or a random generation of a fresh variable or an **event_abort** instruction. In all cases, CryptoVerif checks that this instruction preserves the semantics of the game except when we execute an inserted Shoup event, and rejects it with an error message if it does not.

After this transformation, **insert** calls **auto_SArename** to guarantee Property 4.

When the user inserts **event_abort** e , the transformation **insert** adds a query **event**(e) \Rightarrow false in order to bound the probability of event e .

When the user inserts a **find**[**unique** _{e}] (the user actually types **find**[**unique**] but CryptoVerif automatically generates a fresh event e and inserts **find**[**unique** _{e}] instead, since uniqueness is not proved yet), the transformation **insert** adds a query **event**(e) \Rightarrow false and calls **prove_unique** (Section 5.1.4) in order to try proving uniqueness.

Lemma 54 *The transformation **insert** requires and preserves Properties 1, 2, 3, 4, and 5. If transformation **insert** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D \vee e_1 \vee \dots \vee e_m, \text{EvUsed} \cup \{e_1, \dots, e_m\}$ where e_1, \dots, e_m are the events in the inserted instruction (**event_abort** e_j , **find**[**unique** _{e_j}]) and the probability p comes from **prove_unique**.*

5.1.14 replace

The transformation **replace** μM replaces the term M_0 at program point μ with the term M . (M_0 and M must be simple. The program point μ is designated as explained in Section 5.1.10.) Before performing the replacement, it checks that M is equal to M_0 at that program point (up to a small probability): first, it collects all facts \mathcal{F}_μ that hold at program point μ ; second, it tests equality between M_0 and M using \mathcal{F}_μ and built-in equations (it uses equalities inferred from \mathcal{F}_μ to replace variables with their values, trying to make the terms equal); third, it simplifies M_0 and M using user-defined rewrite rules of Section 3.1, and tests equality between the results using \mathcal{F}_μ and built-in equations; fourth, it rewrites M_0 and M at most **maxReplaceDepth** times using equalities inferred from \mathcal{F}_μ and user-defined rewrite rules of Section 3.1, until it finds a common term modulo the built-in equations. The transformation is performed as soon as the equality between M_0 and M is proved.

The defined conditions of **find** above the program point μ are updated to make sure that Invariant 2 is preserved. This is needed in particular when M makes array accesses that M_0 does not make.

Lemma 55 *The transformation **replace** requires and preserves Properties 1, 2, 3, and 4. It preserves Property 5. If transformation **replace** μ M transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D, \text{EvUsed}$, where p is an upper bound of the probability that M is different from M_0 at μ .*

The variant **assume replace** μ M performs the same replacement, without checking the equality between M_0 and M . This transformation is obviously not sound, but can be used to experiment with modifications in the games. As soon as this transformation is used, CryptoVerif does not claim that any property is proved.

5.1.15 merge_branches [27]

The transformation **merge_branches** performs the following transformations:

1. If some **then** branches of a **find[unique]** execute the same code as the **else** branch (up to renaming of variables defined in these branches and that do not have array accesses, and up to equality of terms proved using facts \mathcal{F}_μ that hold at the program point μ of the considered **find[unique]**), the index variables bound in these **then** branches have no array accesses, and the conditions of these **then** branches do not contain **event_abort** nor unproved **find[unique]_e**, then we remove these **then** branches.

Indeed, these **then** branches have the same effect as the **else** branch. The hypotheses are needed for the following reasons:

- The renamed variables must not have array accesses because renaming variables that have array accesses requires transforming these array accesses. The transformation **merge_arrays** presented in Section 5.1.16 can rename variables with array accesses.
 - The index variables bound in the removed branches must not have array accesses, because removing the definitions of these variables would modify the behavior of the array accesses.
 - The conditions must not contain **event_abort** nor unproved **find[unique]_e**, because if they do, the **find** may abort while code after transformation would not abort.
2. If all branches of **if**, **let** with pattern matching, or **find** execute the same code (up to renaming of variables defined in these branches and that do not have array accesses, and up to equality of terms proved using facts \mathcal{F}_μ that hold at the program point μ of the considered **if**, **let**, or **find**), and in case of **find**, it is not marked **[unique]_e**, the index variables bound in the **then** branches have no array accesses, and the conditions of the **then** branches do not contain **event_abort** nor unproved **find[unique]_e**, then we replace that **if** or **find** with its **else** branch.

In this transformation, we ignore the array accesses that occur in the conditions of the **find** under consideration, since these conditions will disappear after the transformation.

Furthermore, **merge_branches** applies these transformations globally to all **finds** of the game for which the simplification is possible. As a consequence, one can ignore array accesses to all variables in conditions of **find** that will be removed, so more transformations are enabled.

Lemma 56 *The transformation **merge_branches** requires and preserves Properties 1, 2, 3, 4, and 5. If transformation **merge_branches** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D, \text{EvUsed}$, where p is an upper bound on the probability that equalities between terms of merged branches do not hold.*

5.1.16 merge_arrays [27]

The transformation **merge_arrays** $x_{11} \dots x_{1n}, \dots, x_{m1} \dots x_{mn}$ merges the variables x_{j1}, \dots, x_{jn} into a single variable x_{j1} for each $j \leq m$. Each variable x_{jk} must have a single definition. For each $j \leq n$, the variables x_{j1}, \dots, x_{jn} must have the same type and indices of the same type. They must not be defined for the same value of their indices (that is, x_{jk} and $x_{jk'}$ must be defined in different branches of if or find when $k \neq k'$). The arrays x_{j1}, \dots, x_{jn} are merged into a single array x_{j1} for each $j \leq m$. The transformation proceeds as follows:

- If, for each $k \leq n$, x_{1k} is defined above x_{jk} for all $1 < j < m$, we introduce a fresh variable b_k defined by $b_k \leftarrow \text{mark}$ just after the definition of x_{1k} . We call b_k a *branch variable*; it is used to detect that x_{jk} has been defined: $x_{jk}[\tilde{M}]$ is defined before the transformation if and only if $x_{j1}[\tilde{M}]$ and $b_k[\tilde{M}]$ are defined after the transformation, and $x_{j1}[\tilde{M}]$ after the transformation is equal to $x_{jk}[\tilde{M}]$ before the transformation.
- For each find that requires that some variables x_{jk} are defined, we leave the branches that do not require the definition of x_{jk} unchanged and we try to transform the other branches $FB_l = (\tilde{u}_l = \tilde{i}_l \leq \tilde{n}_l \text{ suchthat defined}(\tilde{M}_l) \wedge M_l \text{ then } P_l)$ as follows.
 1. We require that, for each l , there exists a distinct k such that the defined condition of FB_l refers to x_{jk} for some j but not to $x_{jk'}$ for any other k' . (Otherwise, the transformation fails.) We denote by $l(k)$ the value of l that corresponds to k .
 2. We choose a “target” branch $FB_T = (\tilde{u} = \tilde{i} \leq \tilde{n} \text{ suchthat defined}(\tilde{M}) \wedge M \text{ then } P)$: if the defined condition of some branch FB_l refers to x_{j1} for some j , we choose that branch FB_l . Otherwise, we choose any branch FB_l and rename its variables x_{jk} to x_{j1} . We require that the references $x_{j1}[\tilde{M}]$ to the variables x_{j1} in the defined condition of the target branch all have the same indices \tilde{M} . If the transformation succeeds, we will replace all branches FB_l with the target branch.
 3. The branch FB_T after transformation is equivalent to branches $\bigoplus_{k=1}^n FB_T\{x_{jk}/x_{j1}, j = 1, \dots, m\}$ before transformation. We show that these branches are equivalent to the branches FB_l .

For each $k \leq n$,

- if $l(k)$ exists, then we show that $FB_T\{x_{jk}/x_{j1}, j = 1, \dots, m\}$ is equivalent to $FB_{l(k)}$. Let $l = l(k)$. We first rename the variables \tilde{u}_l of FB_l to the variables \tilde{u} of the target branch. For simplicity, we still denote by $FB_l = (\tilde{u}_l = \tilde{i}_l \leq \tilde{n}_l \text{ suchthat defined}(\tilde{M}_l) \wedge M_l \text{ then } P_l)$ the obtained branch. Then we show that, if the variables of \tilde{M}_l are defined, then the variables of $\tilde{M}\{x_{jk}/x_{j1}, j = 1, \dots, m\}$ are defined, and conversely; $M_l = M\{x_{jk}/x_{j1}, j = 1, \dots, m\}$ (knowing the equalities that hold at that program point), and P_l and $P\{x_{jk}/x_{j1}, j = 1, \dots, m\}$ execute the same code up to renaming of variables defined in P_l or $P\{x_{jk}/x_{j1}, j = 1, \dots, m\}$ and that do not have array accesses, and up to equality of terms proved using facts \mathcal{F}_μ that hold at the program point μ of the considered find.
- if $l(k)$ does not exist, then we show that $FB_T\{x_{jk}/x_{j1}, j = 1, \dots, m\}$ can in fact not be executed, because its condition cannot hold: the variables of $\tilde{M}\{x_{jk}/x_{j1}, j = 1, \dots, m\}$ cannot be simultaneously defined or $M\{x_{jk}/x_{j1}, j = 1, \dots, m\}$ cannot hold.

If the transformation above fails and we have introduced branch variables, we replace each condition $\text{defined}(x_{jk}[\tilde{M}])$ with $\text{defined}(x_{j1}[\tilde{M}], b_k[\tilde{M}])$.

If the transformation above fails and we have not introduced branch variables, the whole `merge_arrays` transformation fails.

- The definition of x_{jk} is renamed to x_{j1} and each reference to $x_{jk}[\tilde{M}]$ is renamed to $x_{j1}[\tilde{M}]$.

Lemma 57 *The transformation `merge_arrays` requires and preserves Properties 1, 2, 3, 4, and 5. If transformation `merge_arrays` transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D, \text{EvUsed}$, where p is an upper bound on the probability that required equalities do not hold.*

5.1.17 guess i

When `guessRemoveUnique = true` and some (one-session) secrecy queries are present, the transformation `guess i` first transforms the game G into G_{RU} , by replacing all proved `find[unique]` with `find`. Lemma 58 shows the soundness of this preliminary transformation. It may be advantageous for (one-session) secrecy proofs because the removed `find[unique]` do not need to be proved, while the remaining ones must be reproved after the transformation, as we show below.

Lemma 58 *Let G_{RU} be the game obtained from G by replacing all proved `find[unique]` with `find`. Let sp be a security property (sp is `1-ses.secr.(x)`, `Secrecy(x)`, `bit secr.(x)`, or a correspondence φ , which does not use S , \bar{S} , nor non-unique events). Let D be a disjunction of Shoup events and a subset of non-unique events e_i corresponding to unproved `find[unique] $_{e_i}$` in G (D does not contain S nor \bar{S}). If $\text{Bound}_{G_{\text{RU}}}(V, sp, D, p)$, then $\text{Bound}_G(V, sp, D, p)$.*

Proof Let C' be an evaluation context acceptable for $C_{sp}[G]$ with public variables $V \setminus V_{sp}$ that does not contain events used by sp or D nor non-unique events in G and \mathcal{A} be an attacker in \mathcal{BB} . Let $C = C'[C_{sp}[\cdot]]$. The context C' is also acceptable for $C_{sp}[G_{\text{RU}}]$ with public variables $V \setminus V_{sp}$, and a fortiori does not contain non-unique events in G_{RU} . Since $\text{Bound}_{G_{\text{RU}}}(V, sp, D, p)$, we have $\text{Adv}_{G_{\text{RU}}}(\mathcal{A}, C, sp, D) \leq p(\mathcal{A}, C)$.

First case: sp is a correspondence φ . We have

$$\begin{aligned} \text{Adv}_G(\mathcal{A}, C, sp, D) &= \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D}] \\ &\leq \Pr^{\mathcal{A}}[C[G_{\text{RU}}] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G_{\text{RU}},D}] = \text{Adv}_{G_{\text{RU}}}(\mathcal{A}, C, sp, D) \end{aligned}$$

because traces of G that satisfy $(\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D}$ correspond to similar traces of G_{RU} that satisfy $(\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G_{\text{RU}},D}$. Only traces that satisfy e for some proved `find[unique] $_e$` in G are mapped to different traces in G_{RU} . These traces satisfy `NonUnique $_{G,D}$` .

Second case: sp is `1-ses.secr.(x)`, `Secrecy(x)`, or `bit secr.(x)`. We have

$$\Pr^{\mathcal{A}}[C[G] : S \vee D] \leq \Pr^{\mathcal{A}}[C[G_{\text{RU}}] : S \vee D]$$

since the traces of $C[G]$ that execute an event in $S \vee D$ correspond to similar traces of $C[G_{\text{RU}}]$ that also execute an event in $S \vee D$ ($S \vee D$ does not contain any proved non-unique event of G), and

$$\Pr^{\mathcal{A}}[C[G_{\text{RU}}] : \bar{S} \vee \text{NonUnique}_{G_{\text{RU}},D}] \leq \Pr^{\mathcal{A}}[C[G] : \bar{S} \vee \text{NonUnique}_{G,D}]$$

since the traces of $C[G_{\text{RU}}]$ that execute \bar{S} or an event in `NonUnique $_{G_{\text{RU}},D}$` correspond to either to similar traces of $C[G]$ that execute the same event or to traces that execute a proved non-unique

event in $C[G]$, so an event in $\text{NonUnique}_{G,D}$. Therefore,

$$\begin{aligned} \text{Adv}_G(\mathcal{A}, C, sp, D) &= \Pr^{\mathcal{A}}[C[G] : S \vee D] - \Pr^{\mathcal{A}}[C[G] : \bar{S} \vee \text{NonUnique}_{G,D}] \\ &\leq \Pr^{\mathcal{A}}[C[G_{\text{RU}}] : S \vee D] - \Pr^{\mathcal{A}}[C[G_{\text{RU}}] : \bar{S} \vee \text{NonUnique}_{G_{\text{RU}},D}] \\ &\leq \text{Adv}_{G_{\text{RU}}}(\mathcal{A}, C, sp, D). \end{aligned}$$

In both cases, we obtain $\text{Adv}_G(\mathcal{A}, C, sp, D) \leq p(\mathcal{A}, C)$ and $\text{Bound}_G(V, sp, D, p)$. \square

Next, the main guessing transformation is performed. The transformation **guess** i consists in guessing the tested session of a principal in a protocol, which is a step frequently done in cryptographic proofs. In `CryptoVerif`, we consider a game G and define a transformed game G' by guessing a replication index i : we replace `foreach` $i \leq n$ `do` Q with `foreach` $i \leq n$ `do` Q' where Q' is obtained from Q by replacing the processes P under the first inputs with `if` $i = i_{\text{tested}}$ `then` P' `else` P'' and i_{tested} is a constant. The constant i_{tested} is the index of the tested session. We distinguish the process executed in the tested session, P' , on which we are going to prove security properties, from the process P'' for other sessions which are executed, but for which we do not prove security properties. (In case **diff_constants** = **true**, the constant i_{tested} must not be considered different from other constants of the same type.) The process P' is obtained from P by

- duplicating all events: `event` $e(\widetilde{M})$ is replaced with `event` $e(\widetilde{M})$; `event` $e'(\widetilde{M})$ and similarly `event_abort` e is replaced with `event` e ; `event_abort` e' . We require that in the game G , the same event e cannot occur both under the modified replication `foreach` $i \leq n$ `do` Q and elsewhere in the game. (Otherwise, queries that use e are left unchanged.)
- duplicating definitions of every variable x used in queries for secrecy and one-session secrecy: `let` $x' = x$ `in` is added after each definition of x . We require that in the game G , the same variable x used in queries for secrecy or one-session secrecy cannot be defined both under the modified replication `foreach` $i \leq n$ `do` Q and elsewhere in the game. (Otherwise, the considered query is left unchanged.)

The process P'' is obtained from P by duplicating definitions of every variable x used in queries for secrecy (not one-session secrecy): `let` $x'' = x$ `in` is added after each definition of x .

We replace variables x in secrecy and one-session secrecy queries with their duplicated version x' . For secrecy queries, the duplicated version x'' is added to public variables. In both cases, we prove (one-session) secrecy for the variable x' defined in the tested session. In case of one-session secrecy, that is enough: it shows that x' is indistinguishable from a random value, and that proves one-session secrecy of x for all sessions by symmetry. However, for secrecy, we additionally want to show that the values of x in the various sessions are independent of each other; this is achieved by considering the value of x in sessions other than the tested session (that is, x'') as public: if x' is indistinguishable from random even when x'' is public, then x' is independent of x'' .

In non-injective correspondence queries, we replace *one* non-injective event e before the arrow \Rightarrow with its duplicated version e' . Hence, we prove the query for the tested session, which uses event e' . The proof is valid for all sessions by symmetry.

The probability of attack must basically be multiplied by n for all modified queries. The proof depends on the considered query and is detailed below.

For queries that are left unchanged, i.e. secrecy and one-session secrecy queries for variables not defined under the modified replication, non-injective correspondence queries with no event before the arrow \Rightarrow under the modified replication (the previous queries prove properties about other roles than the one for which we guess the tested session), bit secrecy queries (because

the secret is defined under no replication, so it is not under the guessed replication), as well as injective correspondence queries (see details below), the probability is unchanged. It is clear that these queries are not affected by the transformation.

After this transformation, **guess** calls **auto_SArename** to guarantee Property 4.

Lemma 59 *The transformation **guess** i requires and preserves Properties 1, 2, 3, and 4. It preserves Property 5.*

*Suppose the game G is transformed into G' by the transformation **guess** i , where i is a replication index bounded by n . Below, we consider only the modified queries.*

Let φ and φ' be respectively the semantics of a non-injective correspondence and its transformed correspondence. If $\text{Bound}_{G'}(V, \varphi', D_{\text{false}}, p)$ and p is independent of the value of i_{tested} , then $\text{Bound}_G(V, \varphi, D_{\text{false}}, np)$.

If G' satisfies the one-session secrecy of x' with public variables V ($x, x', x'' \notin V$) up to probability p and p is independent of the value of i_{tested} , then G satisfies the one-session secrecy of x with public variables V up to probability np .

If G' satisfies the secrecy of x' with public variables $V \cup \{x''\}$ ($x, x', x'' \notin V$) up to probability p and p satisfies Property 6, then G satisfies the secrecy of x with public variables V up to probability $n \times p$ (neglecting a small additional runtime of the context).

If $\text{Bound}_{G'}(V \cup \{x'\}, \text{1-ses.secr.}(x'), \text{NonUnique}_{G'}, p)$ and p satisfies Property 6, then we have $\text{Bound}_G(V \cup \{x\}, \text{1-ses.secr.}(x), D_{\text{false}}, np)$.

If $\text{Bound}_{G'}(V \cup \{x', x''\}, \text{Secrecy}(x'), \text{NonUnique}_{G'}, p)$ and p satisfies Property 6, then we have $\text{Bound}_G(V \cup \{x\}, \text{Secrecy}(x), D_{\text{false}}, np)$ (neglecting a small additional runtime of the context).

Property 6 guarantees that p is independent of the value of i_{tested} , as well as other independence conditions needed for secrecy and for the properties on **Bound** because we modify the context in the proof. Since the third argument of **Bound** is always D_{false} in the conclusion of Lemma 59, we cannot use the optimization of considering the disjunction of several properties simultaneously, as outlined in Section 2.8.4: we must consider each property and event separately. Indeed, if we applied guessing to several properties at once, we might need to guess the tested session for each property, which would introduce several factors n . Since the third argument of **Bound** is **NonUnique** $_{G'}$ in the hypothesis of Lemma 59 for (one-session) secrecy properties, uniqueness of $\text{find}[\text{unique}_e]$ must be reproved in game G' after the guess transformation (that is, the probability that these $\text{find}[\text{unique}_e]$ have several successful choices must be bounded again in G').

Proof

Non-injective correspondences We suppose that the events under the transformed replication contain as argument the replication index i of that replication. (CryptoVerif implicitly adds the current program point and replication indices to each event, and uses fresh distinct variables for the added replication indices in the queries. That does not change the meaning of the query.)

Let $\forall i_0 : [1, n], \tilde{x} : \tilde{T}; \text{event}(e(\tilde{M})) \wedge \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$ be the initial query and $\forall i_0 : [1, n], \tilde{x} : \tilde{T}; \text{event}(e'(\tilde{M})) \wedge \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$ be the transformed query, where $\{i_0, \tilde{x}\} = \text{var}(\text{event}(e'(\tilde{M})) \wedge \psi)$, $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\text{event}(e'(\tilde{M})) \wedge \psi)$, and i_0 be the variable for the index of the transformed replication in \tilde{M} .

Let C be an evaluation context acceptable for G with public variables V that does not contain

events used by φ and \mathcal{A} be an attacker in \mathcal{BB} .

$$\begin{aligned}
\text{Adv}_G(\mathcal{A}, C, \varphi, D_{\text{false}}) &= \Pr^{\mathcal{A}}[C[G] : \neg\varphi \wedge \neg\text{NonUnique}_G] \\
&= \Pr^{\mathcal{A}} \left[C[G] : (\exists i_0 \in [1, n], \exists \tilde{x} \in \tilde{T}, \text{event}(e(\tilde{M}))) \right. \\
&\quad \left. \wedge \psi \wedge \neg \exists \tilde{y} \in \tilde{T}', \phi) \wedge \neg\text{NonUnique}_G \right] \\
&= \sum_{i_{\text{val}}=1}^n \Pr^{\mathcal{A}} \left[C[G] : (\exists i_0 \in [1, n], \exists \tilde{x} \in \tilde{T}, i_0 = i_{\text{val}} \wedge \text{event}(e(\tilde{M}))) \right. \\
&\quad \left. \wedge \psi \wedge \neg \exists \tilde{y} \in \tilde{T}', \phi) \wedge \neg\text{NonUnique}_G \right] \\
&= \sum_{i_{\text{val}}=1}^n \Pr^{\mathcal{A}} \left[C[G'] : (\exists i_0 \in [1, n], \exists \tilde{x} \in \tilde{T}, \text{event}(e'(\tilde{M}))) \right. \\
&\quad \left. \wedge \psi \wedge \neg \exists \tilde{y} \in \tilde{T}', \phi) \wedge \neg\text{NonUnique}_{G'} \right] \text{ for } i_{\text{tested}} = i_{\text{val}} \\
&= \sum_{i_{\text{val}}=1}^n \text{Adv}_{G'}(\mathcal{A}, C, \varphi', D_{\text{false}}) \text{ for } i_{\text{tested}} = i_{\text{val}}
\end{aligned}$$

Since $\text{Bound}_{G'}(V, \varphi', D_{\text{false}}, p)$, we have $\text{Adv}_{G'}(\mathcal{A}, C, \varphi', D_{\text{false}}) \leq p(\mathcal{A}, C)$ and the probability p is independent of the value of i_{tested} , so we obtain $\text{Adv}_G(\mathcal{A}, C, \varphi, D_{\text{false}}) \leq n \times p(\mathcal{A}, C)$. Therefore $\text{Bound}_G(V, \varphi, D_{\text{false}}, np)$.

One-session secrecy Let C be an evaluation context acceptable for $C_{1\text{-ses.secr.}(x)}[G]$ with public variables V that does not contain S nor $\bar{\mathsf{S}}$ and \mathcal{A} be an attacker in \mathcal{BB} . Suppose that the modified replication corresponds to the j -th index of variable x . We have

$$\begin{aligned}
\text{Adv}_G^{1\text{-ses.secr.}(x)}(\mathcal{A}, C) &= \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x)}[G]] : \mathsf{S}] - \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x)}[G]] : \bar{\mathsf{S}}] \\
&= \sum_{v=1}^n \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x)}[G]] : \mathsf{S} \wedge u_j = v] \\
&\quad - \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x)}[G]] : \bar{\mathsf{S}} \wedge u_j = v] \\
&= \sum_{v=1}^n \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \mathsf{S} \wedge u_j = v] \\
&\quad - \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \bar{\mathsf{S}} \wedge u_j = v] \text{ for } i_{\text{tested}} = v
\end{aligned}$$

because in G' with $i_{\text{tested}} = v$, $x'[u_1, \dots, u_m] = x[u_1, \dots, u_m]$ when $u_j = v$, so $C_{1\text{-ses.secr.}(x)}[G]$ behaves like $C_{1\text{-ses.secr.}(x')}[G']$ (the events added in G' are not used).

Moreover, $\Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \mathsf{S} \wedge (u_j \text{ not defined} \vee u_j \neq v)] = \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \bar{\mathsf{S}} \wedge (u_j \text{ not defined} \vee u_j \neq v)]$ when $i_{\text{tested}} = v$. Indeed, when u_j is not defined or $u_j \neq v = i_{\text{tested}}$, the query on O_s either is not executed or always yields, independently of the value of b . Hence, changing the value of b just swaps the events S and $\bar{\mathsf{S}}$. So

$$\begin{aligned}
&\Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \mathsf{S} \wedge (u_j \text{ not defined} \vee u_j \neq v) \wedge b = \text{true}] = \\
&\quad \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \bar{\mathsf{S}} \wedge (u_j \text{ not defined} \vee u_j \neq v) \wedge b = \text{false}] \\
&\Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \bar{\mathsf{S}} \wedge (u_j \text{ not defined} \vee u_j \neq v) \wedge b = \text{true}] = \\
&\quad \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \mathsf{S} \wedge (u_j \text{ not defined} \vee u_j \neq v) \wedge b = \text{false}].
\end{aligned}$$

We obtain the announced result by swapping the two sides of the second equality and adding the first equality to it. (The variable b is always defined when S or $\bar{\mathsf{S}}$ is executed.)

Therefore,

$$\begin{aligned} \text{Adv}_G^{1\text{-ses.secr.}(x)}(\mathcal{A}, C) &= \sum_{v=1}^n \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \mathbb{S}] - \Pr^{\mathcal{A}}[C[C_{1\text{-ses.secr.}(x')}[G']] : \bar{\mathbb{S}}] \\ &\quad \text{for } i_{\text{tested}} = v \\ &= \sum_{v=1}^n \text{Adv}_{G'}^{1\text{-ses.secr.}(x')}(\mathcal{A}, C) \text{ for } i_{\text{tested}} = v \end{aligned}$$

Since G' satisfies the one-session secrecy of x' with public variables V up to probability p , we have $\text{Adv}_{G'}^{1\text{-ses.secr.}(x')}(\mathcal{A}, C) \leq p(\mathcal{A}, C)$ and the probability p is independent of the value of i_{tested} , so we obtain $\text{Adv}_G^{1\text{-ses.secr.}(x)}(\mathcal{A}, C) \leq n \times p(\mathcal{A}, C)$. Therefore, G satisfies the one-session secrecy of x with public variables V up to probability np .

Secrecy Let C be an evaluation context acceptable for $C_{\text{Secrecy}(x)}[G]$ with public variables V that does not contain \mathbb{S} nor $\bar{\mathbb{S}}$ and \mathcal{A} be an attacker in \mathcal{BB} . Suppose that the modified replication corresponds to the j -th index of variable x . We have

$$\begin{aligned} \text{Adv}_G^{\text{Secrecy}(x)}(\mathcal{A}, C) &= \frac{1}{2} \left(\Pr^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x)}] : \mathbb{S} \mid b = \text{true}] + \Pr^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x)}] : \mathbb{S} \mid b = \text{false}] \right. \\ &\quad \left. - \Pr^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x)}] : \bar{\mathbb{S}} \mid b = \text{true}] - \Pr^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x)}] : \bar{\mathbb{S}} \mid b = \text{false}] \right) \end{aligned}$$

Let

$$\begin{aligned} Q_{\text{Secrecy}(x), \text{real}} &= O_{s0}() := \text{return } (); \\ &\quad (\text{foreach } i_s \leq n_s \text{ do } O_s[i_s](u_1 : [1, n_1], \dots, u_m : [1, n_m]) := \\ &\quad \quad \text{if defined}(x[u_1, \dots, u_m]) \text{ then} \\ &\quad \quad \text{return } (x[u_1, \dots, u_m]) \\ &\quad \quad | O'_s(b' : \text{bool}) := \text{if } b' \text{ then event_abort } \mathbb{S} \text{ else event_abort } \bar{\mathbb{S}}) \\ Q_{\text{Secrecy}(x), \text{random}} &= O_{s0}() := \text{return } (); \\ &\quad (\text{foreach } i_s \leq n_s \text{ do } O_s[i_s](u_1 : [1, n_1], \dots, u_m : [1, n_m]) := \\ &\quad \quad \text{if defined}(x[u_1, \dots, u_m]) \text{ then} \\ &\quad \quad \text{find } u'_s = i'_s \leq n_s \text{ suchthat defined}(y[i'_s], u_1[i'_s], \dots, u_m[i'_s]) \wedge \\ &\quad \quad \quad u_1[i'_s] = u_1 \wedge \dots \wedge u_m[i'_s] = u_m \\ &\quad \quad \text{then return } (y[u'_s]) \\ &\quad \quad \text{else } y \stackrel{R}{\leftarrow} T; \text{return } (y) \\ &\quad \quad | OS'(b' : \text{bool}) := \text{if } b' \text{ then event_abort } \mathbb{S} \text{ else event_abort } \bar{\mathbb{S}}) \end{aligned}$$

Then

$$\text{Adv}_G^{\text{Secrecy}(x)}(\mathcal{A}, C) = \frac{1}{2} \left(\Pr^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x), \text{real}}] : \mathbb{S}] + \Pr^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x), \text{random}}] : \bar{\mathbb{S}}] \right. \\ \left. - \Pr^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x), \text{real}}] : \bar{\mathbb{S}}] - \Pr^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x), \text{random}}] : \mathbb{S}] \right)$$

Let

$$\begin{aligned}
Q_{\text{Secrecy}(x),v} = O_{s0}() &:= \text{return } (); \\
&\text{(foreach } i_s \leq n_s \text{ do } O_s[i_s](u_1 : [1, n_1], \dots, u_m : [1, n_m]) := \\
&\quad \text{if defined}(x[u_1, \dots, u_m]) \text{ then} \\
&\quad \text{if } u_j \leq v \text{ then return } (x[u_1, \dots, u_m]) \text{ else} \\
&\quad \text{find } u'_s = i'_s \leq n_s \text{ suchthat defined}(y[i'_s], u_1[i'_s], \dots, u_m[i'_s]) \wedge \\
&\quad \quad u_1[i'_s] = u_1 \wedge \dots \wedge u_m[i'_s] = u_m \\
&\quad \text{then return } (y[u'_s]) \\
&\quad \text{else } y \stackrel{R}{\leftarrow} T; \text{return } (y) \\
| O'_s(b' : \text{bool}) &:= \text{if } b' \text{ then event_abort } S \text{ else event_abort } \bar{S}
\end{aligned}$$

The process $Q_{\text{Secrecy}(x),0}$ behaves like $Q_{\text{Secrecy}(x),\text{random}}$ and $Q_{\text{Secrecy}(x),n}$ behaves like $Q_{\text{Secrecy}(x),\text{real}}$ ($n_j = n$). Therefore,

$$\begin{aligned}
\text{Adv}_G^{\text{Secrecy}(x)}(\mathcal{A}, C) &= \frac{1}{2} \left(\begin{aligned} &(\text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),n}] : S] - \text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),0}] : S]) \\ &- (\text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),n}] : \bar{S}] - \text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),0}] : \bar{S}]) \end{aligned} \right) \\
&= \frac{1}{2} \left(\begin{aligned} &\sum_{v=1}^n (\text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),v}] : S] - \text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),v-1}] : S]) \\ &- \sum_{v=1}^n (\text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),v}] : \bar{S}] - \text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),v-1}] : \bar{S}]) \end{aligned} \right)
\end{aligned}$$

We define a context C'_v that returns a random value for $u_j > v$, the real value of x obtained from the public variable x'' in G' for $u_j < v$, and calls $Q_{\text{Secrecy}(x'),\text{real}}$ or $Q_{\text{Secrecy}(x'),\text{random}}$ for $u_j = v$.

$$\begin{aligned}
C'_v = \text{newOracle } O_{s1}; ([\mid \text{foreach } i_s \leq n_s \text{ do } O_s[i_s](u_1 : [1, n_1], \dots, u_m : [1, n_m]) := \\
&\quad \text{if } u_j = v \text{ then (let } z : T = O_{s1}[i_s](u_1, \dots, u_m) \text{ in return } (z)) \text{ else} \\
&\quad \text{if defined}(x''[u_1, \dots, u_m]) \text{ then} \\
&\quad \text{if } u_j < v \text{ then return } (x''[u_1, \dots, u_m]) \text{ else} \\
&\quad \text{find } u'_s = i'_s \leq n_s \text{ suchthat defined}(y[i'_s], u_1[i'_s], \dots, u_m[i'_s]) \wedge \\
&\quad \quad u_1[i'_s] = u_1 \wedge \dots \wedge u_m[i'_s] = u_m \\
&\quad \text{then return } (y[u'_s]) \\
&\quad \text{else } y \stackrel{R}{\leftarrow} T; \text{return } (y)
\end{aligned}$$

where the processes $Q_{\text{Secrecy}(x'),\text{real}}$ and $Q_{\text{Secrecy}(x'),\text{random}}$ use oracle O_{s1} instead of O_s . When $u_j = v$, the query on $O_s[i_s]$ is forwarded to $Q_{\text{Secrecy}(x'),\text{real}}$ (resp. $Q_{\text{Secrecy}(x'),\text{random}}$) on oracle $O_{s1}[i_s]$. The values of x for sessions other than v are collected in x'' by G' ; these are the values returned by the query on $O_s[i_s]$ when $u_j < v$. Finally, when $u_j > v$, the query on $O_s[i_s]$ is answered with a random value y .

Then

$$\begin{aligned}
\text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),v}] : S] &= \text{Pr}^{\mathcal{A}}[C[C'_v[G'] \mid Q_{\text{Secrecy}(x'),\text{real}}] : S], \\
\text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),v}] : \bar{S}] &= \text{Pr}^{\mathcal{A}}[C[C'_v[G'] \mid Q_{\text{Secrecy}(x'),\text{real}}] : \bar{S}], \\
\text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),v-1}] : S] &= \text{Pr}^{\mathcal{A}}[C[C'_v[G'] \mid Q_{\text{Secrecy}(x'),\text{random}}] : S], \text{ and} \\
\text{Pr}^{\mathcal{A}}[C[G \mid Q_{\text{Secrecy}(x),v-1}] : \bar{S}] &= \text{Pr}^{\mathcal{A}}[C[C'_v[G'] \mid Q_{\text{Secrecy}(x'),\text{random}}] : \bar{S}],
\end{aligned}$$

for $i_{\text{tested}} = v$.

Then

$$\begin{aligned} & \text{Adv}_G^{\text{Secrecy}(x)}(\mathcal{A}, C) \\ &= \sum_{v=1}^n \frac{1}{2} \left(\Pr^{\mathcal{A}}[C[C'_v][G' \mid Q_{\text{Secrecy}(x'), \text{real}}] : \mathbf{S}] - \Pr^{\mathcal{A}}[C[C'_v][G' \mid Q_{\text{Secrecy}(x'), \text{random}}] : \mathbf{S}] \right) \\ & \quad \text{for } i_{\text{tested}} = v \\ &= \sum_{v=1}^n \text{Adv}_{G'}^{\text{Secrecy}(x')}(\mathcal{A}, C[C'_v]) \text{ for } i_{\text{tested}} = v \end{aligned}$$

by the link between the advantage for secrecy and $Q_{\text{Secrecy}(x'), \text{real}}$, $Q_{\text{Secrecy}(x'), \text{random}}$ shown above. Since G' satisfies the secrecy of x' with public variables $V \cup \{x''\}$ up to probability p and $C[C'_v]$ is an evaluation context acceptable for $G' \mid Q_{\text{Secrecy}(x')}$ with public variables $V \cup \{x''\}$ that does not contain \mathbf{S} nor $\bar{\mathbf{S}}$, we have $\text{Adv}_{G'}^{\text{Secrecy}(x')}(\mathcal{A}, C[C'_v]) \leq p(\mathcal{A}, C[C'_v])$. Moreover, by Property 6, the probability $p(\mathcal{A}, C[C'_v])$ is independent of $i_{\text{tested}} = v$ (because the type of i_{tested} is bounded) and depends only on the runtime of C , the number of calls C makes to the various oracles (which determine replication bounds), and the length of bitstrings, so we have $p(\mathcal{A}, C[C'_v]) = p(\mathcal{A}, C)$ (the runtime of C'_v can be neglected). Hence, we obtain $\text{Adv}_G^{\text{Secrecy}(x)}(\mathcal{A}, C) \leq n \times p(\mathcal{A}, C)$. Therefore, G satisfies the secrecy of x with public variables V up to probability $n \times p$.

One-session secrecy, secrecy, and bit secrecy In this lemma, bit secrecy queries do not occur. However, we reuse this proof in Lemmas 60 and 61 where bit secrecy queries occur, so we also handle them here. Let sp be 1-ses.secr.(x), Secrecy(x), or bit secr.(x), and sp' be the same property with x' instead of x . Let $V' = V$ when sp is 1-ses.secr.(x) or bit secr.(x), and $V' = V \cup \{x''\}$ when sp is Secrecy(x). Since $\text{Bound}_{G'}(V' \cup \{x'\}, sp', \text{NonUnique}_{G'}, p)$, then by Lemma 28, Property 1, G' satisfies sp' with public variables V' up to probability p' such that $p'(\mathcal{A}, C) = p(\mathcal{A}, C[C_{sp'}])$. So by the previous result for (one-session or bit) secrecy, G satisfies sp with public variables V up to probability np' . Let C be an evaluation context acceptable for $C_{sp}[G]$ with public variables V that does not contain \mathbf{S} nor $\bar{\mathbf{S}}$ and \mathcal{A} be an attacker in \mathcal{BB} . Hence

$$\begin{aligned} & \text{Adv}_G(\mathcal{A}, C[C_{sp}[]], sp, D_{\text{false}}) \\ &= \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathbf{S}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathbf{S}} \vee \text{NonUnique}_G] \\ &\leq \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathbf{S}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathbf{S}}] \\ &= \text{Adv}_G^{sp}(\mathcal{A}, C) \leq n \times p'(\mathcal{A}, C) = n \times p(\mathcal{A}, C[C_{sp}[]]) \end{aligned}$$

Indeed, replacing $C_{sp'}$ with C_{sp} in the argument of p does not change its result, by Property 6. Therefore, we have $\text{Bound}_G(V \cup \{x\}, sp, D_{\text{false}}, np)$. \square

In the **guess** transformation, we cannot modify injective correspondence queries, because two executions of some injective event e with different indices i could be mapped to the same events in the conclusion of the query. In general, it even does not work for non-injective events inside injective queries. As a counter-example, consider the query: $\forall i : [1, n], x : T'; \text{event}(e_1(i, x)) \wedge \text{inj-event}(e_2(x)) \Rightarrow \text{inj-event}(e_3())$ with events $e_1(i_1, x_1)$, $e_1(i_2, x_2)$, $e_2(x_1)$, $e_2(x_2)$ and e_3 each executed once. This query is false: we have two executions of e_2 (with matching executions of e_1) for a single execution of e_3 . That contradicts injectivity. However, it is true if we restrict ourselves to one value of i (the index of the tested session), because we consider $e_1(i_1, x_1)$, $e_2(x_1)$ and e_3 for $i = i_1$ and $e_1(i_2, x_2)$, $e_2(x_2)$ and e_3 for $i = i_2$. Requiring the same value of x in e_1 and

e_2 restricts the events e_2 that we consider when we guess the session for e_1 . Therefore, proving the query for the tested session does not allow us to prove it in the initial game.

Hence, for injective correspondence queries $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$, we proceed as follows: we define a non-injective query $\text{noninj}(\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi)$ simply obtained by replacing injective events with non-injective events, and we try to prove that $\text{noninj}(\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi)$ implies $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$ in the current game. This proof is a modified version of the proof of injective queries (Section 4.2.4): we define the pseudo-formula $\mathcal{C}(\psi, \phi)$ by

$$\mathcal{C}(\psi, M) = \perp$$

$$\mathcal{C}(\psi, \text{event}(e(\tilde{M}))) = \perp$$

$$\mathcal{C}(\psi, \text{inj-event}(e(\tilde{M}))) = \mathcal{S} \text{ such that assuming } \psi = F_1 \wedge \dots \wedge F_m,$$

for every μ_1 that executes F_1, \dots , for every μ_m that executes F_m , letting

$$\mathcal{F} = \theta_1 \mathcal{F}_{F_1, \mu_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m},$$

$$\mathcal{I} = \{j \mapsto (\mu_j, \theta_j I_{\mu_j}) \mid F_j \text{ is an injective event}\},$$

$$\mathcal{V} = \text{var}(\theta_1 I_{\mu_1}) \cup \dots \cup \text{var}(\theta_m I_{\mu_m}) \cup \{\tilde{x}, \tilde{y}\},$$

where for $j \leq m$, θ_j is a renaming of I_{μ_j} to fresh replication indices,

we have $(\mathcal{F}, (\tilde{M}), \mathcal{I}, \mathcal{V}) \in \mathcal{S}$.

$$\mathcal{C}(\psi, \phi_1 \wedge \phi_2) = \mathcal{C}(\psi, \phi_1) \wedge \mathcal{C}(\psi, \phi_2)$$

$$\mathcal{C}(\psi, \phi_1 \vee \phi_2) = \mathcal{C}(\psi, \phi_1) \vee \mathcal{C}(\psi, \phi_2)$$

Given a pseudo-formula \mathcal{C} , we define $\vdash \mathcal{C}$ as in Section 4.2.4.

Proposition 4 *Let $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$ be a correspondence, with $\tilde{x} = \text{var}(\psi)$ and $\tilde{y} = \text{var}(\phi) \setminus \text{var}(\psi)$. Let $\varphi = \llbracket \forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket$ be the semantics the correspondence $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$ (Definition 12), and $\varphi_{\text{ni}} = \llbracket \text{noninj}(\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi) \rrbracket$ be the semantics of the correspondence $\text{noninj}(\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi)$. Let Q_0 be a process that satisfies Properties 3 and 4. Suppose that, in Q_0 , the arguments of the events that occur in ψ are always simple terms.*

Assume that $\vdash \mathcal{C}(\psi, \phi)$ and for all $\mathcal{A} \in \mathcal{BB}$ and all evaluation contexts C acceptable for Q_0 , $\Pr^{\mathcal{A}}[C[Q_0] \preceq \neg \{\vdash \mathcal{C}(\psi, \phi)\}] \leq p(\mathcal{A}, C)$. If $\text{Bound}_{Q_0}(V, \varphi_{\text{ni}}, D_{\text{false}}, p')$, then $\text{Bound}_{Q_0}(V, \varphi, D_{\text{false}}, p' + p)$.

Proposition 4 proves injectivity much like in Section 4.2.4, but using the arguments \tilde{M} of the events in ϕ instead of the program points and replication indices at their execution (which we do not have since in Q_0 we are not able to prove that these events have been executed; otherwise we would simply prove the correspondence in Q_0). Intuitively, $\vdash \mathcal{C}(\psi, \phi)$ shows that, if we have different executions of injective events in ψ , that is, executions of such events with different pairs (program point, replication indices), then the arguments \tilde{M} of each injective event in ϕ must be different, which implies different executions of this event.

Proof Let C be an evaluation context acceptable for Q_0 with public variables V that does not contain events used by φ and \mathcal{A} be an attacker in \mathcal{BB} . Let $\mathcal{C} = \mathcal{C}(\psi, \phi)$. Consider a trace Tr of $C[Q_0]$ for \mathcal{A} such that $Tr \vdash \varphi_{\text{ni}}$, $Tr \vdash \{\vdash \mathcal{C}\}$, Tr does not execute a non-unique event of Q_0 , and the last configuration of Tr cannot be reduced. Let $\mu \mathcal{E}v$ be the sequence of events in the last configuration of Tr . Since $\mu \mathcal{E}v \vdash \varphi_{\text{ni}}$,

$$\mu \mathcal{E}v \vdash \forall \tau_1, \dots, \tau_m \in \mathbb{N}, \forall \tilde{x} \in \tilde{T}, (\psi^\tau \Rightarrow \exists \tilde{y} \in \tilde{T}', \phi)$$

with the notations of Definition 12. By defining functions $f_j \in \mathbb{N}^m \times \prod \tilde{T} \rightarrow \mathbb{N} \cup \{\perp\}$ that map $\tau_1, \dots, \tau_m, \tilde{x}$ to the execution steps of injective events in the proof of ϕ , and to \perp when the

event is not used in the proof of ϕ , we have

$$\mu\mathcal{E}v \vdash \exists f_1, \dots, f_k \in \mathbb{N}^m \times \prod \tilde{T} \rightarrow \mathbb{N} \cup \{\perp\}, \forall \tau_1, \dots, \tau_m \in \mathbb{N}, \forall \tilde{x} \in \tilde{T}, (\psi^\tau \Rightarrow \exists \tilde{y} \in \tilde{T}', \phi^\tau) \quad (32)$$

It remains to show $\text{Inj}(I, f)$ for all $f \in \{f_1, \dots, f_k\}$.

Let $f \in \{f_1, \dots, f_k\}$, and $\text{event}(e(\tilde{M}))@f(\tau_1, \dots, \tau_m, \tilde{x})$ be the event labeled with f in ϕ^τ . Suppose that $f(\tau'_1, \dots, \tau'_m, \tilde{a}') = f(\tau''_1, \dots, \tau''_m, \tilde{a}'') \neq \perp$ and there exists $j \in I$ such that $\tau'_j \neq \tau''_j$, and let us prove a contradiction.

Since $f(\tau'_1, \dots, \tau'_m, \tilde{a}') \neq \perp$, $\text{event}(e(\tilde{M}))@f(\tau'_1, \dots, \tau'_m, \tilde{a}')$ is used in the proof of (32), so letting $\rho_1 = \{\tau_1 \mapsto \tau'_1, \dots, \tau_m \mapsto \tau'_m, \tilde{x} \mapsto \tilde{a}'\}$, $\rho_1, \mu\mathcal{E}v \vdash \psi^\tau$ and there exists an extension ρ'_1 of ρ_1 to \tilde{y} such that $\rho'_1, \mu\mathcal{E}v \vdash \phi^\tau$ and $\rho'_1, \mu\mathcal{E}v \vdash \text{event}(e(\tilde{M}))@f(\tau_1, \dots, \tau_m, \tilde{x})$.

Since $\rho_1, \mu\mathcal{E}v \vdash \psi^\tau$, for all events $F_\ell = \text{event}(e_\ell(\tilde{M}_\ell))@f_\ell$ in ψ^τ , $Tr, \rho_1 \vdash F_\ell$ and $\mu\mathcal{E}v(\tau'_\ell) = (\dots) : e_\ell(\tilde{a}_{\ell,1})$ for $\tilde{a}_{\ell,1}$ such that $\rho_1, \tilde{M}_\ell \Downarrow \tilde{a}_{\ell,1}$. By Lemma 38, there exists a program point $\mu_{\ell,1}$ that executes F_ℓ (in Q_0) and a case $c_{\ell,1}$ such that, for any $\theta_{\ell,1}$ renaming of $I_{\mu_{\ell,1}}$ to fresh replication indices, there exists a mapping $\sigma_{\ell,1}$ with domain $\theta_{\ell,1}I_{\mu_{\ell,1}}$ such that $\mu\mathcal{E}v(\rho_1(\tau_\ell)) = (\mu_{\ell,1}, \sigma_{\ell,1}(\theta_{\ell,1}I_{\mu_{\ell,1}})) : \dots$ and $Tr, \sigma_{\ell,1} \cup \rho_1 \vdash \theta_{\ell,1}\mathcal{F}_{F_\ell, \mu_{\ell,1}, c_{\ell,1}}$. So $\mu\mathcal{E}v(\tau'_\ell) = (\mu_{\ell,1}, \sigma_{\ell,1}(\theta_{\ell,1}I_{\mu_{\ell,1}})) : e_\ell(\tilde{a}_{\ell,1})$.

Let \mathcal{S} be the label of \mathcal{C} at the occurrence corresponding to f , and $\text{inj-event}(e(\tilde{M}))$ be the injective event at that occurrence in ϕ . Let $\mathcal{F}_1 = \bigcup \theta_{\ell,1}\mathcal{F}_{F_\ell, \mu_{\ell,1}, c_{\ell,1}}$, $\mathcal{I}_1 = \{\ell \mapsto (\mu_{\ell,1}, \theta_{\ell,1}I_{\mu_{\ell,1}}) \mid F_\ell \text{ is an injective event}\}$, and $\mathcal{V}_1 = \text{var}(\theta_{1,1}I_{\mu_{1,1}}) \cup \dots \cup \text{var}(\theta_{m,1}I_{\mu_{m,1}}) \cup \{\tilde{x}, \tilde{y}\}$. By construction of \mathcal{C} , we have $(\mathcal{F}_1, (\tilde{M}), \mathcal{I}_1, \mathcal{V}_1) \in \mathcal{S}$.

So we have $Tr, \bigcup \sigma_{\ell,1} \cup \rho_1 \vdash \bigcup \theta_{\ell,1}\mathcal{F}_{F_\ell, \mu_{\ell,1}, c_{\ell,1}}$. Letting $\sigma_1 = \bigcup \sigma_{\ell,1}$, we have $Tr, \sigma_1 \cup \rho'_1 \vdash \mathcal{F}_1$; for ℓ such that F_ℓ is an injective event, $\mu\mathcal{E}v(\tau'_\ell) = \sigma_1\mathcal{I}_1(\ell) : e_\ell(\dots)$; $\mathcal{V}_1 = \text{Dom}(\sigma_1) \cup \{\tilde{x}, \tilde{y}\}$. Since $\rho'_1, \mu\mathcal{E}v \vdash \text{event}(e(\tilde{M}))@f(\tau_1, \dots, \tau_m, \tilde{x})$, we have $\rho'_1, \tilde{M} \Downarrow \tilde{a}_1$ and $\mu\mathcal{E}v(f(\tau'_1, \dots, \tau'_m, \tilde{a}')) = (\dots) : e(\tilde{a}_1)$ for some \tilde{a}_1 .

Since $f(\tau''_1, \dots, \tau''_m, \tilde{a}'') \neq \perp$, we have similarly $(\mathcal{F}_2, (\tilde{M}), \mathcal{I}_2, \mathcal{V}_2) \in \mathcal{S}$, ρ'_2 , and σ_2 such that $Tr, \sigma_2 \cup \rho'_2 \vdash \mathcal{F}_2$; for ℓ such that F_ℓ is an injective event, $\mu\mathcal{E}v(\tau''_\ell) = \sigma_2\mathcal{I}_2(\ell) : e_\ell(\dots)$; $\mathcal{V}_2 = \text{Dom}(\sigma_2) \cup \{\tilde{x}, \tilde{y}\}$; $\rho'_2, \tilde{M} \Downarrow \tilde{a}_2$ and $\mu\mathcal{E}v(f(\tau''_1, \dots, \tau''_m, \tilde{a}'')) = (\dots) : e(\tilde{a}_2)$ for some \tilde{a}_2 .

Let θ'' be a renaming of variables in \mathcal{V}_2 . Then $Tr, \sigma_2\theta''^{-1} \cup \rho'_2\theta''^{-1} \vdash \theta''\mathcal{F}_2$; for ℓ such that F_ℓ is an injective event, $\mu\mathcal{E}v(\tau''_\ell) = \sigma_2\theta''^{-1}\theta''\mathcal{I}_2(\ell) : e_\ell(\dots)$; $\rho'_2\theta''^{-1}, \theta''\tilde{M} \Downarrow \tilde{a}_2$ and $\mu\mathcal{E}v(f(\tau''_1, \dots, \tau''_m, \tilde{a}'')) = (\dots) : e(\tilde{a}_2)$ for some \tilde{a}_2 .

Then $Tr, \sigma_1 \cup \sigma_2\theta''^{-1} \cup \rho'_1 \cup \rho'_2\theta''^{-1} \vdash \mathcal{F}_1 \cup \theta''\mathcal{F}_2$.

There exists $j \in I$ such that $\tau'_j \neq \tau''_j$, so $\sigma_1\mathcal{I}_1(j) \neq \sigma_2\theta''^{-1}\theta''\mathcal{I}_2(j)$ (distinct events have distinct pairs (program point, replication indices) by Lemma 41), so there exists $j \in \text{Dom}(\mathcal{I}_1) = I$ such that $Tr, \sigma_1 \cup \sigma_2\theta''^{-1} \cup \rho'_1 \cup \rho'_2\theta''^{-1} \vdash \mathcal{I}_1(j) \neq \theta''\mathcal{I}_2(j)$.

Since $f(\tau'_1, \dots, \tau'_m, \tilde{a}') = f(\tau''_1, \dots, \tau''_m, \tilde{a}'')$, $\mu\mathcal{E}v(f(\tau'_1, \dots, \tau'_m, \tilde{a}')) = \mu\mathcal{E}v(f(\tau''_1, \dots, \tau''_m, \tilde{a}''))$, so $\tilde{a}_1 = \tilde{a}_2$, so $Tr, \sigma_1 \cup \sigma_2\theta''^{-1} \cup \rho'_1 \cup \rho'_2\theta''^{-1} \vdash \tilde{M} = \theta''\tilde{M}$.

So $Tr, \sigma_1 \cup \sigma_2\theta''^{-1} \cup \rho'_1 \cup \rho'_2\theta''^{-1} \vdash \mathcal{F}_1 \cup \theta''\mathcal{F}_2 \cup \{\bigvee_{j \in \text{Dom}(\mathcal{I}_1)} \mathcal{I}_1(j) \neq \theta''\mathcal{I}_2(j), \tilde{M} = \theta''\tilde{M}\}$.

Since the trace satisfies $\{\vdash \mathcal{C}\}$, this is a contradiction. Therefore, we conclude that the considered trace satisfies φ . Hence, every full trace of $C[Q_0]$ that satisfies φ_{ni} , $\{\vdash \mathcal{C}\}$, and does not execute a non-unique event of Q_0 also satisfies φ . Therefore, every full trace of $C[Q_0]$ that satisfies $\neg\varphi$ satisfies $\neg(\varphi_{\text{ni}} \wedge \{\vdash \mathcal{C}\} \wedge \neg\text{NonUnique}_{Q_0})$, so every full trace of $C[Q_0]$ that satisfies $\neg\varphi \wedge \neg\text{NonUnique}_{Q_0}$ satisfies $\neg(\varphi_{\text{ni}} \wedge \{\vdash \mathcal{C}\} \wedge \neg\text{NonUnique}_{Q_0}) \wedge \neg\text{NonUnique}_{Q_0} = (\neg\varphi_{\text{ni}} \vee \neg\{\vdash \mathcal{C}\}) \wedge \neg\text{NonUnique}_{Q_0} = (\neg\varphi_{\text{ni}} \wedge \neg\text{NonUnique}_{Q_0}) \vee (\neg\{\vdash \mathcal{C}\} \wedge \neg\text{NonUnique}_{Q_0, D_{\text{false}}})$, so it satisfies

$(\neg\varphi_{\text{ni}} \wedge \neg\text{NonUnique}_{Q_0}) \vee \neg\{\vdash \mathcal{C}\}$. So

$$\begin{aligned}
\text{Adv}_{Q_0}(\mathcal{A}, C, \varphi, D_{\text{false}}) &= \Pr^{\mathcal{A}}[C[Q_0] : \neg\varphi \wedge \neg\text{NonUnique}_{Q_0}] \\
&\leq \Pr^{\mathcal{A}}[C[Q_0] : (\neg\varphi_{\text{ni}} \wedge \neg\text{NonUnique}_{Q_0}) \vee \neg\{\vdash \mathcal{C}\}] \\
&\leq \Pr^{\mathcal{A}}[C[Q_0] : \neg\{\vdash \mathcal{C}\}] + \Pr^{\mathcal{A}}[C[Q_0] : \neg\varphi_{\text{ni}} \wedge \neg\text{NonUnique}_{Q_0}] \\
&\leq \Pr^{\mathcal{A}}[C[Q_0] : \neg\{\vdash \mathcal{C}\}] + \text{Adv}_{Q_0}(\mathcal{A}, C, \varphi_{\text{ni}}, D_{\text{false}}) \\
&\leq \Pr^{\mathcal{A}}[C[Q_0] : \neg\{\vdash \mathcal{C}\}] + p(\mathcal{A}, C) \quad \text{since } \text{Bound}_{Q_0}(V, \varphi_{\text{ni}}, D_{\text{false}}, p) \\
&\leq p'(\mathcal{A}, C).
\end{aligned}$$

Therefore $\text{Bound}_{Q_0}(V, \varphi, D_{\text{false}}, p')$. \square

If the proof that $\text{noninj}(\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi)$ implies $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$ works, we just have to prove $\text{noninj}(\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi)$ and we can apply the **guess** transformation for non-injective correspondences. Otherwise, we simply leave the query $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$ unchanged.

A transformation **guess i && above**, similar to **guess i** , can be used to guess the whole sequence \tilde{i} of replication indices above and including the modified replication, by testing the equality $\tilde{i} = \tilde{i}_{\text{tested}}$ instead of $i = i_{\text{tested}}$.

5.1.18 **guess** $x[c_1, \dots, c_m]$

Like the transformation **guess i** , when **guessRemoveUnique = true** and some (one-session or bit) secrecy queries are present, the transformation **guess** $x[c_1, \dots, c_m]$ first transforms the game G into G_{RU} , by replacing all proved **find[unique]** with **find**. Lemma 58 shows the soundness of this preliminary transformation.

Next, the transformation **guess** $x[c_1, \dots, c_m]$ transforms a game G into a game G' by guessing the value of a variable $x[c_1, \dots, c_m]$: it replaces the processes P under the definition of $x[c_1, \dots, c_m]$ with

$$\text{if } x[c_1, \dots, c_m] = v_{\text{tested}} \text{ then } P \text{ else event_abort bad_guess}$$

and v_{tested} is a constant, which is the guessed value of $x[c_1, \dots, c_m]$. (At each definition of x , CryptoVerif must be able to determine whether it is a definition of $x[c_1, \dots, c_m]$ or not. The variable x must not be defined inside a term. In case **diff_constants = true**, the constant v_{tested} must not be considered different from other constants of the same type.)

In case there is a (one-session or bit) secrecy query, it uses instead

$$\text{let } \text{guess_x_defined} = \text{true in if } x[c_1, \dots, c_m] = v_{\text{tested}} \text{ then } P \text{ else event_abort bad_guess}$$

where **guess_x_defined** is a fresh variable, and we add **guess_x_defined** to the public variables of (one-session or bit) secrecy queries. That gives the adversary knowledge of whether the guessed variable is defined or not. This is useful because the adversary may need to swap its answer differently depending on whether the guessed variable is defined or not, so that the cases in which this variable is not defined always increase the probability of breaking (one-session or bit) secrecy.

When there are only correspondence queries, we can actually execute any code when $x[c_1, \dots, c_m]$ is different from the guessed value v_{tested} . In particular, we can execute P with $x[c_1, \dots, c_m]$ set to v_{tested} , which has the effect of replacing the definition of $x[c_1, \dots, c_m]$ with **let** $x = v_{\text{tested}}$ and removing the test $x[c_1, \dots, c_m] = v_{\text{tested}}$.

To sum up, we also define a transformation **guess** $x[c_1, \dots, c_m]$ **no_test** that can be applied when there are only correspondence queries and when $x[c_1, \dots, c_m]$ is defined only by definitions of the form `let` $x = M$. (This is the most useful case, since the definition of x can then be simplified.) This transformation replaces these definitions `let` $x = M$ with `let` $x = v_{\text{tested}}$ when M is a simple term and with `let` $ignore = M$ in `let` $x = v_{\text{tested}}$ otherwise, where $ignore$ is a fresh variable whose value is not used.

The transformation **guess** $x[c_1, \dots, c_m]$ **no_test** would not be valid in the presence of secrecy queries (at least not with the same probability), because before transformation the value of $x[c_1, \dots, c_m]$ may contain part of the secret variable and this value may leak, while after transformation, that leaks disappears and the variable may be perfectly secret for all values $x[c_1, \dots, c_m] = v_{\text{tested}}$.

Lemma 60 *The transformations **guess** $x[c_1, \dots, c_m]$ and **guess** $x[c_1, \dots, c_m]$ **no_test** require and preserve Properties 1, 2, 3, and 4. They preserve Property 5.*

*Suppose the game G is transformed into G' by the transformation **guess** $x[c_1, \dots, c_m]$ or **guess** $x[c_1, \dots, c_m]$ **no_test**, where x is of type T and $T \neq \emptyset$.*

Let φ be the semantics of a correspondence. Let D be a disjunction of Shoup and non-unique events that does not contain S nor \bar{S} . If $\text{Bound}_{G'}(V, \varphi, D, p)$ and p is independent of the value of v_{tested} , then $\text{Bound}_G(V, \varphi, D, |T|p)$.

*Let sp be $1\text{-ses.secr.}(y)$, $\text{Secrecy}(y)$, or $\text{bit secr.}(x)$. Then G is transformed into G' by the transformation **guess** $x[c_1, \dots, c_m]$. If G' satisfies sp with public variables $V \cup \{\text{guess_x_defined}\}$ ($y \notin V$) up to probability p and p satisfies Property 6, then G satisfies sp with public variables V up to probability $|T| \times p$ (neglecting a small additional runtime of the context). If $\text{Bound}_{G'}(V' \cup \{\text{guess_x_defined}\}, sp, \text{NonUnique}_{G'}, p)$ and p satisfies Property 6, then $\text{Bound}_G(V', sp, D_{\text{false}}, |T|p)$ (neglecting a small additional runtime of the context).*

Proof

Correspondences Let C be an evaluation context acceptable for G with any public variables that does not contain events used by φ or D and \mathcal{A} be an attacker in \mathcal{BB} .

$$\begin{aligned} & \text{Adv}_G(\mathcal{A}, C, \varphi, D) \\ &= \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D}] \\ &= \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D} \wedge x[c_1, \dots, c_m] \text{ not defined}] + \\ & \quad \sum_{v \in T} \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D} \wedge x[c_1, \dots, c_m] = v] \\ &\leq \sum_{v \in T} \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D} \wedge (x[c_1, \dots, c_m] = v \vee x[c_1, \dots, c_m] \text{ not defined})] \end{aligned}$$

Moreover,

$$\begin{aligned} & \Pr^{\mathcal{A}}[C[G'] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G',D}] \geq \\ & \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D} \wedge (x[c_1, \dots, c_m] = v_{\text{tested}} \vee x[c_1, \dots, c_m] \text{ not defined})] \end{aligned}$$

This property holds because G' behaves like G when $x[c_1, \dots, c_m]$ is not defined or $x[c_1, \dots, c_m] = v_{\text{tested}}$, in both transformations **guess** $x[c_1, \dots, c_m]$ and **guess** $x[c_1, \dots, c_m]$ **no_test**. So

$$\begin{aligned} \text{Adv}_G(\mathcal{A}, C, \varphi, D) &\leq \sum_{v \in T} \Pr^{\mathcal{A}}[C[G'] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G',D}] \text{ for } v_{\text{tested}} = v \\ &\leq \sum_{v \in T} \text{Adv}_{G'}(\mathcal{A}, C, \varphi, D) \text{ for } v_{\text{tested}} = v \end{aligned}$$

Since $\text{Bound}_{G'}(V, \varphi, D, p)$, we have $\text{Adv}_{G'}(\mathcal{A}, C, \varphi, D) \leq p(\mathcal{A}, C)$ and p is independent of the value of v_{tested} , so we obtain $\text{Adv}_G(\mathcal{A}, C, \varphi, D) \leq |T| \times p(\mathcal{A}, C)$. Therefore $\text{Bound}_G(V, \varphi, D, |T|p)$.

(One-session or bit) secrecy Let C be an evaluation context acceptable for $C_{sp}[G]$ with public variables V that does not contain S nor $\bar{\mathsf{S}}$ and \mathcal{A} be an attacker in \mathcal{BB} . We have

$$\begin{aligned} \text{Adv}_G^{sp}(\mathcal{A}, C) &= \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}}] \\ &= \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge x[c_1, \dots, c_m] \text{ not defined}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge x[c_1, \dots, c_m] \text{ not defined}] \\ &\quad + \sum_{v=1}^{|T|} \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge x[c_1, \dots, c_m] = v] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge x[c_1, \dots, c_m] = v] \end{aligned}$$

From the adversary C , we define four adversaries C' that call oracle O'_s (O''_s for bit secrecy) with b'' instead of b' , where $b'' = \text{if defined}(\text{guess}_x.\text{defined})$ then $f_1(b')$ else $f_2(b')$ where $f_1(b')$ is either b' or $-b'$, and similarly for f_2 , and consider the adversary $C'_{\max, v_{\text{tested}}}$ among those four that yields the maximum $\text{Adv}_{G'}^{sp}(\mathcal{A}, C')$. Changing b' into $-b'$ swaps the events S and $\bar{\mathsf{S}}$, and therefore swaps their probabilities. Hence, for this adversary $C'_{\max, v_{\text{tested}}}$,

$$\begin{aligned} \text{Adv}_{G'}^{sp}(\mathcal{A}, C'_{\max, v_{\text{tested}}}) &= \\ &|\Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge x[c_1, \dots, c_m] = v_{\text{tested}}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge x[c_1, \dots, c_m] = v_{\text{tested}}]| + \\ &|\Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge x[c_1, \dots, c_m] \text{ not defined}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge x[c_1, \dots, c_m] \text{ not defined}]| \end{aligned}$$

Since G' satisfies sp with public variables $V \cup \{\text{guess}_x.\text{defined}\}$ ($y \notin V$) up to probability p and $C'_{\max, v_{\text{tested}}}$ is an evaluation context acceptable for $C_{sp}[G']$ with public variables $V \cup \{\text{guess}_x.\text{defined}\}$ that does not contain S nor $\bar{\mathsf{S}}$, we have $\text{Adv}_{G'}^{sp}(\mathcal{A}, C'_{\max, v_{\text{tested}}}) \leq p(\mathcal{A}, C'_{\max, v_{\text{tested}}})$. So we have

$$\begin{aligned} \sum_{v_{\text{tested}}=1}^{|T|} p(\mathcal{A}, C'_{\max, v_{\text{tested}}}) &\geq \sum_{v_{\text{tested}}=1}^{|T|} \text{Adv}_{G'}^{sp}(\mathcal{A}, C'_{\max, v_{\text{tested}}}) \\ &\geq \sum_{v_{\text{tested}}=1}^{|T|} \left| \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge x[c_1, \dots, c_m] = v_{\text{tested}}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge x[c_1, \dots, c_m] = v_{\text{tested}}] \right| \\ &\quad + |T| \times \left| \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge x[c_1, \dots, c_m] \text{ not defined}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge x[c_1, \dots, c_m] \text{ not defined}] \right| \\ &\geq \text{Adv}_G^{sp}(\mathcal{A}, C) \end{aligned}$$

Moreover, by Property 6, p is independent of the value of v_{tested} (since the type T is bounded) and $p(\mathcal{A}, C)$ depends only on the runtime of C , the number of calls C makes to the various oracles

(which determine replication bounds), and the length of bitstrings, so we have $p(\mathcal{A}, C'_{\max, v_{\text{tested}}}) = p(\mathcal{A}, C)$. (The additional runtime of the context can be neglected.) So $\text{Adv}_G^{sp}(\mathcal{A}, C) \leq |T|p(\mathcal{A}, C)$. Therefore, G satisfies sp with public variables V up to probability $|T| \times p$. The proof of the second property for (one-session or bit) secrecy proceeds as for **guess** i in Lemma 59. \square

5.1.19 guess_branch

Like the transformation **guess** i , when **guessRemoveUnique** = **true** and some (one-session or bit) secrecy queries are present, the transformation **guess_branch** μ first transforms the game G into G_{RU} , by replacing all proved **find**[**unique**] with **find**. Lemma 58 shows the soundness of this preliminary transformation.

Next, the transformation **guess_branch** μ guesses the branch taken by a branching instruction (**if**, **let**, **find**) at program point μ . The program point μ is designated as explained in Section 5.1.10. The instruction at μ must be executed at most once (either because it is not under replication or because this is proved by CryptoVerif, showing that two executions with distinct replication indices lead to a contradiction: $\mathcal{F}_\mu \cup \mathcal{F}_\mu\{\tilde{i}'/\tilde{i}\} \cup \{\tilde{i}' \neq \tilde{i}\}$ yields a contradiction, where \tilde{i} are the current replication indices at μ and \tilde{i}' are fresh replication indices, using a mode of the equational prover of Section 3.4 that does not allow elimination of collisions, so that this property is proved without probability loss). Suppose this instruction has k branches.

We consider a game G and define transformed games G'_j ($0 \leq j < k$) in which branch j of the instruction at μ is kept and all other branches are replaced with **event_abort** **bad_guess**.

In case there is a (one-session or bit) secrecy query, let **guess_br_defined** = **true** is added before the instruction at μ where **guess_br_defined** is a fresh variable, and we add **guess_br_defined** to the public variables of (one-session or bit) secrecy queries. That gives the adversary knowledge of whether the instruction at μ is executed or not. This is useful because the adversary may need to swap its answer differently depending on whether that instruction is executed or not, so that the cases in which that instruction is not executed always increase the probability of breaking (one-session or bit) secrecy.

When there are only correspondence queries, we can actually execute any code when the taken branch is different from the guessed one. In particular, we can execute the same code as in the tested branch, which has the effect of removing the test at μ when that test is **if**. (The tests **find** and **let** with pattern-matching have additional effects: guaranteeing the definition of variables for **find**; defining variables for **let**. In general, that prevents their removal.)

To sum up, we also define a transformation **guess_branch** μ **no_test** that can be applied when there are only correspondence queries and the instruction at μ is **if** M **then** P_1 **else** P_0 . This transformation defines two transformed games G'_j ($j \in \{0, 1\}$) in which the instruction at μ is replaced with P_j when M is a simple term and with **let** $ignore = M$ in P_j otherwise, where $ignore$ is a fresh variable whose value is not used.

The transformation **guess_branch** μ **no_test** would not be valid in the presence of secrecy queries (at least not with the same probability), because before transformation the test at μ may make the value of M leak, which can reveal for instance one bit of the secret variable, while after transformation, that leaks disappears and the variable may be perfectly secret both when P_0 and when P_1 are executed.

Lemma 61 *The transformations **guess_branch** μ and **guess_branch** μ **no_test** require and preserve Properties 1, 2, 3, and 4. They preserve Property 5.*

*Suppose the game G is transformed into games G'_j ($0 \leq j < k$) by the transformation **guess_branch** μ or **guess_branch** μ **no_test**.*

Let φ be the semantics of a correspondence. Let D be a disjunction of Shoup and non-unique events that does not contain S nor $\bar{\mathsf{S}}$. If for all $0 \leq j < k$, $\text{Bound}_{G'_j}(V, \varphi, D, p_j)$, then $\text{Bound}_G(V, \varphi, D, \sum_{j=0}^{k-1} p_j)$.

Let sp be $\text{1-ses.secr.}(y)$, $\text{Secrecy}(y)$, or $\text{bit.secr.}(x)$. Then G is transformed into G'_j by the transformation **guess_branch** μ . If G'_j satisfies sp with public variables $V \cup \{\text{guess_br_defined}\}$ ($y \notin V$) up to probability p_j for $0 \leq j < k$ and the probabilities p_j satisfy Property 6, then G satisfies sp with public variables V up to probability $\sum_{j=0}^{k-1} p_j$ (neglecting a small additional runtime of the context). If for all $0 \leq j < k$, $\text{Bound}_{G'_j}(V' \cup \{\text{guess_br_defined}\}, sp, \text{NonUnique}_{G'_j}, p_j)$ and the probabilities p_j satisfy Property 6, then $\text{Bound}_G(V', sp, D_{\text{false}}, \sum_{j=0}^{k-1} p_j)$ (neglecting a small additional runtime of the context).

Proof

Correspondences Let C be an evaluation context acceptable for G with any public variables that does not contain events used by φ or D and \mathcal{A} be an attacker in \mathcal{BB} .

$$\begin{aligned}
\text{Adv}_G(\mathcal{A}, C, \varphi, D) &= \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D}] \\
&= \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D} \wedge \mu \text{ not executed}] + \\
&\quad \sum_{j=0}^{k-1} \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D} \wedge \text{branch } j \text{ is taken at } \mu] \\
&\leq \sum_{j=0}^{k-1} \Pr^{\mathcal{A}} \left[C[G] : \begin{array}{l} (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D} \wedge \\ (\mu \text{ not executed} \vee \text{branch } j \text{ is taken at } \mu) \end{array} \right] \\
&\leq \sum_{j=0}^{k-1} \Pr^{\mathcal{A}}[C[G'_j] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G'_j,D}] \tag{33} \\
&\leq \sum_{j=0}^{k-1} \text{Adv}_{G'_j}(\mathcal{A}, C, \varphi, D) \\
&\leq \sum_{j=0}^{k-1} p_j(\mathcal{A}, C) \qquad \text{since } \text{Bound}_{G'_j}(V, \varphi, D, p_j)
\end{aligned}$$

The step (33) is valid because G'_j behaves like G when μ is not executed or branch j is taken at μ , in both transformations **guess_branch** μ and **guess_branch** μ **no_test**. Therefore, we obtain $\text{Bound}_G(V, \varphi, D, \sum_{j=0}^{k-1} p_j)$.

(One-session or bit) secrecy Let C be an evaluation context acceptable for $C_{sp}[G]$ with public variables V that does not contain S nor $\bar{\mathsf{S}}$ and \mathcal{A} be an attacker in \mathcal{BB} . We have

$$\begin{aligned}
&\text{Adv}_G^{sp}(\mathcal{A}, C) \\
&= \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}}] \\
&= \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge \mu \text{ not executed}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge \mu \text{ not executed}] \\
&\quad + \sum_{j=0}^{k-1} \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge \text{branch } j \text{ is taken at } \mu] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge \text{branch } j \text{ is taken at } \mu]
\end{aligned}$$

From the adversary C , we define four adversaries C' that call oracle O'_s (O''_s for bit secrecy) with b'' instead of b' , where $b'' = \text{if defined(guess_br_defined) then } f_1(b') \text{ else } f_2(b')$ where $f_1(b')$ is either b' or $\neg b'$, and similarly for f_2 , and consider the adversary $C'_{\max,j}$ among those four that yields the maximum $\text{Adv}_{G'_j}^{sp}(\mathcal{A}, C')$. Changing b' into $\neg b'$ swaps the events S and $\bar{\mathsf{S}}$, and therefore swaps their probabilities. Hence, for this adversary $C'_{\max,j}$,

$$\begin{aligned} & \text{Adv}_{G'_j}^{sp}(\mathcal{A}, C'_{\max,j}) \\ &= |\Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge \text{branch } j \text{ is taken at } \mu] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge \text{branch } j \text{ is taken at } \mu]| \\ &+ |\Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge \mu \text{ not executed}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge \mu \text{ not executed}]| \end{aligned}$$

So

$$\begin{aligned} & \sum_{j=0}^{k-1} \text{Adv}_{G'_j}^{sp}(\mathcal{A}, C'_{\max,j}) \\ & \geq \sum_{j=0}^{k-1} \left| \begin{array}{c} \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge \text{branch } j \text{ is taken at } \mu] \\ - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge \text{branch } j \text{ is taken at } \mu] \end{array} \right| \\ & \quad + \sum_{j=0}^{k-1} |\Pr^{\mathcal{A}}[C[C_{sp}[G]] : \mathsf{S} \wedge \mu \text{ not executed}] - \Pr^{\mathcal{A}}[C[C_{sp}[G]] : \bar{\mathsf{S}} \wedge \mu \text{ not executed}]| \\ & \geq \text{Adv}_G^{sp}(\mathcal{A}, C) \end{aligned}$$

Since G'_j satisfies sp with public variables $V \cup \{\text{guess_br_defined}\}$ up to probability p_j and $C'_{\max,j}$ is an evaluation context acceptable for $C_{sp}[G'_j]$ with public variables $V \cup \{\text{guess_br_defined}\}$ that does not contain S nor $\bar{\mathsf{S}}$, we have $\text{Adv}_{G'_j}^{sp}(\mathcal{A}, C'_{\max,j}) \leq p_j(\mathcal{A}, C'_{\max,j})$. Moreover, by Property 6, $p_j(C)$ depends only on the runtime of C , the number of calls C makes to the various oracles (which determine replication bounds), and the length of bitstrings, so we have $p_j(C'_{\max,j}) = p_j(C)$. (The additional runtime of the context can be neglected.) So $\text{Adv}_G^{sp}(\mathcal{A}, C) \leq \sum_{j=0}^{k-1} p_j(C)$. Therefore, G satisfies sp with public variables V up to probability $\sum_{j=0}^{k-1} p_j$. The proof of the second property for (one-session or bit) secrecy proceeds as for **guess** i in Lemma 59. \square

5.1.20 global_dep_anal [26]

Extension: A few minor extensions have been implemented with respect to the following description.

The global dependency analysis **global_dep_anal** x tries to find a set of variables S such that only variables in S depend on x . In particular, when the global dependency analysis succeeds, the control flow and the view of the adversary do not depend on x , except in cases of negligible probability.

Let x be a variable defined only by random choices $x \stackrel{R}{\leftarrow} T$ where T is a large type. Let S_{def} be a set of variables defined only by assignments. Let S_{dep} be a set of variables containing x . (Intuitively, S_{dep} will be a superset of variables that depend on x .)

We say that a function $f : T \rightarrow T'$ is *uniform* when each element of T' has at most $|T|/|T'|$ antecedents by f . In particular, this is true in the following two cases:

- f is such that $f(x)$ is uniformly distributed in T' if x is uniformly distributed in T .

- f is the restriction to the image of f' of an inverse of f' , where f' is a poly-injective function. (We consider that $f(x)$ is undefined when x is not in the image of f' . Here, in contrast to the rest of the paper, we allow $f : T \rightarrow T'$ to be defined only on a subset of T .) Precisely, when $x_k \in S_{\text{def}}$ is defined by a pattern-matching $\text{let } f'(x_1, \dots, x_n) = M \text{ in } P \text{ else } P'$, we have $x_k = f'^{-1}(M)$, but furthermore when x_k is defined we know that the value of M is in the image of f' , so we have $x_k = f(M)$ where $f = f'^{-1}|_{\text{im } f'}$.

We say that M characterizes a part of x with $S_{\text{def}}, S_{\text{dep}}$ when for all M_0 obtained from M by substituting variables of S_{def} with their definition (when there is a dependency cycle among variables of S_{def} , we do not substitute a variable inside its definition), $\alpha M_0 = M_0$ implies $f_1(\dots f_k((\alpha x)[\widetilde{M}']))) = f_1(\dots f_k(x[\widetilde{M}])))$ for some uniform functions f_1, \dots, f_k and for some \widetilde{M} and \widetilde{M}' , where α is a renaming of variables of S_{dep} to fresh variables, $x[\widetilde{M}]$ is a subterm of M_0 , $(\alpha x)[\widetilde{M}']$ is a subterm of αM_0 , the variables in S_{dep} do not occur in \widetilde{M} or \widetilde{M}' , T is the type of the result of f_1 (or of x when $k = 0$), and T is a large type. In that case, the value of M uniquely determines the value of $f_1(\dots f_k(x[\widetilde{M}])))$.

We use a simple rewriting prover to determine that. We consider the set of terms $\mathcal{M}_0 = \{\alpha M_0 = M_0\}$, and we rewrite elements of \mathcal{M}_0 using the first kind of user-defined rewrite rules mentioned in Section 3.1 and the rule $\{M_1 \wedge M_2\} \cup \mathcal{M}' \rightarrow \{M_1, M_2\} \cup \mathcal{M}'$.

When \mathcal{M}_0 can be rewritten to a set that contains an equality of the form $f_1(\dots f_k(x[\widetilde{M}']))) = f_1(\dots f_k((\alpha x)[\widetilde{M}'])))$ or $f_1(\dots f_k((\alpha x)[\widetilde{M}']))) = f_1(\dots f_k(x[\widetilde{M}'])))$ for some \widetilde{M} and \widetilde{M}' such that the variables in S_{dep} do not occur in \widetilde{M} or \widetilde{M}' , we have that M characterizes a part of x with $S_{\text{def}}, S_{\text{dep}}$.

We say that M characterizes a part of x when M characterizes a part of x with \emptyset, S' where S' is $\{x\}$ union the set of all variables except those defined by random choices. (We know that variables different from x and defined by random choices do not depend on x , so in the absence of more precise information, we can set $S_{\text{dep}} = S'$.)

We say that $\text{only_dep}(\mathcal{A}, x) = S$ when intuitively, only variables in S depend on x , and the adversary cannot see the value of x . Formally, $\text{only_dep}(\mathcal{A}, x) = S$ when

- $S \cap V = \emptyset$.
- Variables of S do not occur in a term N of a return value $\text{return}(N)$.
- Variables of S except x are defined only by assignments.
- If a variable $y \in S$ occurs in M in $\text{let } z : T = M \text{ in } P$, then $z \in S$.
- Variables in S may occur in defined conditions of find but only at the root of them.
- All terms M_j in processes $\text{find}(\bigoplus_{j=1}^m \widetilde{u}_j[\widetilde{v}] \leq \widetilde{n}_j \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P'$ are combinations by \wedge, \vee , or \neg of terms that either do not contain variables in S or are of the form $M_1 = M_2$ or $M_1 \neq M_2$ where M_1 characterizes a part of x with $S \setminus \{x\}, S$ and no variable of S occurs in M_2 , or M_2 characterizes a part of x with $S \setminus \{x\}, S$ and no variable of S occurs in M_1 .

The last item implies that the result of tests does not depend on the values of variables in S , except in cases of negligible probability. Indeed, the tests $M_1 = M_2$ with M_1 characterizes a part of x with $S \setminus \{x\}, S$ and M_2 does not depend on variables in S are false except in cases of negligible probability, since the value of M_1 uniquely determines the value of $f_1(\dots f_k(x[\widetilde{M}'])))$ and M_2 does not depend on $f_1(\dots f_k(x[\widetilde{M}'])))$, so the equality $M_1 = M_2$ happens for a single value of $f_1(\dots f_k(x[\widetilde{M}'])))$, which yields a negligible probability because f_1, \dots, f_k are uniform, x

is chosen with uniform probability, and the type of the result of f_1 is large. Similarly, the tests $M_1 \neq M_2$ are true except in cases of negligible probability.

In checking the conditions of $\text{only_dep}(\mathcal{A}, x) = S$, we do not consider the parts of the code that are unreachable due to tests whose result is known by the conditions above.

The set S is computed by a fixpoint iteration, starting from $\{x\}$ and adding variables defined by assignments that depend on variables already in S .

If we manage to show that $\text{only_dep}(\mathcal{A}, x) = S$, we transform the game as follows:

- We replace with false terms $M_1 = M_2$ in conditions of `find` where M_1 characterizes a part of x with $S \setminus \{x\}$, S and no variable of S occurs in M_2 , or symmetrically.
- We replace with true terms $M_1 \neq M_2$ in conditions of `find` where M_1 characterizes a part of x with $S \setminus \{x\}$, S and no variable of S occurs in M_2 , or symmetrically.

Lemma 62 *The transformation **global_dep_anal** requires and preserves Properties 1, 2, 3, 4, and 5. If transformation **global_dep_anal** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D, \text{EvUsed}$, where p is an upper bound on the probability that required equalities do not hold.*

5.1.21 simplify [26, 27]

Extension: *Some extensions have been implemented with respect to the following description.*

We use the following transformations in order to simplify games. These transformations exploit the information collected as explained in Section 3.

1. Each term M in the game is replaced with a simplified term M' obtained by reducing M by user-defined rewrite rules knowing \mathcal{F}_{P_M} (see Sections 3.1 and 3.4) and the rewrite rules obtained from \mathcal{F}_{P_M} by the above equational prover where P_M is the smallest process containing M . The replacement is performed only when at least one user-defined rewrite rule has been used, to avoid complicating the game by substituting all variables with their value.
2. When setting **inferUnique** is true, CryptoVerif tries to prove uniqueness of `find[uniquee]`, as in transformation **prove_unique** (Section 5.1.4).
3. If $P = \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P'$, M_j does not contain `event_abort` nor unproved `find[uniquee]`, $u_{jk}[\tilde{i}]$ reduces into M' by user-defined rewrite rules knowing \mathcal{F}_{P_j} (see Sections 3.1 and 3.4) and the rewrite rules obtained from \mathcal{F}_{P_j} , and u_{jk} does not occur in M' , then u_{jk} is removed from the j -th branch of this `find`, i_{jk} is replaced with $M' \{i_{j'k'}/u_{j'k'}, j' \leq m, k' \leq m_j\}$ in $M_{j1}, \dots, M_{jl_j}, M_j$ and P_j is replaced with `let $u_{jk}[\tilde{i}] : [1, n_{jk}] = M'$ in P_j` . (Intuitively, $u_{jk}[\tilde{i}] = M'$, so the value of $u_{jk}[\tilde{i}]$ can be computed by evaluating M' instead of performing an array lookup. We remove $u_{jk}[\tilde{i}]$ from the variables looked up by `find` and replace $u_{jk}[\tilde{i}]$ with its value M' .)
4. Suppose that $P = \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P'$, there exists a term M such that `defined`(M) $\in \mathcal{F}_{P_j}$, $x[N_1, \dots, N_l]$ is a subterm of M , $x \neq u_{jk}$ for all $k \leq m_j$, and none of the following conditions holds: a) P is under a definition of x in Q_0 ; b) Q_0 contains $Q_1 \mid Q_2$ such that a definition of x occurs in Q_1 and P is under Q_2 or a definition of x occurs in Q_2 and P is under Q_1 ; c) Q_0 contains $lp + 1$ replications above a process Q

that contains a definition of x and P , where lp is the length of the longest common prefix between N_1, \dots, N_l and the current replication indices at the definitions of x . Then the j -th branch of the `find` is removed. (In this case, $x[N_1, \dots, N_l]$ cannot be defined at P , so the j -th branch of the `find` cannot be taken.)

5. Suppose that $P = \text{find}[\text{unique?}] (\bigoplus_{j=1}^m u_{j1}[\tilde{i}] = i_{j1} \leq n_{j1}, \dots, u_{jm_j}[\tilde{i}] = i_{jm_j} \leq n_{jm_j} \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P'$, there exist terms M, M' such that $\text{defined}(M) \in \mathcal{F}_{P_j}$, $x[N_1, \dots, N_l]$ is a subterm of M , $\text{defined}(M') \in \mathcal{F}_{P_j}$, $x'[N'_1, \dots, N'_l]$ is a subterm of M' , $N_k = N'_k$ for all $k \leq \min(l, l')$, $x \neq x'$, and x and x' are incompatible, then the j -th branch of the `find` is removed. Two variables x and x' are said to be compatible when either there exists $Q_1 \mid Q_2$ in the game such that x is defined in Q_1 and x' is defined in Q_2 , or there is a definition of x' under a definition of x , or symmetrically.
6. If $P = \text{find}[\text{unique?}] (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P'$ and \mathcal{F}_{P_j} yields a contradiction, then the j -th branch of the `find` is removed if M_j does not contain `event_abort` nor unproved `find[uniquee]`, and P_j is replaced with `yield` if M_j contains `event_abort` or some unproved `find[uniquee]`.
7. If $P = \text{find}[\text{unique?}] \text{ else } P'$, then P is replaced with P' .
8. If $\text{find}[\text{unique?}] (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then } P_j) \text{ else } P'$ and $\mathcal{F}_{P'}$ yields a contradiction, then P' is replaced with `yield`.
9. If $P = \text{find}[\text{unique?}] \tilde{u}[\tilde{i}] = \tilde{i}' \leq \tilde{n} \text{ suchthat defined}(M_1, \dots, M_l) \wedge M \text{ then } P_1 \text{ else } P'$, `[unique?]` is not `[uniquee]` for some non-unique event e that is not proved yet to have negligible probability, $\mathcal{F}_{P'}$ yields a contradiction, M is simple (so M never aborts), and the variables in \tilde{u} are not used outside P and are not in V , then P is replaced with P_1 . (When the `find` defines variables \tilde{u} used elsewhere, we cannot remove it.)
10. If $P = \text{find}[\text{unique?}] (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] \leq \tilde{n}_j \text{ suchthat defined}(M_{j1}, \dots, M_{jl_j}) \wedge M_j \text{ then yield}) \text{ else yield}$, `[unique?]` is not `[uniquee]` for some non-unique event e that is not proved yet to have negligible probability, the terms M_j do not contain `event_abort` nor unproved `find[uniquee]`, and the variables in \tilde{u}_j are not used outside P and are not in V , then P is replaced with `yield`.
11. The `defined` conditions of `find` are updated so that Invariant 2 is satisfied. (When such a `defined` condition guarantees that M is defined, $\text{defined}(M)$ implies $\text{defined}(M')$, and after simplification M' appears in the scope of this condition, then M' has to be added to this condition if it is not already present.)
12. If $P = x \stackrel{R}{\leftarrow} T; P'$ or let $x : T = M$ in P' and x is not used in the game and is not in V , then P is replaced with P' .
13. If one of the `then` branches of a `find[unique]` always succeeds and the conditions of this `find[unique]` do not contain `event_abort` nor unproved `find[uniquee]`, then we keep only that branch.

Indeed, the other branches are never taken: the conditions of this `find[unique]` never abort in traces counted in the probability, and the `find` itself aborts when there are several successful choices.

14. We reorganize a `find[unique]` that occurs in a `then` branch of a `find[unique]`: we transform

$$\text{find[unique]} \left(\bigoplus_{j=1}^k \tilde{u}_j = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat } c_j \text{ then } P_j \right) \text{ else } P$$

where $P_{j_0} = \text{find[unique]} \left(\bigoplus_{j'=1}^{k'} FB_{j'} \right) \text{ else } P''_{j_0}$ into

$$\begin{aligned} & \text{find[unique]} \left(\bigoplus_{j=1, \dots, k; j \neq j_0} \tilde{u}_j = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat } c_j \text{ then } P_j \right) \\ & \oplus \left(\bigoplus_{j'=1}^{k'} \bigoplus_{(\tilde{u}' = \tilde{i}' \leq \tilde{n}' \text{ suchthat } c' \text{ then } P') \in \mathbf{b}(FB_{j'}, L_{j'})} \right. \\ & \quad \left. \tilde{u}_{j_0} = \tilde{i}_{j_0} \leq \tilde{n}_{j_0}, \tilde{u}' = \tilde{i}' \leq \tilde{n}' \text{ suchthat } c_{j_0} \wedge c' \text{ then } P' \right) \\ & \text{else find[unique]} \tilde{u}_{j_0} = \tilde{i}_{j_0} \leq \tilde{n}_{j_0} \text{ suchthat } c_{j_0} \text{ then } P''_{j_0} \text{ else } P \end{aligned}$$

where

- either for all $j \leq k$, c_j does not contain `event_abort` nor unproved `find[unique]e` or c_{j_0} never aborts (this is true in particular when c_{j_0} does not contain `event_abort` nor proved or unproved `find[unique]e`)
- \tilde{i} are the current replication indices at the transformation point
- for all j' , $FB_{j'} = (\tilde{u}'_{j'} = \tilde{i}'_{j'} \leq \tilde{n}'_{j'} \text{ suchthat } c'_{j'} \text{ then } P'_{j'})$, $c'_{j'} = \text{defined}(\tilde{M}_{j'}) \wedge \dots$, $c'_{j'}$ never aborts, and $L_{j'}$ is the list of $x[\tilde{N}]$ subterm of $\tilde{M}_{j'}$ with $x \in \tilde{u}_{j_0}$ and $\tilde{N} \neq \tilde{i}$, ordered by increasing size
- $\mathbf{b}(\tilde{u}' = \tilde{i}' \leq \tilde{n}' \text{ suchthat } c' \text{ then } P', []) = \{\tilde{u}' = \tilde{i}' \leq \tilde{n}' \text{ suchthat } c' \{\tilde{i}_{j_0}/\tilde{u}_{j_0}\} \text{ then } P'\}$
- $\mathbf{b}(FB, x[\tilde{N}] :: L) = \mathbf{b}(FB, L) \cup \{\tilde{u}' = \tilde{i}' \leq \tilde{n}' \text{ suchthat } c' \{i/x[\tilde{N}]\} \wedge \tilde{N} = \tilde{i} \text{ then } P' \mid (\tilde{u}' = \tilde{i}' \leq \tilde{n}' \text{ suchthat } c' \text{ then } P') \in \mathbf{b}(FB, L)\}$ where $i = x\{\tilde{i}_{j_0}/\tilde{u}_{j_0}\}$.

The function \mathbf{b} takes into account that, before the transformation, $\tilde{u}_{j_0}[\tilde{i}]$ is defined when we test `defined`($\tilde{M}_{j'}$) in $FB_{j'}$, while after the transformation, $\tilde{u}_{j_0}[\tilde{i}]$ is not defined yet when we perform this test. Furthermore, the value of $\tilde{u}_{j_0}[\tilde{i}]$ will be \tilde{i}_{j_0} . Therefore, 1) when we access $x[\tilde{i}]$ in $c'_{j'} = \text{defined}(\tilde{M}_{j'}) \wedge \dots$, we replace this access with i , where $i = x\{\tilde{i}_{j_0}/\tilde{u}_{j_0}\}$; this is done in $\mathbf{b}(FB, [])$ by the substitution $\{\tilde{i}_{j_0}/\tilde{u}_{j_0}\}$; and 2) when we access $x[\tilde{N}]$ in $c'_{j'} = \text{defined}(\tilde{M}_{j'}) \wedge \dots$ for \tilde{N} not syntactically equal to \tilde{i} , we need to distinguish two cases: either at runtime $\tilde{N} = \tilde{i}$ and we replace this access with i (second part of the union in $\mathbf{b}(FB, x[\tilde{N}] :: L)$), or at runtime $\tilde{N} \neq \tilde{i}$ and we continue using $x[\tilde{N}]$ (first part of the union in $\mathbf{b}(FB, x[\tilde{N}] :: L)$). The array accesses $x[\tilde{N}]$ in $L_{j'}$ are ordered by increasing size because, in case of nested array accesses, we need to handle the bigger array access first (so it must occur last in the list), because after substitution of the smaller one with i , we would not recognize the bigger one.

This transformation cannot be performed when the outer `find` is not unique because it might change the probability of taking each branch. Moreover, we tried performing such a transformation when the inner `find` is not unique (in this case, after transformation, the outer `find` is not unique), but it had a negative impact in some examples. Furthermore, in the latter case, the transformation can be performed manually by inserting the desired outer `find` and simplifying the game: CryptoVerif will remove the useless branches of `find`.

The conditions c_{j_0} and $c'_{j'}$ are conjunctions of `defined` conditions and a term. In the current implementation, the transformation is not performed when the term in c_{j_0} is false (the branch of `find` will be removed by another transformation), and the branches such that the terms in $c'_{j'}$ are false are first removed. Furthermore, the transformation is performed only when one of the following conditions holds: the terms in c_{j_0} and all $c'_{j'}$ are simple, or

the term in c_{j_0} is true and all $c'_{j'}$ never abort, or the terms in $c'_{j'}$ are all true and c_{j_0} never aborts. With the usual simplification of \wedge true, this guarantees that the transformed game satisfies Property 5. Moreover, this implies the abortion conditions (c_{j_0} and all $c'_{j'}$ never abort).

After this transformation, we advise renaming the variables \tilde{u}_{j_0} to distinct names, since they now have multiple definitions.

15. We reorganize a `find[unique]` that occurs in a condition of a `find`: we transform

$$\text{find}[\text{unique?}] \left(\bigoplus_{j=1}^k \tilde{u}_j = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat } c_j \text{ then } P_j \right) \text{ else } P$$

where

$$c_{j_0} = \text{defined}(\tilde{M}') \wedge \text{find}[\text{unique}] \left(\bigoplus_{j'=1}^{k'} \tilde{u}'_{j'} = \tilde{i}'_{j'} \leq \tilde{n}'_{j'} \text{ suchthat } c'_{j'} \text{ then } M'_{j'} \right) \text{ else false}$$

for all $j' \leq k'$, $M'_{j'}$ never aborts and either for all $j \leq k$, c_j does not contain `event_abort` nor unproved `find[uniquee]` or for all $j' \leq k'$, $c'_{j'}$ never aborts, into

$$\begin{aligned} & \text{find}[\text{unique?}] \left(\bigoplus_{j=1, \dots, k; j \neq j_0} \tilde{u}_j = \tilde{i}_j \leq \tilde{n}_j \text{ suchthat } c_j \text{ then } P_j \right) \\ & \oplus \left(\bigoplus_{j'=1}^{k'} \tilde{u}_{j_0} = \tilde{i}_{j_0} \leq \tilde{n}_{j_0}, \tilde{u}'_{j'} = \tilde{i}'_{j'} \leq \tilde{n}'_{j'} \text{ suchthat} \right. \\ & \quad \left. \text{defined}(\tilde{M}') \wedge c'_{j'} \wedge M'_{j'} \{ \tilde{i}'_{j'} / \tilde{u}'_{j'} \} \text{ then } P_{j_0} \right) \\ & \text{else } P \end{aligned}$$

The indication `[unique?]` corresponds to either `[unique]` or empty. The `find` is marked `[unique]` after transformation if the outer `find` was `[unique]` before transformation.

The variables $\tilde{u}'_{j'}$ are defined inside the condition of a `find` so by Invariant 3, they have no array accesses. The transformation performed by function `b` above is therefore not needed here.

The conditions $c'_{j'}$ are conjunctions of `defined` conditions and a term. The current implementation first removes the branches such that one of the following two conditions holds: the terms in $c'_{j'}$ are false or $M'_{j'}$ is false and $c'_{j'}$ never aborts. Furthermore, the transformation is performed only when for all j' , one of the following conditions holds: the term in $c'_{j'}$ and $M'_{j'}$ are simple, or the term in $c'_{j'}$ is true and $M'_{j'}$ never aborts, or $M'_{j'}$ is true or false and the term in $c'_{j'}$ never aborts. With the usual simplification of \wedge true and \wedge false, this guarantees that the transformed game satisfies Property 5. Moreover, this implies the abortion conditions ($c'_{j'}$ and $M'_{j'}$ never abort).

The simplification is iterated at most `maxIterSimplif` times. The iteration stops earlier in case a fixpoint is reached.

Lemma 63 *The transformation **simplify** requires and preserves Properties 1, 2, 3, and 4. It preserves 5. If transformation **simplify** transforms G into G' , then $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D, \text{EvUsed}$, where p is an upper bound on the probability that required equalities do not hold.*

Extension: *When the game is not expanded, **simplify** is actually a weaker simplification transformation. We plan to improve that.*

5.1.22 all_simplify

all_simplify perform several simplifications on the game, as if

- **simplify**,
- **move all** if **autoMove = true**,
- **remove_assign useless** if **autoRemoveAssignFindCond = false**,
remove_assign findcond if **autoRemoveAssignFindCond = true**,
- **SArename random** if **autoSARename = true**,
- and **merge_branches** if **autoMergeBranches = true**

had been called.

5.1.23 success simplify

The transformation **success simplify** is a combination of **success** (Section 4) and **simplify** (Section 5.1.21), with the following addition. First, in the **success** step, the command **success simplify** collects information that is known to be true when the adversary manages to break at least one of the desired properties. Then, the first iteration of the **simplify** step removes parts of the game that contradict this information and replaces them with **event_abort** **adv_loses**.

In more detail, **success simplify** collects a set \mathcal{L}^- of $(\mathcal{V}, \mathcal{F})$ and a set of formulas \mathcal{L}^+ . If the adversary breaks a desired security property, then either there exists a set of facts \mathcal{F} in \mathcal{L}^- that holds or there exists a formula in \mathcal{L}^+ that holds. The sets of facts \mathcal{L}^- correspond to cases in which the proof of the security property failed; their probability may be high. The associated variables \mathcal{V} are replication indices and non-process variables that occur in \mathcal{F} . The formulas in \mathcal{L}^+ correspond to cases in which the security property was proved (up to a certain probability); these formulas are negations of the formulas that prove the security property in the considered case; the probability that they hold is bounded by the equational prover of CryptoVerif. The contents of \mathcal{L}^+ does not influence the game obtained after the transformation. It is useful to compute the probability difference coming from the transformation. The sets \mathcal{L}^- and \mathcal{L}^+ are computed as follows:

- In case there is an indistinguishability query, or a (one-session or bit) secrecy query on a variable x not defined only by $\stackrel{R}{\leftarrow}$ or by assignments of variables defined by $\stackrel{R}{\leftarrow}$, or a correspondence query $\forall \tilde{x} : \tilde{T}; \psi \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$ with some event e in ψ such that the game contains **event** $e(\tilde{M})$ and some term in \tilde{M} is not simple, no information is collected at all and **simplify** is not performed.
- For each correspondence query **event**(e) \Rightarrow false where e is a non-unique event, for every μ that executes **event**(e), for every c , $(\theta' I_\mu, \theta' \mathcal{F}_{\text{event}(e), \mu, c}^0)$ is added to \mathcal{L}^- , for some θ' renaming of I_μ to fresh replication indices. (Indeed, in order to break **event**(e) \Rightarrow false, event e must be executed, so the facts $\mathcal{F}_{\text{event}(e), \mu, c}^0$ at some execution of event e hold.)
- For each other correspondence query $\forall \tilde{x} : \tilde{T}; F_1 \wedge \dots \wedge F_m \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$, let $\varphi = \llbracket \forall \tilde{x} : \tilde{T}; F_1 \wedge \dots \wedge F_m \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi \rrbracket$ and $\mathcal{S}_0 = \{(\mu_1, c_1, \dots, \mu_m, c_m) \mid \forall j \leq m, \mu_j \text{ executes } F_j \text{ and } c_j \text{ is a case for } \mathcal{F}_{\mu_j, c_j}\}$; by trying to prove the correspondence, we build a subset \mathcal{S}_1 of \mathcal{S}_0 , a family of substitutions θ , and a pseudo-formula \mathcal{C} such that $\text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S}_1)$. (θ and \mathcal{C} are computed incrementally on the successful cases in the proof of the correspondence.)

For all $(\mu_1, c_1, \dots, \mu_m, c_m) \in \mathcal{S}_0 \setminus \mathcal{S}_1$ (**success** fails to prove the correspondence for those cases):

- If the query is $\forall \tilde{x} : \tilde{T}; \text{event}(e(\tilde{N})) \Rightarrow \text{false}$ and the process or term at μ_1 is $\text{event } e(\tilde{M}); \dots$, then $(\tilde{x} \cup \theta' I_{\mu_1}, \theta' \mathcal{F}_{\mu_1, c_1} \cup \{\text{lastdefprogrampoint}(\mu_1, \theta' I_{\mu_1}), \tilde{N} = \theta' \tilde{M}\})$ is added to \mathcal{L}^- , where θ' is a renaming of I_{μ_1} to fresh replication indices. (We can stop the trace just after event e without changing the truth of the query, and that is more precise because we can use the *elsefind* facts at e .)
- Otherwise, $(\tilde{x} \cup \bigcup_{j=1}^m \theta_j I_{\mu_j}, \bigcup_{j=1}^m \theta_j \mathcal{F}_{F_j, \mu_j, c_j})$ is added to \mathcal{L}^- , where for $j \leq m$, θ_j is a renaming of I_{μ_j} to fresh replication indices. (Indeed, in order to break the correspondence $\forall \tilde{x} : \tilde{T}; F_1 \wedge \dots \wedge F_m \Rightarrow \exists \tilde{y} : \tilde{T}'; \phi$, the events F_1, \dots, F_m must be executed, so the facts $\mathcal{F}_{F_j, \mu_j, c_j}$ for $j \leq m$ that hold when F_1, \dots, F_m are executed certainly hold when the correspondence is broken. In principle, we could add $\neg \exists \tilde{y} : \tilde{T}'; \phi$ to the facts $\bigcup_{j=1}^m \theta_j \mathcal{F}_{F_j, \mu_j, c_j}$ added to \mathcal{L}^- . However, we have no way to express universal quantification in general in known facts, so when \tilde{y} is not empty, we could not add $\forall \tilde{y} : \tilde{T}'; \neg \phi$ but would end up adding $\neg \phi$ which in fact means $\exists \tilde{y} : \tilde{T}'; \neg \phi$. That would remain sound assuming the types in \tilde{T}' are not empty, but would be weaker. Moreover, in practice, we end up having to distinguish precisely the case in which $\exists \tilde{y} : \tilde{T}'; \phi$ can be proved from the case in which it cannot, which can typically be done by inserting an appropriate *find*. We generally insert a Shoup event e in the else branch of that *find*, triggered when the correspondence cannot be proved, a case for which we want to bound the probability. After that, it remains to prove $\text{event}(e) \Rightarrow \text{false}$: we apply **success simplify** to that correspondence. Adding $\neg \phi = \text{true}$ would not change anything for that correspondence, and we exploit that the condition of the inserted *find* is false at event e , which gives us more precise information than having added $\neg \phi$ for the initial correspondence.)

Moreover, $\neg \{\text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S}_1)\}$ is added to \mathcal{L}^+ . (**success** proves the correspondence for the cases in \mathcal{S}_1 .)

- For each secrecy, one-session secrecy, or bit secrecy query on a variable x defined only by $\stackrel{R}{\leftarrow}$ or by assignments of variables defined by $\stackrel{R}{\leftarrow}$, let $\mathcal{S}_0 = \{\mu \mid \mu \text{ follows a definition of } x\}$ and $\mathcal{S}_1 = \{\mu \in \mathcal{S}_0 \mid \text{prove}^{1\text{-ses.secr.}(x)}(\mu)\}$.

For each $\mu \in \mathcal{S}_0 \setminus \mathcal{S}_1$ (**success** fails to prove one-session secrecy for those cases), $(\theta I_\mu, \theta \mathcal{F}_\mu)$ is added to \mathcal{L}^- , where θ is a renaming of I_μ to fresh replication indices. (Indeed, if secrecy, one-session secrecy, or bit secrecy of x is broken, a definition of x must have been executed, so the facts \mathcal{F}_μ at that definition hold.)

Moreover, $\neg \{\text{prove}^{1\text{-ses.secr.}(x)}(\mathcal{S}_1)\}$ is added to \mathcal{L}^+ . (**success** proves one-session secrecy for the cases in \mathcal{S}_1 .)

Additionally, if the considered query is a secrecy query, then for each $\mu_1, \mu_2 \in \mathcal{S}_0$, let $z_1[\tilde{M}_1] = \text{defRand}_{\mu_1}(x)$ and $z_2[\tilde{M}_2] = \text{defRand}_{\mu_2}(x)$. If $z_1 \neq z_2$, then the definitions at μ_1 and μ_2 are proved to be independent, and nothing is added to \mathcal{L}^- nor \mathcal{L}^+ . If $z_1 = z_2$, let \tilde{i} be the current replication indices at the definition of x , let θ_1 and θ_2 be two distinct renamings of \tilde{i} to fresh replication indices, let $\tilde{i}_1 = \theta_1 \tilde{i}$ and $\tilde{i}_2 = \theta_2 \tilde{i}$, let $\mathcal{F} = \theta_1 \mathcal{F}_{\mu_1} \cup \theta_2 \mathcal{F}_{\mu_2} \cup \{\theta_1 \tilde{M}_1 = \theta_2 \tilde{M}_2, \tilde{i}_1 \neq \tilde{i}_2\}$. If \mathcal{F} yields a contradiction, then the definitions at μ_1 and μ_2 are proved to be independent up to a small probability, $\exists \tilde{i}_1, \exists \tilde{i}_2, \wedge \mathcal{F}$ is added to \mathcal{L}^+ . Otherwise, $(\tilde{i}_1 \cup \tilde{i}_2, \mathcal{F})$ is added to \mathcal{L}^- .

In the **simplify** step, the set \mathcal{L}^- is used as follows: for each program point μ not in a condition of `find`, if for all $(\mathcal{V}, \mathcal{F}) \in \mathcal{L}^-$, $\mathcal{F}_\mu \cup \mathcal{F}$ yields a contradiction, then the code at μ is replaced with `event_abort adv_loses`. (The reason why \mathcal{V} is needed in the implementation is for the optimization of probabilities of collisions: we determine using which indices in \mathcal{V} we get the smaller bound for the number of collisions.)

The probability that a security property is broken before the transformation and not after is then bounded by the probability that a modified program point μ is reached and the adversary breaks the property. If that breach corresponds to a case in \mathcal{L}^+ , the probability of the breach itself is bounded by construction of \mathcal{L}^+ . If that breach corresponds to a case in \mathcal{L}^- , the probability of the breach and reaching μ is bounded because for all $(\mathcal{V}, \mathcal{F}) \in \mathcal{L}^-$, $\mathcal{F}_\mu \cup \mathcal{F}$ yields a contradiction, which bounds the probability that the facts $\mathcal{F}_\mu \cup \mathcal{F}$ hold for some $(\mathcal{V}, \mathcal{F}) \in \mathcal{L}$, and \mathcal{F}_μ holds when μ is reached while some \mathcal{F} in \mathcal{L} holds when the adversary breaks the property. This is formalized by the following lemma.

Lemma 64 *The transformation **success simplify** requires and preserves Properties 1, 2, 3, 4, and 5.*

If transformation **success simplify** transforms G into G' , the distinguisher D is a disjunction of Shoup and non-unique events, the property sp and the disjuncts in D correspond to active queries, $\mathcal{L}^- = \{(\mathcal{V}_j, \mathcal{F}_j) \mid j \in J\}$, the modified program points are μ_k for $k \in K$, $\mathcal{F}_k^{\text{mod}} = \mathcal{F}_{\mu_k}$, for all $\mathcal{A} \in \mathcal{BB}$ and all evaluation contexts C acceptable for G , $\Pr^{\mathcal{A}}[C[G] \preceq (\bigvee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \wedge \mathcal{F}_j \wedge \wedge \mathcal{F}_k^{\text{mod}}) \vee (\bigvee \mathcal{L}^+)] \leq p(\mathcal{A}, C)$, and $\text{Bound}_{G'}(V, sp, D, p')$, then we have $\text{Bound}_G(V, sp, D, p + p')$.

Unfortunately, we cannot prove $\mathcal{D}, \mathcal{D}_{\text{SNU}} : G, D, \text{EvUsed} \xrightarrow{V}_p G', D, \text{EvUsed} \cup \{\text{adv_loses}\}$ for the transformation **success simplify**, because, in case of (one-session or bit) secrecy of a variable x , the inequality needed for this property may not hold: we need to take into account that, when x is not defined, traces that execute \mathbf{S} and those that execute $\bar{\mathbf{S}}$ compensate in the computation of $\text{Adv}_G(\mathcal{A}, C, sp, D)$ in order to prove the soundness of this transformation. We cannot prove this soundness independently for $\Pr^{\mathcal{A}}[C[G] : \mathbf{S}]$ and for $\Pr^{\mathcal{A}}[C[G] : \bar{\mathbf{S}}]$.

Moreover, in the implementation, for (one-session or bit) secrecy properties, a probability $2p(\mathcal{A}, C)$ is added instead of just $p(\mathcal{A}, C)$ as shown by the lemma above. The factor 2 is difficult to avoid because other simplifications are performed at the same time as described in the transformation **simplify** (Section 5.1.21), and the factor 2 is needed for these transformations.

Proof

Fact 1. Let φ be a correspondence not of the form $\forall \tilde{x} \in \tilde{T}, \text{event}(e(\tilde{N})) \Rightarrow \text{false}$, or the correspondence $\text{event}(e) \Rightarrow \text{false}$ for some Shoup event e . Let C be any evaluation context acceptable for G with public variables V that does not contain events used by φ and \mathcal{A} be an attacker in \mathcal{BB} . Let Tr be any full trace of $C[G]$ for \mathcal{A} that does not execute any non-unique event of G and such that $Tr \vdash \neg\varphi$. Then $Tr \vdash (\bigvee_{j \in J} \exists \mathcal{V}_j, \wedge \mathcal{F}_j) \vee (\bigvee \mathcal{L}^+)$.

Proof of Fact 1. By Lemma 42, for any substitutions $\theta(\mu_1, c_1, \dots, \mu_m, c_m)$ equal to the identity on \tilde{x} , for any pseudo-formula \mathcal{C} ,

$$Tr \vdash \neg\{\{\text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S}_0)\}\}$$

So

$$Tr \vdash (\neg\{\{\text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S}_1)\}\}) \vee \left(\bigvee_{(\mu_1, c_1, \dots, \mu_m, c_m) \in \mathcal{S}_0 \setminus \mathcal{S}_1} \neg\{\{\text{prove}^\varphi(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m)\}\} \right).$$

Moreover, $\neg\{\mathcal{F} \xrightarrow{\theta} \mathcal{I}, \mathcal{V}, \mathcal{C}} \phi\} \Rightarrow \exists \tilde{z} \in \tilde{T}'', \wedge \mathcal{F}$ where $\tilde{z} = \mathcal{V}$ and \tilde{T}'' are the types of these variables, by an easy induction on ϕ , so $\neg\{\text{prove}^\varphi(\mathcal{C}, \theta, \mu_1, c_1, \dots, \mu_m, c_m)\} \Rightarrow \exists \theta_1 I_{\mu_1}, \dots, \exists \theta_m I_{\mu_m}, \exists \tilde{x} \in \tilde{T}, \wedge \theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m}$ where, for $j \leq m$, θ_j is a renaming of I_{μ_j} to fresh replication indices. So

$$\text{Tr} \vdash (\neg\{\text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S}_1)\}) \vee \left(\bigvee_{(\mu_1, c_1, \dots, \mu_m, c_m) \in \mathcal{S}_0 \setminus \mathcal{S}_1} \exists \theta_1 I_{\mu_1}, \dots, \exists \theta_m I_{\mu_m}, \exists \tilde{x} \in \tilde{T}, \wedge \theta_1 \mathcal{F}_{F_1, \mu_1, c_1} \cup \dots \cup \theta_m \mathcal{F}_{F_m, \mu_m, c_m} \right).$$

The formula $\neg\{\text{prove}^\varphi(\mathcal{C}, \theta, \mathcal{S}_1)\}$ is added to \mathcal{L}^+ and $(\tilde{x} \cup \bigcup_{j=1}^m \theta_j I_{\mu_j}, \bigcup_{j=1}^m \theta_j \mathcal{F}_{F_j, \mu_j, c_j})$ is added to \mathcal{L}^- when $(\mu_1, c_1, \dots, \mu_m, c_m) \in \mathcal{S}_0 \setminus \mathcal{S}_1$ (recall that, when e is Shoup event, e is always executed by `event_abort e`) so $\text{Tr} \vdash \left(\bigvee_{j \in J} \exists \mathcal{V}_j, \wedge \mathcal{F}_j \right) \vee (\vee \mathcal{L}^+)$.

Fact 2. Let e be a non-unique event. Let C be any evaluation context acceptable for G with public variables V that does not contain e and \mathcal{A} be an attacker in \mathcal{BB} . Let Tr be any full trace of $C[G]$ for \mathcal{A} such that $\text{Tr} \vdash e$. Then $\text{Tr} \vdash \left(\bigvee_{j \in J} \exists \mathcal{V}_j, \wedge \mathcal{F}_j \right) \vee (\vee \mathcal{L}^+)$.

Proof of Fact 2. By Lemma 38, there exist a program point μ (in G) and a case c such that, for any θ' renaming of I_μ to fresh replication indices, there exists a mapping σ with domain $\theta' I_\mu$ such that $\text{Tr}, \sigma \vdash \theta' \mathcal{F}_{\text{event}(e), \mu, c}^0$. Let $(\mathcal{V}, \mathcal{F}) = (\theta' I_\mu, \theta' \mathcal{F}_{\text{event}(e), \mu, c}^0)$ be the element of \mathcal{L}^- for μ and c , in the treatment of correspondence $\text{event}(e) \Rightarrow \text{false}$. We have $\text{Tr} \vdash \exists \mathcal{V}, \wedge \mathcal{F}$. Therefore, $\text{Tr} \vdash \bigvee_{j \in J} \exists \mathcal{V}_j, \wedge \mathcal{F}_j$.

Fact 3. Let C be any evaluation context acceptable for G with public variables V and \mathcal{A} be an attacker in \mathcal{BB} . Let Tr be any trace of $C[G]$ for \mathcal{A} . If $\text{Tr} \vdash \neg \bigvee_{k \in K} \exists I_{\mu_k}, \wedge \mathcal{F}_k^{\text{mod}}$, then there is no configuration in Tr at a modified program point μ_k .

Proof of Fact 3. By contraposition, if there is a configuration $\text{Conf} = E, \sigma, \mu_k M, \mathcal{T}, \mu \mathcal{E}v$ or $\text{Conf} = E, (\sigma, \mu_k P), \mathcal{Q}, \mathcal{O}r, \mathcal{T}, \mu \mathcal{E}v$ at a modified program point μ_k in trace Tr , then let θ be a renaming of I_{μ_k} to fresh indices and $\rho = \{\theta I_{\mu_k} \mapsto \sigma I_{\mu_k}\}$; by Corollary 31, $\text{Tr}, \rho \vdash \theta \mathcal{F}_{\mu_k}$. So $\text{Tr}, \rho \vdash \theta \mathcal{F}_k^{\text{mod}}$, so $\text{Tr} \vdash \exists I_{\mu_k}, \wedge \mathcal{F}_k^{\text{mod}}$.

We perform the proof for each query separately.

Case 1: sp is some correspondence φ different from $\forall \tilde{x} \in \tilde{T}, \text{event}(e(\tilde{N})) \Rightarrow \text{false}$ (including sp is true). Let C be any evaluation context acceptable for G with public variables V that does not contain events used by φ , D , nor non-unique events of G , and \mathcal{A} be an attacker in \mathcal{BB} .

Consider any full trace Tr of $C[G]$ for \mathcal{A} such that $\text{Tr} \vdash (\neg \varphi \vee D) \wedge \neg \text{NonUnique}_{G, D}$. Let us show that $\text{Tr} \vdash \left(\bigvee_{j \in J} \exists \mathcal{V}_j, \wedge \mathcal{F}_j \right) \vee (\vee \mathcal{L}^+)$.

Case 1.1: Tr does not execute any non-unique event of G . Then $\text{Tr} \vdash \neg \varphi \vee D_s$ where D_s is the disjunction of Shoup events in D .

Case 1.1.1: $\text{Tr} \vdash \neg \varphi$ We conclude by Fact 1.

Case 1.1.2: $\text{Tr} \vdash e$ for some Shoup event e in D_s . We conclude by Fact 1 for the correspondence $\text{event}(e) \Rightarrow \text{false}$.

Case 1.2: Tr executes a non-unique event of G . Then $\text{Tr} \vdash e$ for some non-unique event e in G and in D . We conclude by Fact 2.

If $Tr \vdash \bigvee_{k \in K} \exists I_{\mu_k}, \wedge \mathcal{F}_k^{\text{mod}}$, then $Tr \vdash \left(\bigvee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \wedge \mathcal{F}_j \wedge \wedge \mathcal{F}_k^{\text{mod}} \right) \vee (\bigvee \mathcal{L}^+)$.

Otherwise, by Fact 3, there is no configuration in Tr at a modified program point μ_k , so Tr has a matching trace in $C[G']$ that also satisfies $(\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G',D}$.

We conclude that

$$\begin{aligned}
 \text{Adv}_G(\mathcal{A}, C, \varphi, D) &\leq \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D}] \\
 &\leq \Pr^{\mathcal{A}}[C[G] : \left(\bigvee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \wedge \mathcal{F}_j \wedge \wedge \mathcal{F}_k^{\text{mod}} \right) \vee (\bigvee \mathcal{L}^+)] \\
 &\quad + \Pr^{\mathcal{A}}[C[G'] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G',D}] \\
 &\leq \Pr^{\mathcal{A}}[C[G] \preceq \left(\bigvee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \wedge \mathcal{F}_j \wedge \wedge \mathcal{F}_k^{\text{mod}} \right) \vee (\bigvee \mathcal{L}^+)] \\
 &\quad + \Pr^{\mathcal{A}}[C[G'] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G',D}] \\
 &\hspace{20em} \text{by Lemma 1} \\
 &\leq p(\mathcal{A}, C) + \text{Adv}_{G'}(\mathcal{A}, C, \varphi, D) \\
 &\leq p(\mathcal{A}, C) + p'(\mathcal{A}, C) \hspace{10em} \text{since Bound}_{G'}(V, sp, D, p')
 \end{aligned}$$

Hence, we have $\text{Bound}_G(V, sp, D, p + p')$.

Case 2: sp is some correspondence $\varphi = \forall \tilde{x} \in \tilde{T}, \text{event}(e(\tilde{N})) \Rightarrow \text{false}$. Let C be any evaluation context acceptable for G with public variables V that does not contain events used by φ , D , nor non-unique events of G , and \mathcal{A} be an attacker in \mathcal{BB} .

Consider a full trace Tr of $C[G]$ for \mathcal{A} such that $Tr \vdash (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D}$. Let us show that either some prefix of Tr satisfies $\left(\bigvee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \wedge \mathcal{F}_j \wedge \wedge \mathcal{F}_k^{\text{mod}} \right) \vee (\bigvee \mathcal{L}^+)$ or there is a matching trace in $C[G']$ that satisfies $(\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G',D}$.

We have $Tr \vdash \exists \tilde{x} \in \tilde{T}, \text{event}(e(\tilde{N})) \vee D$.

Case 2.1: Tr executes a Shoup event e' of D . Then $Tr \vdash e'$ and Tr is actually a full trace that does not execute any non-unique event of G . By Fact 1, $Tr \vdash \left(\bigvee_{j \in J} \exists \mathcal{V}_j, \wedge_{F \in \mathcal{F}_j} F \right) \vee (\bigvee \mathcal{L}^+)$.

Case 2.2: Tr executes a non-unique event of G . Then $Tr \vdash e'$ for some non-unique event e' in G and in D . Then Tr is actually a full trace. By Fact 2, $Tr \vdash \left(\bigvee_{j \in J} \exists \mathcal{V}_j, \wedge_{F \in \mathcal{F}_j} F \right) \vee (\bigvee \mathcal{L}^+)$.

In cases 2.1 and 2.2,

- If $Tr \vdash \bigvee_{k \in K} \exists I_{\mu_k}, \wedge \mathcal{F}_k^{\text{mod}}$, then $Tr \vdash \left(\bigvee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \wedge \mathcal{F}_j \wedge \wedge \mathcal{F}_k^{\text{mod}} \right) \vee (\bigvee \mathcal{L}^+)$.
- Otherwise, by Fact 3, there is no configuration in Tr at a modified program point μ_k , so Tr has a matching trace in $C[G']$ that also satisfies $(\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G',D}$.

Case 2.3: Tr does not execute any non-unique event of G nor any Shoup event of D . Then $Tr \vdash \exists \tilde{x} \in \tilde{T}, \text{event}(e(\tilde{N}))$.

Let $\mu\mathcal{E}v$ be the sequence of events and E be the environment in the last configuration of Tr . There is a mapping ρ of the variables \tilde{x} to their values such that $Tr, \rho \vdash e(\tilde{N})$. As in the proof of Lemma 38, there exist \tilde{a} and $\tau \in \mathbb{N}$ such that $\rho, \tilde{N} \Downarrow \tilde{a}$ and $\mu\mathcal{E}v(\tau) = (\mu, \tilde{a}_0) : e(\tilde{a})$ for some μ and \tilde{a}_0 . The rule of the semantics that may have added this element to $\mu\mathcal{E}v$ is (Event), (EventAbort), (CtxEvent), (FindE) or (EventT). (It cannot be (Find3) nor (Get3) because e is not a non-unique event. It cannot be (GetE) because G does not contain `get` by Property 3.)

- Case 2.3.1: In case (Event), we have reductions

$$\begin{aligned} Conf &= E_0, (\sigma_0, {}^\mu\text{event } e(\widetilde{M}); P) :: St, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{p_0}_{t_0} \dots \xrightarrow{p_1}_{t_1} E_1, (\sigma_1, {}^\mu\text{event } e(\widetilde{a}); P) :: St, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_1, \mu\mathcal{E}v_1 \\ &\xrightarrow{1} Conf' = E_1, (\sigma_1, P) :: St, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_1, (\mu\mathcal{E}v_1, (\mu, \text{Im}(\sigma_1))) : e(\widetilde{a}) \end{aligned}$$

where ${}^\mu\text{event } e(\widetilde{M}); P$ is a subprocess of $C[G]$ up to renaming of oracles, by any number of applications of (Ctx) and a final application of (Event). The terms \widetilde{M} are simple terms (when some term in \widetilde{M} is not simple, no information is collected at all and **simplify** is not performed), so in fact

$$\begin{aligned} Conf &= E_0, (\sigma_0, {}^\mu\text{event } e(\widetilde{M}); P) :: St, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{1}^* E_0, (\sigma_0, {}^\mu\text{event } e(\widetilde{a}); P) :: St, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{1} Conf' = E_0, (\sigma_0, P) :: St, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, (\mu\mathcal{E}v_0, (\mu, \text{Im}(\sigma_0))) : e(\widetilde{a}) \end{aligned}$$

- Case 2.3.2: In case (EventAbort), we have a reduction

$$\begin{aligned} Conf &= E, (\sigma_0, {}^\mu P) :: St, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{p}_t Conf' = E, (\sigma_0, \text{abort}) :: St, \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, (\mu\mathcal{E}v_0, (\mu, \text{Im}(\sigma_0))) : e \end{aligned}$$

where ${}^\mu P$ is a subprocess of $C[G]$ up to renaming of oracles, by (EventAbort).

- Case 2.3.3: In case (EventT), we have reductions

$$\begin{aligned} Conf &= E_0, \sigma_0, {}^\mu\text{event } e(\widetilde{M}); N, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{p_0}_{t_0} \dots \xrightarrow{p_1}_{t_1} E_1, \sigma_1, {}^\mu\text{event } e(\widetilde{a}); N, \mathcal{T}_1, \mu\mathcal{E}v_1 \\ &\xrightarrow{1} E_1, \sigma_1, N, \mathcal{T}_1, (\mu\mathcal{E}v_1, (\mu, \text{Im}(\sigma_1))) : e(\widetilde{a}) \end{aligned}$$

where ${}^\mu\text{event } e(\widetilde{M}); N$ is a subterm of $C[G]$, by any number of applications of (CtxT) and a final application of (EventT). The terms \widetilde{M} are simple terms, so in fact

$$\begin{aligned} Conf &= E_0, \sigma_0, {}^\mu\text{event } e(\widetilde{M}); N, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{1}^* E_0, \sigma_0, {}^\mu\text{event } e(\widetilde{a}); N, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{1} E_0, \sigma_0, N, \mathcal{T}_0, (\mu\mathcal{E}v_0, (\mu, \text{Im}(\sigma_0))) : e(\widetilde{a}) \end{aligned}$$

By Invariant 4, μ is not inside a condition of **find** or **get**, so by Lemma 5, $Conf$ is not in the derivation of an hypothesis of a rule for **find** or **get**. The only rule for processes other than those for **find** or **get** that evaluates a non-simple term is (Ctx) and similarly, the only rule for terms other than those for **find** or **get** that evaluates a term is (CtxT), so we have

$$\begin{aligned} &E_0, (\sigma_0, C_0[C_1[\dots C_k[{}^\mu\text{event } e(\widetilde{M}); N] \dots]]) , \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, \mu\mathcal{E}v_0 \\ &\xrightarrow{1}^* Conf' = E_0, (\sigma_0, C_0[C_1[\dots C_k[N] \dots]]) , \mathcal{Q}_0, \mathcal{O}r_0, \mathcal{T}_0, (\mu\mathcal{E}v_0, (\mu, \text{Im}(\sigma_0))) : e(\widetilde{a}) \end{aligned}$$

for some C_0 context defined in Figure 12, $k \in \mathbb{N}$, and C_1, \dots, C_k contexts defined in Figure 7, by k applications of (CtxT) and one application of (Ctx).

- Case 2.3.4: In case (CtxEvent), we have

$$Conf = E_0, \sigma_0, {}^\mu\text{event_abort } e, \mathcal{T}_0, \mu\mathcal{E}v_0 \xrightarrow{p}_t E_0, \sigma_0, \text{event_abort } (\mu, \text{Im}(\sigma_0)) : e, \mathcal{T}_0, \mu\mathcal{E}v_0$$

by (EventAbortT) where ${}^\mu\text{event_abort } e$ is a subterm of $C[G]$, followed by applications (FindTE), (CtxT), (CtxEventT), (FindE), (Ctx), and (CtxEvent). We let $Conf'$ be the last configuration of Tr .

Let θ' be a renaming of I_μ to fresh replication indices. We have

$$\begin{aligned} \text{prove}^\varphi(\mathcal{C}, \theta_0, \mu, c) &= (\theta' \mathcal{F}_{\text{event}(e(\tilde{N})), \mu, c} \text{ yields a contradiction}), \\ \{\{\text{prove}^\varphi(\mathcal{C}, \theta_0, \mu, c)\}\} &= \forall \theta' I_\mu, \forall \tilde{x} \in \tilde{T}, \neg \bigwedge \theta' \mathcal{F}_{\text{event}(e(\tilde{N})), \mu, c}, \\ \mathcal{C} \text{ is a pseudo-formula with all leaves } \perp, &\text{ so } \{\{\vdash \mathcal{C}\}\} = \text{true}. \end{aligned}$$

If $(\mu, c) \in \mathcal{S}_0 \setminus \mathcal{S}_1$, then in cases 2.3.1 and 2.3.3, $(\tilde{x} \cup \theta' I_\mu, \theta' \mathcal{F}_{\mu, c} \cup \{\text{lastdefprogrampoint}(\mu, \theta' I_\mu), \tilde{N} = \theta' \tilde{M}\})$ is added to \mathcal{L}^- and in cases 2.3.2 and 2.3.4, $(\theta' I_\mu, \theta' \mathcal{F}_{\text{event}(e), \mu, c})$ is added to \mathcal{L}^- . Moreover, $\neg\{\{\text{prove}^\varphi(\mathcal{C}, \theta_0, \mathcal{S}_1)\}\} = \bigvee_{(\mu, c) \in \mathcal{S}_1} \exists \theta' I_\mu, \exists \tilde{x} \in \tilde{T}, \bigwedge \theta' \mathcal{F}_{\text{event}(e(\tilde{N})), \mu, c}$ is added to \mathcal{L}^+ .

As in the proof of Lemma 38, we have $\sigma_0 = [I_\mu \mapsto \tilde{a}_0]$. Let $\sigma = \{\theta' I_\mu \mapsto \tilde{a}_0\}$. As in the proof of Lemma 38, we have $Tr, \sigma \cup \rho \vdash \theta' \mathcal{F}_{\text{event}(e(\tilde{N})), \mu, c}$, since Tr does not execute any non-unique event of G .

- In cases 2.3.1 and 2.3.3 when $(\mu, c) \in \mathcal{S}_1$ and in cases 2.3.2 and 2.3.4, we have $Tr \vdash \left(\bigvee_{j \in J} \exists \mathcal{V}_j, \bigwedge \mathcal{F}_j \right) \vee (\bigvee \mathcal{L}^+)$.

– If $Tr \vdash \bigvee_{k \in K} \exists I_{\mu_k}, \bigwedge \mathcal{F}_k^{\text{mod}}$, then $Tr \vdash \left(\bigvee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \bigwedge \mathcal{F}_j \wedge \bigwedge \mathcal{F}_k^{\text{mod}} \right) \vee (\bigvee \mathcal{L}^+)$.

– Otherwise, by Fact 3, there is no configuration in Tr at a modified program point μ_k , so Tr has a matching trace in $C[G']$ that also satisfies $(\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G', D}$.

- In cases 2.3.1 and 2.3.3 when $(\mu, c) \in \mathcal{S}_0 \setminus \mathcal{S}_1$, let Tr' be the prefix of Tr that stops at $Conf'$. We have $Conf \preceq_{Tr} Conf'$, so by Corollary 3, $Tr \preceq Conf', \sigma \vdash \theta' \mathcal{F}_{\mu, c}$, that is, $Tr', \sigma \vdash \theta' \mathcal{F}_{\mu, c}$. We have $E_0, \sigma_0, \tilde{M} \Downarrow \tilde{a}$, so $E_0, \sigma, \theta' \tilde{M} \Downarrow \tilde{a}$. The environment $E_{Tr'}$ extends E_0 , so $E_{Tr'}, \sigma, \theta' \tilde{M} \Downarrow \tilde{a}$, so $Tr', \sigma \cup \rho \vdash \theta' \tilde{M} = \tilde{N}$. Since $E_{Tr \preceq Conf'} = E_{Tr \preceq Conf}$ and $\sigma_{Conf'} = \sigma_{Conf}$, $Tr', \sigma \vdash \text{lastdefprogrampoint}(\mu, \theta' \tilde{i})$. Hence $Tr' \vdash \bigvee_{j \in J} \exists \mathcal{V}_j, \bigwedge \mathcal{F}_j$.

– If $Tr' \vdash \bigvee_{k \in K} \exists I_{\mu_k}, \bigwedge \mathcal{F}_k^{\text{mod}}$, then $Tr' \vdash \left(\bigvee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \bigwedge \mathcal{F}_j \wedge \bigwedge \mathcal{F}_k^{\text{mod}} \right) \vee (\bigvee \mathcal{L}^+)$ and Tr' is a prefix of Tr .

– Otherwise, by Fact 3, there is no configuration in Tr' at a modified program point μ_k , so Tr has a matching trace in $C[G']$ that also satisfies $(\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G', D}$. (Note that Tr may still be modified by **success simplify**: the matching trace in $C[G']$ may execute a modified program point μ_k after $Conf'$, but still the matching trace executes $e(\tilde{M})$, so satisfies $\neg\varphi$, and it does not execute any non-unique event of G' .)

We conclude that

$$\begin{aligned}
\text{Adv}_G(\mathcal{A}, C, \varphi, D) &= \Pr^{\mathcal{A}}[C[G] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G,D}] \\
&\leq \Pr^{\mathcal{A}}[C[G] \preceq \left(\bigvee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \bigwedge \mathcal{F}_j \wedge \bigwedge \mathcal{F}_k^{\text{mod}} \right) \vee \left(\bigvee \mathcal{L}^+ \right)] \\
&\quad + \Pr^{\mathcal{A}}[C[G'] : (\neg\varphi \vee D) \wedge \neg\text{NonUnique}_{G',D}] \\
&\leq p(\mathcal{A}, C) + \text{Adv}_{G'}(\mathcal{A}, C, \varphi, D) \\
&\leq p(\mathcal{A}, C) + p'(\mathcal{A}, C) \qquad \text{since Bound}_{G'}(V, sp, D, p')
\end{aligned}$$

Hence, we have $\text{Bound}_G(V, sp, D, p + p')$.

Case 3: sp is 1-ses.secr. (x) , Secrecy (x) , or bit secr. (x) with $C = C'[C_{sp}[\]]$ and x is defined only by $\stackrel{R}{\leftarrow}$ or by assignments of variables defined by $\stackrel{R}{\leftarrow}$. Let C' be any evaluation context acceptable for $C_{sp}[G]$ with public variables $V \setminus V_{sp}$ that does not contain S, \bar{S} , any event in D , nor any non-unique event of G , and \mathcal{A} be an attacker in \mathcal{BB} .

$$\begin{aligned}
&\text{Adv}_G(\mathcal{A}, C, sp, D) \\
&= \Pr^{\mathcal{A}}[C[G] : S \vee D] - \Pr^{\mathcal{A}}[C[G] : \bar{S} \vee \text{NonUnique}_{G,D}] \\
&= \Pr^{\mathcal{A}}[C[G] : S] + \Pr^{\mathcal{A}}[C[G] : D] - \Pr^{\mathcal{A}}[C[G] : \bar{S}] - \Pr^{\mathcal{A}}[C[G] : \text{NonUnique}_{G,D}] \\
&\qquad \qquad \qquad \text{since these events are mutually exclusive} \\
&= \Pr^{\mathcal{A}}[C[G] : S \wedge \neg sp] + \Pr^{\mathcal{A}}[C[G] : D] \\
&\quad - \Pr^{\mathcal{A}}[C[G] : \bar{S} \wedge \neg sp] - \Pr^{\mathcal{A}}[C[G] : \text{NonUnique}_{G,D}] \qquad \text{by Lemma 36} \\
&= \Pr^{\mathcal{A}}[C[G] : (S \wedge \neg sp) \vee D] - \Pr^{\mathcal{A}}[C[G] : \bar{S} \wedge \neg sp] - \Pr^{\mathcal{A}}[C[G] : \text{NonUnique}_{G,D}]
\end{aligned}$$

We have

- $\Pr^{\mathcal{A}}[C[G'] : \text{NonUnique}_{G',D}] \leq \Pr^{\mathcal{A}}[C[G] : \text{NonUnique}_{G,D}]$,
- $\Pr^{\mathcal{A}}[C[G'] : \bar{S} \wedge \neg sp] \leq \Pr^{\mathcal{A}}[C[G] : \bar{S} \wedge \neg sp]$

since a trace that executes a non-unique event or \bar{S} in G' cannot execute event adv_loses , so it executed without change in G . Let us show that $\Pr^{\mathcal{A}}[C[G] : (S \wedge \neg sp) \vee D] \leq p(\mathcal{A}, C) + \Pr^{\mathcal{A}}[C[G'] : (S \wedge \neg sp) \vee D]$. Let Tr be a full trace of $C[G]$ for \mathcal{A} such that $Tr \vdash (S \wedge \neg sp) \vee D$. Let us show that $Tr \vdash \left(\bigvee_{j \in J} \exists \mathcal{V}_j, \bigwedge \mathcal{F}_j \right) \vee \left(\bigvee \mathcal{L}^+ \right)$.

Case 3.1: Tr executes a non-unique event of G . Then $Tr \vdash e$ for some non-unique event e in G and in D . We conclude by Fact 2.

Case 3.2: $Tr \vdash e$ for some Shoup event e in D . Then Tr does not execute any non-unique event of G . We conclude by Fact 1 for the correspondence $\text{event}(e) \Rightarrow \text{false}$.

Case 3.3: $Tr \vdash S \wedge \neg sp$. Since $Tr \vdash \neg sp$, we have $Tr \vdash \neg\{\text{prove}^{sp}(\text{Tpp}(Tr))\}$, so we are in one of the following two cases:

- There exists $\mu \in \text{Tpp}(Tr)$ such that $Tr \vdash \neg\{\text{prove}^{1\text{-ses.secr.}(x)}(\mu)\}$.
 - If $\mu \in \mathcal{S}_1$, then $Tr \vdash \bigvee \mathcal{L}^+$ since $\neg\{\text{prove}^{1\text{-ses.secr.}(x)}(\mathcal{S}_1)\}$ is added to \mathcal{L}^+ .
 - If $\mu \in \mathcal{S}_0 \setminus \mathcal{S}_1$, then $(\theta I_\mu, \theta \mathcal{F}_\mu)$ is added to \mathcal{L}^- , where θ is a renaming of I_μ to fresh replication indices. Moreover, we have $\neg\{\text{noleak}(z[\widetilde{M}'], \mathcal{I}, \mathcal{F})\} \Rightarrow \exists \mathcal{I}, \bigwedge \mathcal{F}$, so $\neg\{\text{prove}^{1\text{-ses.secr.}(x)}(\mu)\} \Rightarrow \exists \theta I_\mu, \bigwedge \theta \mathcal{F}_\mu$ since $\text{defRand}_\mu(x)$ is defined. Therefore, $Tr \vdash \bigvee_{j \in J} \exists \mathcal{V}_j, \bigwedge \mathcal{F}_j$.

- In case $sp = \text{Secrecy}(x)$, there exist $\mu_1, \mu_2 \in \text{Tpp}(Tr)$ such that $Tr \vdash \neg\{\text{prove}^{\text{distinct}(x)}(\mu_1, \mu_2)\}$. Let $z_1[\widetilde{M}_1] = \text{defRand}_{\mu_1}(x)$, $z_2[\widetilde{M}_2] = \text{defRand}_{\mu_2}(x)$, \widetilde{i} be the current replication indices at the definition of x , θ_1 and θ_2 be two distinct renamings of \widetilde{i} to fresh replication indices, $\widetilde{i}_1 = \theta_1\widetilde{i}$, $\widetilde{i}_2 = \theta_2\widetilde{i}$, and $\mathcal{F} = \theta_1\mathcal{F}_{\mu_1} \cup \theta_2\mathcal{F}_{\mu_2} \cup \{\theta_1\widetilde{M}_1 = \theta_2\widetilde{M}_2, \widetilde{i}_1 \neq \widetilde{i}_2\}$. Then $z_1 = z_2$ and $Tr \vdash \exists\widetilde{i}_1, \exists\widetilde{i}_2, \wedge\mathcal{F}$.
 - If \mathcal{F} yields a contradiction, then $\exists\widetilde{i}_1, \exists\widetilde{i}_2, \wedge\mathcal{F}$ is added to \mathcal{L}^+ , so $Tr \vdash \vee\mathcal{L}^+$.
 - Otherwise, $(\widetilde{i}_1 \cup \widetilde{i}_2, \mathcal{F})$ is added to \mathcal{L}^- , so $Tr \vdash \vee_{j \in J} \exists\mathcal{V}_j, \wedge\mathcal{F}_j$.

If $Tr \vdash \vee_{k \in K} \exists I_{\mu_k}, \wedge \mathcal{F}_k^{\text{mod}}$, then $Tr \vdash \left(\vee_{j \in J, k \in K} \exists I_{\mu_k}, \exists \mathcal{V}_j, \wedge \mathcal{F}_j \wedge \wedge \mathcal{F}_k^{\text{mod}} \right) \vee (\vee \mathcal{L}^+)$.

Otherwise, by Fact 3, there is no configuration in Tr at a modified program point μ_k , so Tr has a matching trace in $C[G']$ that also satisfies $(S \wedge \neg sp) \vee D$.

Therefore, $\Pr^{\mathcal{A}}[C[G] : (S \wedge \neg sp) \vee D] \leq p(\mathcal{A}, C) + \Pr^{\mathcal{A}}[C[G'] : (S \wedge \neg sp) \vee D]$, so

$\text{Adv}_G(\mathcal{A}, C, sp, D)$

$$\leq p(\mathcal{A}, C) + \Pr^{\mathcal{A}}[C[G'] : (S \wedge \neg sp) \vee D] - \Pr^{\mathcal{A}}[C[G'] : \bar{S} \wedge \neg sp] - \Pr^{\mathcal{A}}[C[G'] : \text{NonUnique}_{G', D}]$$

Moreover, by applying on G' the same steps as on G at the beginning of this proof, we have

$$\begin{aligned} & \text{Adv}_{G'}(\mathcal{A}, C, sp, D) \\ &= \Pr^{\mathcal{A}}[C[G'] : (S \wedge \neg sp) \vee D] - \Pr^{\mathcal{A}}[C[G'] : \bar{S} \wedge \neg sp] - \Pr^{\mathcal{A}}[C[G'] : \text{NonUnique}_{G', D}] \end{aligned}$$

so

$$\text{Adv}_G(\mathcal{A}, C, sp, D) \leq p(\mathcal{A}, C) + \text{Adv}_{G'}(\mathcal{A}, C, sp, D) \leq p(\mathcal{A}, C) + p'(\mathcal{A}, C)$$

since $\text{Bound}_{G'}(V, sp, D, p')$. Therefore, we have $\text{Bound}_G(V, sp, D, p + p')$. \square

5.2 crypto: Applying the Security Assumptions on Primitives

Extension: This description is based on the first version of the *crypto* transformation. Some extensions and improvements have been implemented since then.

5.2.1 Basic Transformation

The security of cryptographic primitives is defined using observational equivalences given as axioms. Importantly, this formalism allows us to specify many different primitives in a generic way. Such equivalences are then used by the prover in order to transform a game into another, observationally equivalent game, as explained below in this section.

The primitives are specified using equivalences of the form $G_1 \mid \dots \mid G_m \approx G'_1 \mid \dots \mid G'_m$ where G is defined by the following grammar, with $l \geq 0$ and $m \geq 1$:

| | |
|--|-----------------------------|
| $G ::=$ | group of oracles |
| foreach $i \leq n$ do $y_1 \stackrel{R}{\leftarrow} T_1; \dots; y_l \stackrel{R}{\leftarrow} T_l; (G_1 \mid \dots \mid G_m)$ | replication, random choices |
| $O(x_1 : T_1, \dots, x_l : T_l) := FP$ | oracle |
| $FP ::=$ | oracle body |
| return (M) | return |
| $x[\widetilde{i}] \stackrel{R}{\leftarrow} T; FP$ | random number |

$\text{let } x[\tilde{i}] : T = M \text{ in } FP$ assignment
 $\text{find } (\bigoplus_{j=1}^m \tilde{u}_j[\tilde{i}] \leq \tilde{n}_j \text{ suchthat defined}(M_{j_1}, \dots, M_{j_{l_j}}) \wedge M_j \text{ then } FP_j) \text{ else } FP$ array lookup

Intuitively, $O(x_1 : T_1, \dots, x_l : T_l) := FP$ represents an oracle that takes as argument values x_1, \dots, x_l of types T_1, \dots, T_l respectively and returns a result computed by FP . The observational equivalence $G_1 \mid \dots \mid G_m \approx G'_1 \mid \dots \mid G'_m$ expresses that the adversary has a negligible probability of distinguishing oracles in the left-hand side from corresponding oracles in the right-hand side. Formally, groups of oracles can be encoded as oracle definitions that take as input their arguments and return their result, as shown in Figure 23: $\llbracket G \rrbracket_{\tilde{i}}^{\tilde{j}}$ denotes the translation of the group of oracles G into an oracle definition. The translation of $\text{foreach } i \leq n \text{ do } y_1 \stackrel{R}{\leftarrow} T_1; \dots; y_l \stackrel{R}{\leftarrow} T_l; (G_1 \mid \dots \mid G_m)$ defines oracle $O_{\tilde{j}}$ so that the context can trigger the generation of random numbers y_1, \dots, y_l . (In the left-hand side of equivalences, the result FP of oracles must simply be a return $\text{return } (M)$ where M is a simple term.) The observational equivalence $G_1 \mid \dots \mid G_m \approx G'_1 \mid \dots \mid G'_m$ is then an abbreviation for $\llbracket G_1 \mid \dots \mid G_m \rrbracket \approx \llbracket G'_1 \mid \dots \mid G'_m \rrbracket$.

For example, the security of a MAC (Definition 2) is represented by the equivalence $L \approx R$ where:

$$\begin{aligned}
L &= \text{foreach } i'' \leq n'' \text{ do } k \stackrel{R}{\leftarrow} T_{mk}; (\\
&\quad \text{foreach } i \leq n \text{ do } O_m(x : \text{bitstring}) := \text{return } (\text{mac}(x, k)) \\
&\quad | \text{foreach } i' \leq n' \text{ do } O_v(m : \text{bitstring}, ma : T_{ms}) := \text{return } (\text{verify}(m, k, ma))) \\
R &= \text{foreach } i'' \leq n'' \text{ do } k \stackrel{R}{\leftarrow} T_{mk}; (\\
&\quad \text{foreach } i \leq n \text{ do } O_m(x : \text{bitstring}) := \text{return } (\text{mac}'(x, k)), \\
&\quad | \text{foreach } i' \leq n' \text{ do } O_v(m : \text{bitstring}, ma : T_{ms}) := \\
&\quad \quad \text{find } u \leq n \text{ suchthat defined}(x[u]) \wedge (m = x[u]) \\
&\quad \quad \wedge \text{verify}'(m, k, ma) \text{ then return } (\text{true}) \text{ else return } (\text{false})
\end{aligned} \tag{mac_{eq}}$$

where mac' and verify' are function symbols with the same types as mac and verify respectively. (We use different function symbols on the left- and right-hand sides, just to prevent a repeated application of the transformation induced by this equivalence. Since we add these function symbols, we also add the equation

$$\forall k : T_{mk}, \forall m : \text{bitstring}, \text{verify}'(m, k, \text{mac}'(m, k)) = \text{true} \tag{mac'}$$

which restates (mac) for mac' and verify' .) Intuitively, the equivalence $L \approx R$ leaves MAC computations unchanged (except for the use of primed function symbols in R), and allows one to replace a MAC checking $\text{verify}(m, k, ma)$ with a lookup in the array x of messages whose mac has been computed with key k : if m is found in the array x and $\text{verify}(m, k, ma)$, we return true; otherwise, the check fails (up to negligible probability), so we return false. (If the check succeeded with m not in the array x , the adversary would have forged a MAC.) Obviously, the form of L requires that r is used only to compute or check MACs, for the equivalence to be correct. Formally, the following result shows the correctness of our modeling. It is a fairly easy consequence of Definition 2, and is proved in [26, Appendix E.3].

Proposition 5 *If (mac, verify) is a UF-CMA message authentication code, $\text{mac}' = \text{mac}$, and $\text{verify}' = \text{verify}$, then $\llbracket L \rrbracket \approx \llbracket R \rrbracket$.*

$$\begin{aligned}
\llbracket G_1 \mid \dots \mid G_m \rrbracket &= \llbracket G_1 \rrbracket^1 \mid \dots \mid \llbracket G_m \rrbracket^m \\
\llbracket \text{foreach } i \leq n \text{ do } y_1 \stackrel{R}{\leftarrow} T_1; \dots; y_l \stackrel{R}{\leftarrow} T_l; (G_1 \mid \dots \mid G_m) \rrbracket_{\tilde{i}}^{\tilde{j}} &= \\
\text{foreach } i \leq n \text{ do } O_{\tilde{j}}[\tilde{i}, i]() &:= y_1 \stackrel{R}{\leftarrow} T_1; \dots; y_l \stackrel{R}{\leftarrow} T_l; \text{return } (); (\llbracket G_1 \rrbracket_{\tilde{i}, i}^{\tilde{j}, 1} \mid \dots \mid \llbracket G_m \rrbracket_{\tilde{i}, i}^{\tilde{j}, m}) \\
\llbracket O(x_1 : T_1, \dots, x_l : T_l) := FP \rrbracket_{\tilde{i}}^{\tilde{j}} &= O[\tilde{i}](x_1 : T_1, \dots, x_l : T_l) := FP
\end{aligned}$$

where $O_{\tilde{j}}$ are pairwise distinct oracles, $\tilde{i} = i_1, \dots, i_{l'}$, and $\tilde{j} = j_0, \dots, j_{l'}$.

Figure 23: Translation from groups of oracles to oracle definitions

Similarly, if (enc, dec) is an IND-CPA symmetric encryption scheme (Definition 3), then we have the following equivalence:

$$\begin{aligned}
&\text{foreach } i' \leq n' \text{ do } k \stackrel{R}{\leftarrow} T_k; \text{foreach } i \leq n \text{ do} \\
&\quad O_e(x : \text{bitstring}) := r \stackrel{R}{\leftarrow} T_r; \text{return } (\text{enc}(x, k, r)) \\
&\approx \text{foreach } i' \leq n' \text{ do } k \stackrel{R}{\leftarrow} T_k; \text{foreach } i \leq n \text{ do} \\
&\quad O_e(x : \text{bitstring}) := r \stackrel{R}{\leftarrow} T_r; \text{return } (\text{enc}'(\mathbf{Z}(x), k, r))
\end{aligned} \tag{enc_{eq}}$$

where enc' is a function symbol with the same type as enc , and $\mathbf{Z} : \text{bitstring} \rightarrow \text{bitstring}$ is the function that returns a bitstring of the same length as its argument, consisting only of zeroes. Using equations such as $\forall x : T, \mathbf{Z}(\text{T2b}(x)) = \mathbf{Z}_T$, we can prove that $\mathbf{Z}(\text{T2b}(x))$ does not depend on x when x is of a fixed-length type and $\text{T2b} : T \rightarrow \text{bitstring}$ is the natural injection. The equivalences that formalize the security assumptions on primitives are designed and proved correct by hand from security assumptions in a more standard form, as in the MAC example. Importantly, these manual proofs are done only once for each primitive, and the obtained equivalence can be reused for proving many different protocols automatically.

We use such equivalences $L \approx R$ in order to transform a process Q_0 observationally equivalent to $C[\llbracket L \rrbracket]$ into a process Q'_0 observationally equivalent to $C[\llbracket R \rrbracket]$, for some evaluation context C . In order to check that $Q_0 \approx^V C[\llbracket L \rrbracket]$, the prover uses sufficient conditions, which essentially guarantee that all uses of certain secret variables of Q_0 , in a set S , can be implemented by calling oracles of L . Let \mathcal{M} be a set of occurrences of terms, corresponding to uses of variables of S . Informally, the prover shows the following properties.

- For each $M \in \mathcal{M}$, there exist a term N_M , which is the return value of an oracle of L , and a substitution σ_M such that $M = \sigma_M N_M$. (Precisely, σ_M applies to the abbreviated form of N_M in which we write x instead of $x[\tilde{i}]$.) Intuitively, the evaluation of M in Q_0 will correspond to a call to the oracle returning N_M in $C[\llbracket L \rrbracket]$.
- The variables of S do not occur in V , are bound by random choices in Q_0 , and occur only in terms $M = \sigma_M N_M \in \mathcal{M}$ in Q_0 , at occurrences that are images by σ_M of variables bound by random choices in L . (To be precise, the variables of S are also allowed to occur at the root of defined conditions; in that case, their value does not matter, just the fact that they are defined.)
- Let \tilde{i} and \tilde{i}' be the sequences of current replication indices at N_M in L and at M in Q_0 , respectively. The prover shows that there exists a function mapIdx_M that maps the array indices at M in Q_0 to the array indices at N_M in L : the evaluation of M when

$\tilde{i}' = \tilde{a}$ will correspond in $C[[L]]$ to the evaluation of N_M when $\tilde{i} = \text{mapIdx}_M(\tilde{a})$. Thus, σ_M and mapIdx_M induce a correspondence between terms and variables of Q_0 and variables of L : for all $M \in \mathcal{M}$, for all $x[\tilde{i}'']$ that occur in N_M , $(\sigma_M x)\{\tilde{a}/\tilde{i}'\}$ corresponds to $x[\tilde{i}'']\{\text{mapIdx}_M(\tilde{a})/\tilde{i}\}$, that is, $(\sigma_M x)\{\tilde{a}/\tilde{i}'\}$ in a trace of Q_0 has the same value as $x[\tilde{i}'']\{\text{mapIdx}_M(\tilde{a})/\tilde{i}\}$ in the corresponding trace of $C[[L]]$ (\tilde{i}'' is a prefix of \tilde{i}). We detail below conditions that this correspondence has to satisfy.

For example, consider a process Q_0 that contains $M_1 = \text{enc}(M'_1, x_k, x'_r[i_1])$ under a replication `foreach $i_1 \leq n_1$ do` and $M_2 = \text{enc}(M'_2, x_k, x''_r[i_2])$ under a replication `foreach $i_2 \leq n_2$ do`, where x_k, x'_r, x''_r are bound by random choices. Let $S = \{x_k, x'_r, x''_r\}$, $\mathcal{M} = \{M_1, M_2\}$, and $N_{M_1} = N_{M_2} = \text{enc}(x[i', i], k[i'], r[i', i])$. The functions mapIdx_{M_1} and mapIdx_{M_2} are defined by

$$\begin{aligned} \text{mapIdx}_{M_1}(a_1) &= (1, a_1) \text{ for } a_1 \in [1, n_1] \\ \text{mapIdx}_{M_2}(a_2) &= (1, a_2 + n_1) \text{ for } a_2 \in [1, n_2] \end{aligned}$$

Then $M'_1\{a_1/i_1\}$ corresponds to $x[1, a_1]$, x_k to $k[1]$, $x'_r[a_1]$ to $r[1, a_1]$, $M'_2\{a_2/i_2\}$ to $x[1, a_2 + n_1]$, and $x''_r[a_2]$ to $r[1, a_2 + n_1]$. The functions mapIdx_{M_1} and mapIdx_{M_2} are such that $x'_r[a_1]$ and $x''_r[a_2]$ never correspond to the same cell of r ; indeed, $x'_r[a_1]$ and $x''_r[a_2]$ are independent random numbers in Q_0 , so their images in $C[[L]]$ must also be independent random numbers.

The above correspondence must satisfy the following soundness conditions:

- when x is an oracle argument in L , the term that corresponds to $x[\tilde{a}']$ must have the same type as $x[\tilde{a}']$, and when two terms correspond to the same $x[\tilde{a}']$, they must evaluate to the same value;
- when x is bound by $x \stackrel{R}{\leftarrow} T$ in L , the term that corresponds to $x[\tilde{a}']$ must evaluate to $z[\tilde{a}'']$ where $z \in S$ and z is bound by $z \stackrel{R}{\leftarrow} T$ in Q_0 , and the relation that associates $z[\tilde{a}'']$ to $x[\tilde{a}']$ is an injective function (so that independent random numbers in L correspond to independent random numbers in Q_0).

It is easy to check that, in the previous example, these conditions are satisfied.

The transformation of Q_0 into Q'_0 consists in two steps. First, we replace the random choices that define variables of S with random choices that define fresh variables corresponding to variables bound by $\stackrel{R}{\leftarrow}$ in R . The correspondence between variables of Q_0 and variables of $C[[L]]$ is extended to include these fresh variables. Second, we reorganize Q_0 so that each evaluation of a term $M \in \mathcal{M}$ first stores the values of the arguments x_1, \dots, x_m of the oracle $O(x_1 : T_1, \dots, x_m : T_m) := C[\text{return}(N_M)]$ in fresh variables, then computes N_M and stores its result in a fresh variable, and uses this variable instead of M ; then we simply replace the computation of N_M with the corresponding oracle call of R , taking into account the correspondence of variables.

Next, we formalize this transformation. We require the following conditions for the equivalences $L \approx R$ that model cryptographic primitives:

- H0. $[[L]]$ and $[[R]]$ satisfy Invariants 1, 2, and 8. The oracle names in L (resp. R) are pairwise distinct. Furthermore, the return value of each oracle in R has the same type as the return value of the corresponding oracle of L .
- H1. In L , the oracles bodies FP are simply returns of terms M ; all their array accesses use the current replication indices. (Allowing `let` or `find` in L is difficult, because we need to recognize the terms M in a context and in a possibly syntactically modified form.)

- H2. L and R have the same structure: same replications, same number of oracles, same oracle names, same number of arguments with the same types for each oracle.
- H3. The variables y_j defined by $\stackrel{R}{\leftarrow}$ and x_j defined by oracle arguments in L and R are distinct from other variables defined in R .
- H4. Under `foreach $i \leq n$ do` with no random choice in L , one can have only a single oracle $O(x_1 : T_1, \dots, x_l : T_l) := FP$. (One can transform `foreach $i \leq n$ do ($O_1(\widetilde{x}_1 : \widetilde{T}_1) := FP_1 \mid \dots \mid O_m(\widetilde{x}_m : \widetilde{T}_m) := FP_m$)` `foreach $i_1 \leq n_1$ do ...` `foreach $i_{m'} \leq n_{m'}$ do ...`) into `(foreach $i \leq n$ do $O_1(\widetilde{x}_1 : \widetilde{T}_1) := FP_1 \mid \dots \mid O_m(\widetilde{x}_m : \widetilde{T}_m) := FP_m$)` `foreach $i_1 \leq n'_1$ do ...` `foreach $i_{m'} \leq n'_{m'}$ do ...`) in order to eliminate situations that do not satisfy this requirement.)
- H5. Replications in L (resp. R) must have pairwise distinct bounds n . (This strengthens the typing: the typing then guarantees that, if several variables are accessed with the same oracle indices, then these variables are defined under the same replication.)
- H6. For all random choices $y \stackrel{R}{\leftarrow} T$ that occur above a term M in L , y occurs in M . (This guarantees that, in Hypothesis H'3.1 below, $z_{jk}[M_{j1}, \dots, M_{jq_j}]$ is defined for all $j \leq l$ and $k \leq m_j$. With Hypothesis H4, this guarantees that index_j is well-defined in Hypothesis H'3.1.3 below.)
- H7. Finds in R are of the form

$$\text{find } (\bigoplus_{j=1}^m \widetilde{u}_j \leq \widetilde{n}_j \text{ suchthat defined}(z_{j1}[\widetilde{u}_{j1}], \dots, z_{jl_j}[\widetilde{u}_{jl_j}]) \wedge M_j \text{ then } FP_j) \text{ else } FP'$$

where the following conditions are satisfied:

- For all $1 \leq k \leq l_j$, \widetilde{u}_{jk} is the concatenation of a prefix of the current replication indices (the same prefix for all k) and a non-empty prefix of \widetilde{u}_j .
- When \widetilde{u}_j is non-empty, at least one \widetilde{u}_{jk} for $1 \leq k \leq l_j$ is the concatenation of a prefix of the current replication indices with the whole sequence \widetilde{u}_j .
- When $l_j \neq 0$, there exists $k \in \{1, \dots, l_j\}$ such that for all $k' \neq k$, $z_{jk'}$ is defined syntactically above all definitions of z_{jk} and $\widetilde{u}_{jk'}$ is a prefix of \widetilde{u}_{jk} . (This implies that the same find cannot access variables defined in different oracles under the same replication in R .)
- Finally, variables z_{jk} are not defined by a find in R . (Otherwise, the transformation would be considerably more complicated.)

Such equivalences $L \approx R$ are used by the prover by replacing a process Q_0 observationally equivalent to $C[[L]]$ with a process Q'_0 observationally equivalent to $C[[R]]$, for some evaluation context C . We now give sufficient conditions for a process to be equivalent to $C[[L]]$. These conditions essentially guarantee that all uses of certain secret variables of Q_0 , in a set S , can be implemented by calling oracles of L . These conditions are explained in more detail below.

We first define the function `extract` used in order to extract information from the left- or right-hand sides of the equivalence.

$$\begin{aligned} \text{extract}(O(x_1 : T_1, \dots, x_l : T_l) := FP, ()) = \\ O(x_1 : T_1, \dots, x_l : T_l) := FP \end{aligned}$$

$$\begin{aligned} & \text{extract}(\text{foreach } i \leq n \text{ do } y_1 \stackrel{R}{\leftarrow} T_1; \dots; y_l \stackrel{R}{\leftarrow} T_l; (G_1 \mid \dots \mid G_m), (j_1, \dots, j_k)) = \\ & \quad (y_1 : T_1, \dots, y_l : T_l), \text{extract}(G_{j_1}, (j_2, \dots, j_k)) \\ & \text{extract}((G_1 \mid \dots \mid G_m), (j_0, \dots, j_k)) = \\ & \quad \text{extract}(G_{j_0}, (j_1, \dots, j_k)) \end{aligned}$$

We rename the variables of Q_0 such that variables of L and R do not occur in Q_0 . Assume that there exist a set of variables S and a set \mathcal{M} of occurrences of terms in Q_0 such that:

H'1. $S \cap V = \emptyset$.

H'2. No term in \mathcal{M} occurs in the condition part of a find ($\text{defined}(M_1, \dots, M_l) \wedge M$).

H'3. For each $M \in \mathcal{M}$, there exist a sequence $BL(M) = (j_0, \dots, j_l)$ such that $\text{extract}(L, BL(M)) = (y_{11} : T_{11}, \dots, y_{1m_1} : T_{1m_1}), \dots, (y_{l1} : T_{l1}, \dots, y_{lm_l} : T_{lm_l}), O(x_1 : T_1, \dots, x_m : T_m) := \text{return}(N)$ and a substitution σ such that $M = \sigma N$ (σ applies to the abbreviated form of N in which we write x instead of $x[\tilde{i}]$) and the following conditions hold:

H'3.1. For all $j \leq l$ and $k \leq m_j$, σy_{jk} is a variable access $z_{jk}[M_{j_1}, \dots, M_{j_{q_j}}]$, with $z_{jk} \in S$. We define $z_{jk} = \text{varImL}(y_{jk}, M)$.

H'3.1.1. All definitions of z_{jk} in Q_0 are of the form $z_{jk}[\dots] \stackrel{R}{\leftarrow} T_{jk}$, and for all $k \leq m_j$, they occur under the same replications (but they may occur under different replications for different values of j).

H'3.1.2. When $j \neq j'$ or $k \neq k'$, $z_{jk} \neq z_{j'k'}$.

H'3.1.3. The sequence of array indices $M_{j_1}, \dots, M_{j_{q_j}}$ is the same for all $k \leq m_j$ (but may depend on j). We denote by $\text{index}_j(M)$ a substitution that maps the current replication indices at the definition of z_{jk} to $M_{j_1}, \dots, M_{j_{q_j}}$ respectively. If $m_l = 0$, $\text{index}_l(M)$ is not set by the previous definition, so we set $\text{index}_l(M)$ to map the current replication indices at M to themselves. For each $j < l$, there exists a substitution $\rho_j(M)$ such that $\text{index}_j(M) = \text{index}_{j+1}(M) \circ \rho_j(M)$ and the image of $\rho_j(M)$ does not contain the current replication indices at M . We denote by $\text{im index}_j(M)$ the sequence image by $\text{index}_j(M)$ of the sequence of current replication indices at the definition of z_{jk} (so, $\text{im index}_j(M) = (M_{j_1}, \dots, M_{j_{q_j}})$). We define $\text{im } \rho_j(M)$ similarly.

H'3.2. For all $j \leq m$, σx_j is a term of type T_j .

H'3.3. All occurrences in Q_0 of a variable in S are either as z_{jk} above or at the root of an argument of a **defined** test in a **find** process.

To make it precise which term M each element refers to, we add M as a subscript, writing $y_{jk,M}$ for y_{jk} , $z_{jk,M}$ for z_{jk} , $T_{jk,M}$ for T_{jk} , $x_{j,M}$ for x_j , $T_{j,M}$ for T_j , N_M for N , and σ_M for σ . We also define $\text{nNew}_{j,M} = m_j$, $\text{nNewSeq}_M = l$, and $\text{nInput}_M = m$.

H'4. We say that two terms $M, M' \in \mathcal{M}$ share the first l' sequences of random variables when $y_{jk,M} = y_{jk,M'}$ and $z_{jk,M} = z_{jk,M'}$ for all $j \leq l'$ and $k \leq \text{nNew}_{j,M} = \text{nNew}_{j,M'} \neq 0$. Let l' be the greatest integer such that M and M' share the first l' sequences of random variables. Then we require:

H'4.1. The sets of variables $\{z_{jk,M} \mid j > l' \text{ and } k \leq \text{nNew}_{j,M}\}$ and $\{z_{jk,M'} \mid j > l' \text{ and } k \leq \text{nNew}_{j,M'}\}$ must be disjoint.

H'4.2. $\rho_j(M) = \rho_j(M')$ for all $j < l'$.

H'4.3. If $l' = \text{nNewSeq}_M$ and $N_M = N_{M'}$, then there exists M_0 such that we have $M = (\text{index}_{l'}(M))M_0$, $M' = (\text{index}_{l'}(M'))M_0$, and M_0 does not contain the current replication indices at M or M' .

When these conditions are satisfied, there exists a context C such that $Q_0 \approx_0^V C[[L]]$.

Terms in \mathcal{M} must not occur in conditions of `find` (Hypothesis H'2) because such terms may refer to variables defined by `find`, and by the transformation, these variables might be moved outside their scope, thus violating Invariant 2.

In Hypothesis H'3, the sequence $BL(M)$ indicates which branch of L corresponds to the term M .

Hypothesis H'3.2 checks that the values received as arguments in L are of the proper type. Hypothesis H'3.1.1 checks that variables $z_{jk,M}$ that correspond to variables defined by $\stackrel{R}{\leftarrow}$ in L are of the proper type. The variables y_{jk} defined by $\stackrel{R}{\leftarrow}$ in L are used only in terms N in L . Correspondingly, Hypothesis H'3.3 checks that the corresponding variables $z_{jk,M} \in S$ are not used elsewhere in Q_0 and Hypothesis H'1 checks that they cannot be used directly by the context.

In L , for distinct j, k , the variables y_{jk} correspond to independent random numbers. Correspondingly, Hypothesis H'3.1.2 requires that the variables $z_{jk,M}$ are created by different random choices for distinct j, k . In L , the variables y_{jk} are accessed with the same indices for any k (but a fixed j). Correspondingly, Hypothesis H'3.1.3 requires that the variables $z_{jk,M}$ are accessed with the same indices in $\text{index}_j(M)$ for any k . When instances of N and N' both refer to y_{jk} with the same indices, then they also refer to $y_{j'k'}$ with the same indices when $j' \leq j$. Correspondingly, if M and M' refer to the same z_{jk} , by Hypothesis H'4.1, they also refer to the same $z_{j'k'}$ for $j' \leq j$. Moreover, if $\text{index}_j(M)$ and $\text{index}_j(M')$ evaluate to the same bitstrings, then $\text{index}_{j'}(M)$ and $\text{index}_{j'}(M')$ also evaluate to the same bitstrings, since $\text{index}_{j'}(M) = \text{index}_j(M) \circ \rho_{j-1}(M) \circ \dots \circ \rho_{j'}(M)$ by Hypothesis H'3.1.3 and $\rho_k(M) = \rho_k(M')$ for $k < j$ by Hypothesis H'4.2. These conditions guarantee that we can establish a correspondence from the array cells of variables of S in Q_0 to the array cells of variables defined by $\stackrel{R}{\leftarrow}$ in L , and that this correspondence is an injective function, as required in Section 5.2.

Finally, a term N in L is evaluated at most once for each value of the indices of y_{l_1}, \dots, y_{lm_l} , so N is computed for a single value of the arguments x_1, \dots, x_m . Correspondingly, by Hypothesis H'4.3, when M and M' share the $l = \text{nNewSeq}_M$ sequences of random variables and $\text{index}_l(M)$ and $\text{index}_l(M')$ evaluate to the same bitstring, then M and M' evaluate to the same bitstring.

We compute the possible values of the sets S and \mathcal{M} by fixpoint iteration. We start with $\mathcal{M} = \emptyset$ and S containing a single variable of Q_0 bound by a random choice. (We try all possible variables.) When a term M of Q_0 contains a variable in S , we try to find an oracle in L that corresponds to M , and if we succeed, we add M to \mathcal{M} , and add to S variables in M that correspond to variables bound by random choices in L . (If we fail, the transformation is not possible.) We continue until a fixpoint is reached, in which case all occurrences of variables of S are in terms of \mathcal{M} .

We now describe how we construct a process Q'_0 such that $Q'_0 \approx_0^V C[[R]]$.

1. We first move random choices in the right-hand side of the equivalence, so that they occur above oracle definitions instead of inside oracle bodies. As explained below, this is necessary for the correctness of the subsequent transformation of Q_0 , when random choices appear in the corresponding part of the left-hand side. More precisely, we transform the right-hand side of the equivalence, R , as follows: for each j_1, \dots, j_l , if $\text{extract}(L, (j_1, \dots, j_l)) = (y_{11} : T_{11}, \dots, y_{1m_1} : T_{1m_1}), \dots, (y_{l1} : T_{l1}, \dots, y_{lm_l} : T_{lm_l}), O(x_1 : T_1, \dots, x_m : T_m) := \text{return}(N)$ with $m_l \neq 0$ and $\text{extract}(R, (j_1, \dots, j_l)) = (y'_{11} : T'_{11}, \dots, y'_{1m'_1} : T'_{1m'_1}), \dots, (y'_{l1} : T'_{l1}, \dots,$

$y'_{lm'_l} : T'_{lm'_l}), O(x_1 : T_1, \dots, x_m : T_m) := FP$, for each $z \stackrel{R}{\leftarrow} T$ in FP ,

- we add $z : T$ in the sequence of random variables $y'_{l_1} : T'_{l_1}, \dots, y'_{lm'_l} : T'_{lm'_l}$;
- if z does not occur in defined conditions of find in R , we remove $z \stackrel{R}{\leftarrow} T$ from FP ;
- otherwise, we replace $z \stackrel{R}{\leftarrow} T$ with $\text{let } z' : T = \text{cst}$ for some constant cst and add $z'[\widetilde{M}]$ to each defined condition of R that contains $z[\widetilde{M}]$.

This transformation is needed, because in the right-hand side, a new random number must be chosen exactly for each different call to the oracle $O(x_1 : T_1, \dots, x_m : T_m) := FP$. This would not be guaranteed without that transformation, because when the left-hand side N is evaluated at several occurrences with the same random numbers $y_{l_1} : T_{l_1}, \dots, y_{lm_l} : T_{lm_l}$ ($m_l \neq 0$), these occurrences all correspond to a single call to $O(x_1 : T_1, \dots, x_m : T_m) := \text{return}(N)$, so a single call to $O(x_1 : T_1, \dots, x_m : T_m) := FP$, but we create a copy of FP for each occurrence. After the transformation, FP contains no choice of random numbers, so we can evaluate it several times without changing the result. When $m_l = 0$, evaluations of N at several occurrences can correspond to different calls to $O(x_1 : T_1, \dots, x_m : T_m) := \text{return}(N)$, so the transformation is not necessary.

2. Next, we create fresh variables corresponding to variables of the right-hand side of the equivalence. For each $M \in \mathcal{M}$, let $\text{extract}(R, BL(M)) = (y'_{11,M} : T'_{11,M}, \dots, y'_{1m'_1,M} : T'_{1m'_1,M}), \dots, (y'_{l1,M} : T'_{l1,M}, \dots, y'_{lm'_l,M} : T'_{lm'_l,M}), O(x_{1,M} : T_{1,M}, \dots, x_{m,M} : T_{m,M}) := FP_M$ with $l = \text{nNewSeq}_M$, $m = \text{nInput}_M$ and we define $\text{nNew}'_{j,M} = m'_j$. We create fresh variables $z'_{jk,M} = \text{varImR}(y'_{jk,M}, M)$ for each $j \leq \text{nNewSeq}_M$, $k \leq \text{nNew}'_{j,M}$, and $M \in \mathcal{M}$, such that if M and M' share the first l' sequences of random variables, then $z'_{jk,M} = z'_{jk,M'}$ for $j \leq l'$ and $k \leq \text{nNew}'_{j,M}$. All variables $z'_{jk,M}$ are otherwise pairwise distinct.

We also create a fresh variable $\text{varImR}(x_{j,M}, M)$ for each $j \leq \text{nInput}_M$ and each $M \in \mathcal{M}$, and a fresh variable $\text{varImR}(z, M)$ for each variable z defined by let or $\stackrel{R}{\leftarrow}$ in FP_M and each $M \in \mathcal{M}$.

3. We update the defined conditions of finds, in order to preserve Invariant 2. More precisely, if a defined condition of a find contains $z_{j1,M}[M_1, \dots, M_{l'}]$ for some M , we add $\text{defined}(z'_{jk',M}[M_1, \dots, M_{l'}])$ for all $k' \leq \text{nNew}'_{j,M}$ to this condition. (So that accesses to $z'_{jk',M}[M_1, \dots, M_{l'}]$ created when transforming term M satisfy Invariant 2, since accesses to $z_{j1,M}[M_1, \dots, M_{l'}]$ occur in M and satisfy Invariant 2.)
4. We update random choices corresponding to random choices of the left-hand side of the equivalence: we either remove them or replace them with random choices corresponding to the right-hand side of the equivalence. More precisely, when $x \in S$ occurs at the root of a term M_k in a condition $\text{defined}(M_1, \dots, M_l)$, we replace its definition $x \stackrel{R}{\leftarrow} T; Q$ with $\text{let } x : T = \text{cst}$ in Q for some constant cst ; when it does not occur in defined tests, we remove its definition. If $x = z_{j1,M}$ for some M , we add $z'_{jk,M} \stackrel{R}{\leftarrow} T'_{jk,M}$ for each $k \leq \text{nNew}'_{j,M}$ where $x \stackrel{R}{\leftarrow} T$ was.
5. Finally, we transform the terms $M \in \mathcal{M}$ corresponding to oracles of the left-hand side of the equivalence into their corresponding oracle in the right-hand side. For each term $M \in \mathcal{M}$, let $l = \text{nNewSeq}_M$. We replace M with $(z'_{lk,M} \stackrel{R}{\leftarrow} T'_{lk,M})_{k \leq \text{nNew}'_{l,M}} M'$ if $\text{nNew}_{l,M} = 0$ and $\text{nNew}'_{l,M} > 0$, and with P'_M otherwise, where

- $M' = (\text{let } \text{varImR}(x_{k,M}, M) : T_{k,M} = \sigma_M x_{k,M} \text{ in})_{k \leq \text{nInput}_M} \text{transf}_{\phi_0}(FP_M)$.
- ϕ_0 is defined as follows:

$$\begin{aligned}\phi_0(x_{j,M}[i_1, \dots, i_l]) &= \text{varImR}(x_{j,M}, M)[i'_1, \dots, i'_l] \\ \phi_0(z[i_1, \dots, i_l]) &= \text{varImR}(z, M)[i'_1, \dots, i'_l] \\ \phi_0(y'_{jk,M}[i_1, \dots, i_j]) &= \text{varImR}(y'_{jk,M}, M)[\text{im index}_j(M)]\end{aligned}$$

where i_1, \dots, i_l are the current replication indices at the definition of $x_{j,M}$ in R , i'_1, \dots, i'_l are the current replication indices at M in Q_0 , and z is a variable defined by `let` or $\stackrel{R}{\leftarrow}$ in FP_M .

- A function ϕ from array accesses to array accesses is extended to terms as a substitution, by $\phi(f(M_1, \dots, M_m)) = f(\phi(M_1), \dots, \phi(M_m))$.
- $\text{transf}_{\phi}(FP)$ is defined recursively as follows:

$$\begin{aligned}\text{transf}_{\phi}(\text{return } (M')) &= \phi(M') \\ \text{transf}_{\phi}(z \stackrel{R}{\leftarrow} T; FP') &= \text{varImR}(z, M) \stackrel{R}{\leftarrow} T; \text{transf}_{\phi}(FP') \\ \text{transf}_{\phi}(\text{let } z : T = M' \text{ in } FP') &= \text{let } \text{varImR}(z, M) : T = \phi(M') \text{ in } \text{transf}_{\phi}(FP') \\ \text{transf}_{\phi}(\text{find}(\bigoplus_{j=1}^m FB_j) \text{ else } FP') &= \text{find}(\bigoplus_{j=1}^m \text{transf}_{\phi}(FB_j)) \text{ else } \text{transf}_{\phi}(FP')\end{aligned}$$

and for `find` branches FB , $\text{transf}_{\phi}(FB)$ is defined as follows:

$$\begin{aligned}\text{transf}_{\phi}(\text{suchthat } M' \text{ then } FP') &= \text{suchthat } \phi(M') \text{ then } \text{transf}_{\phi}(FP') \\ \text{transf}_{\phi}(\tilde{u} \leq \tilde{n} \text{ suchthat defined}(z_k[M_{k1}, \dots, M_{kl'_k}]_{1 \leq k \leq l}) \wedge M_1 \text{ then } FP') &= \\ \bigoplus_{M' \in \mathcal{M}'} \tilde{u}' \leq \tilde{n}' \text{ suchthat defined}(\phi'(z_k[M_{k1}, \dots, M_{kl'_k}]_{1 \leq k \leq l}) \wedge & \\ \text{im index}_{j_1}(M')\{\tilde{u}'/i'\}) = \text{im index}_{j_1}(M) \wedge \phi'(M_1) \text{ then } \text{transf}_{\phi'}(FP') &\end{aligned}$$

where $l \neq 0$; j_1 is the length of the prefix of the current replication indices that occurs in $M_{k1}, \dots, M_{kl'_k}$ (by Hypothesis H7); \mathcal{M}' is the set of $M' \in \mathcal{M}$ such that $\text{varImR}(z_k, M')$ is defined for $k \leq l$ and M' and M share the first j_1 sequences of random variables; \tilde{i}' is the sequence of current replication indices at M' ; \tilde{u}' is a sequence formed with a fresh variable for each variable in \tilde{i}' ; \tilde{n}' is the sequence of bounds of replications above M' ; ϕ' is an extension of ϕ with $\phi'(z_k[M_{k1}, \dots, M_{kl'_k}]) = \text{varImR}(z_k, M')[\text{im index}_j(M')\{\tilde{u}'/i'\}]$ if $z_k = y'_{jk',M'}$ for some k' , and $\phi'(z_k[M_{k1}, \dots, M_{kl'_k}]) = \text{varImR}(z_k, M')[\tilde{u}']$ if z_k is defined by `let` or by an oracle input. Optimizations for the definition of $\text{transf}_{\phi}(FB)$ are presented in Section 5.2.2.

The two essential parts of the transformation are the last two ones, numbered 4 and 5. In step 4, we add the random choices to create random variables that correspond to random variables of R . We create the variables $z'_{jk,M}$ at the place where $z_{j1,M}$ was created in the initial game (We could have chosen $z_{jk',M}$ for any k'), or when there is no $z_{j1,M}$, we have $j = \text{nNewSeq}_M$ and we create $z'_{jk,M}$ just before evaluating M . In step 5, we transform the term M itself into the corresponding oracle process of R , FP_M . The only delicate part for evaluating FP_M is the case of `find`: instead of looking up arrays of R , we look up the corresponding arrays of Q'_0 given by the mapping ϕ .

This transformation may replace a simple term M with a term that is not simple, since the oracle bodies in R may contain random number generations, assignments and `find`. Therefore, it may break Property 5. Hence, by default, CryptoVerif calls **expand** (Section 5.1.3) after the **crypto** transformation, in order to recover Property 5.

5.2.2 Optimizations for $\text{transf}_\phi(FB)$

We can apply two optimizations to the definition of $\text{transf}_\phi(FB)$:

- When $\text{im index}_{j_1}(M')$ is a prefix of \tilde{v}' , $\text{im index}_{j_1}(M')\{\tilde{u}'/\tilde{v}'\}$ is a prefix of \tilde{u}' , so the equality $\text{im index}_{j_1}(M')\{\tilde{u}'/\tilde{v}'\} = \text{im index}_{j_1}(M)$ defines the value of a prefix of \tilde{u}' . We simply substitute the fixed elements of \tilde{u}' with their value, and remove them from the sequence of variables to be looked up by `find`.
- When all variables z_k are $y_{jk',M'}$ for some j, k' , and M' , with $\max j = j_0$, we use the following definition instead:

$$\begin{aligned} & \text{transf}_\phi(\tilde{i} \leq \tilde{n} \text{ suchthat defined}(z_k[M_{k_1}, \dots, M_{k_{l'_k}}]_{1 \leq k \leq l}) \wedge M_1 \text{ then } FP') = \\ & \bigoplus_{M' \in \mathcal{M}'} \tilde{u}' \leq \tilde{n}' \text{ suchthat defined}(\phi'(z_k[M_{k_1}, \dots, M_{k_{l'_k}}]_{1 \leq k \leq l}) \wedge \\ & \text{im}(\rho_{j_0-1}(M') \circ \dots \circ \rho_{j_1}(M'))\{\tilde{u}'/\tilde{v}'\} = \text{im index}_{j_1}(M) \wedge \phi'(M_1) \text{ then } \text{transf}_{\phi'}(FP')) \end{aligned}$$

where j_1 is the length of the prefix of the current replication indices that occurs in $M_{k_1}, \dots, M_{k_{l'_k}}$ (by Hypothesis H7); \mathcal{M}' is the set of $M' \in \mathcal{M}$ such that $\text{varImR}(z_k, M')$ is defined for $k \leq l$ and M' and M share the j_1 first sequences of random variables; \tilde{v}' is the sequence of current replication indices at the definition of $z_{j_0k, M'}$; \tilde{u}' is a sequence formed with a fresh variable for each variable in \tilde{v}' ; \tilde{n}' is the sequence of bounds of replications above the definition of $z_{j_0k, M'}$; ϕ' is an extension of ϕ with $\phi'(z_k[M_{k_1}, \dots, M_{k_{l'_k}}]) = \text{varImR}(z_k, M')[\text{im}(\rho_{j_0-1}(M') \circ \dots \circ \rho_j(M'))\{\tilde{u}'/\tilde{v}'\}]$ if $z_k = y'_{jk, M'}$.

The composition $\rho_{j_0-1}(M') \circ \dots \circ \rho_j(M')$ computes the indices of $z'_{jk', M'}$ for any k' from the indices of $z'_{j_0k'', M'}$ for any k'' .

When several terms $M' \in \mathcal{M}$ share the first j_0 sequences of random variables, they generate the same ϕ' , so only one `find` branch needs to be added for all of them, which can reduce considerably the number of `find` branches to add.

An optimization similar to the first one above also applies to this case, when $\text{im}(\rho_{j_0-1}(M') \circ \dots \circ \rho_{j_1}(M'))$ is a prefix of \tilde{v}' .

5.2.3 Guiding the Application of Equivalences

We introduce a small extension to the equivalences $G_1 \mid \dots \mid G_m \approx G'_1 \mid \dots \mid G'_m$ described in Section 5.2. These equivalences become $G_1 \text{ mode}_1 \mid \dots \mid G_m \text{ mode}_m \approx G'_1 \mid \dots \mid G'_m$, where mode_j is either empty or $[all]$. The mode $[all]$ is an indication for the prover, to guide the application of the equivalence without changing its semantics. When $\text{mode}_j = [all]$, \mathcal{M} must contain all occurrences in the initial game Q of the root function symbols of terms M inside G_j . When mode_j is empty, at least one variable defined by $\stackrel{R}{\leftarrow}$ in G_j must correspond to a variable in S .

The following hypotheses guarantee the good usage of modes:

- H8. At most one mode_j can be empty. (Otherwise, when several sets of random variables can be chosen for each G_j , there are many possible combinations for applying the transformation.)
- H9. If G_j is of the form `foreach $i \leq n$ do $O(x_1 : T_1, \dots, x_l : T_l) := FP$` without any random choice, then $\text{mode}_j = [all]$. (A random choice is needed in the definition of empty mode.)

An equivalence can be declared *[manual]*. In this case, this equivalence is not applied by the automatic proof strategy. It can be used only by manual proof indications. This is useful, for instance, if applying the equivalence would yield an infinite loop.

Each oracle in the left-hand side can also be labeled with an integer $[n]$, which represents a priority: CryptoVerif preferably uses oracles labeled with a lower integer. (The absence of label corresponds to the label $[0]$.) Each oracle in the left-hand side can finally be labeled with *[useful_change]*. This indication is also used for the proof strategy: if at least one *[useful_change]* indication is present, CryptoVerif applies the transformation defined by the equivalence only when at least one *[useful_change]* oracle is called in the game.

5.2.4 Relaxing Hypothesis H6

Hypothesis H6 requires that for all random choices $y \stackrel{R}{\leftarrow} T$ that occur above a term N in the left-hand side of an equivalence, y occurs in N . We can relax this hypothesis, by allowing that some random variables y do not occur in N , provided that the missing variables can be determined using Hypothesis H'4.1: when some term M shares some variable y in the l' -th sequence of random variables with some other term M' , we know that it must also share with M' all random variables in sequences above and including the l' -th sequence; so, knowing the random variables associated to M' , we can determine some of those associated to M . The transformation simply fails when the algorithm described above cannot fully determine the random variables associated to some term M .

With this extension, we additionally need to make sure that, when an expression of the right-hand side uses a variable y defined by a random choice, this variable will be defined in the transformed game before the expression is evaluated. To do that, we establish a correspondence between the random choices before the transformation and the random choices after, such that, as far as possible, if a variable is used in a transformed expression, the corresponding variable before transformation is also used in the initial expression. For variables that appear in a transformed expression but are not used in the initial expression, we check when performing the game transformation that they will correctly be defined. If they are not correctly defined, the transformation fails.

5.2.5 Relaxing Hypothesis H'2

Hypothesis H'2 requires that no term N transformed by the equivalence occurs in the condition part of a find (defined(M_1, \dots, M_l) \wedge M). We can relax this hypothesis by allowing N to occur in M (but not in the defined test), provided

- the variables \tilde{u} bound by this find do not occur in the following terms in the transformed expression of N :
 - N' in processes of the form `let $x : T = N'$ in ...`;
 - N'_{jk} and N'_j in processes of the form `find ($\bigoplus_{j=1}^m u_{j1}[\tilde{u}] \leq n_{j1}, \dots, u_{jm_j}[\tilde{u}] \leq n_{jm_j}$ suchthat defined($N'_{j1}, \dots, N'_{jl_j}$) \wedge N'_j then ...) else ...`

(If the variables \tilde{u} bound by find occurred in such terms, the transformation would move them outside the scope of their definition.)

- if N depends on the variables \tilde{u} bound by this find, then the corresponding left-hand side of the equivalence N' is under a replication `foreach $i \leq n$ do` with no random choice in L . (Otherwise, L contains $y_1 \stackrel{R}{\leftarrow} T_1; \dots; y_l \stackrel{R}{\leftarrow} T_l; (\dots \mid O(x_1 : T'_1, \dots, x_{l'} : T'_{l'}) := \text{return } (N') \mid$

...) and L allows a single evaluation of N' for each execution of the random choices $y_1 \stackrel{R}{\leftarrow} T_1; \dots; y_l \stackrel{R}{\leftarrow} T_l$, while the term N is evaluated several times, once for each value of \tilde{u} , so the transformation is impossible. In other words, the variables \tilde{u} must be considered as replication indices in this transformation.)

5.2.6 Soundness and Example

The following proposition shows the soundness of the transformation.

Proposition 6 *Let Q_0 be a process that satisfies Invariants 1, 2, and 8 and Q'_0 the process obtained from Q_0 by the above transformation. Then Q'_0 satisfies Invariants 1, 2, and 8 and, if $\llbracket L \rrbracket \approx \llbracket R \rrbracket$ for all polynomials $\text{maxlen}_\eta(c_{j_0, \dots, j_l})$ and n where n is any replication bound of L or R , then $Q_0 \approx^V Q'_0$.*

Example 10 In order to treat Example 1, the prover is given as input the indication that T_{mk} , T_k , and T_r are fixed-length types; the type declarations for the functions $\text{mac}, \text{mac}' : \text{bitstring} \times T_{mk} \rightarrow T_{ms}$, $\text{verify}, \text{verify}' : \text{bitstring} \times T_{mk} \times T_{ms} \rightarrow \text{bool}$, $\text{enc}, \text{enc}' : \text{bitstring} \times T_k \times T_r \rightarrow T_e$, $\text{dec} : T_e \times T_k \rightarrow \text{bitstring}_\perp$, $\text{k2b} : T_k \rightarrow \text{bitstring}$, $\text{i}_\perp : \text{bitstring} \rightarrow \text{bitstring}_\perp$, $\text{Z} : \text{bitstring} \rightarrow \text{bitstring}$, and the constant $\text{Z}_k : \text{bitstring}$; the equations (mac) , (mac') , (enc) , and $\forall x : T_k, \text{Z}(\text{k2b}(x)) = \text{Z}_k$ (which expresses that all keys have the same length); the indication that k2b and i_\perp are poly-injective (which generates the equations (k2b) and similar equations for i_\perp); equivalences $L \approx R$ for MAC (mac_{eq}) and encryption (enc_{eq}); and the process Q_0 of Example 1.

The process Q_0 of Example 1 can be transformed using the security of the MAC. Let $S = \{x_{mk}\}$, $M_1 = \text{mac}(x_m[i], x_{mk})$, $M_2 = \text{verify}(x'_m[i'], x_{mk}, x_{ma}[i'])$, and $\mathcal{M} = \{M_1, M_2\}$. We have $N_{M_1} = \text{mac}(x[i''], i, k[i''])$, $N_{M_2} = \text{verify}(m[i''], i', k[i''], ma[i''], i')$, $\text{mapIdx}_{M_1}(a_1) = (1, a_1)$, and $\text{mapIdx}_{M_2}(a_2) = (1, a_2)$, so $x_m[a_1]$ corresponds to $x[1, a_1]$, x_{mk} to $k[1]$, $x'_m[a_2]$ to $m[1, a_2]$, and $x_{ma}[a_2]$ to $ma[1, a_2]$.

After transformation, we get the following process Q'_0 :

$$\begin{aligned} Q'_0 &= \text{Start}() := x_k \stackrel{R}{\leftarrow} T_k; x_{mk} \stackrel{R}{\leftarrow} T_{mk}; \text{return } (); (Q'_A \mid Q'_B) \\ Q'_A &= \text{foreach } i \leq n \text{ do } O_A[i]() := x'_k \stackrel{R}{\leftarrow} T_k; x_r \stackrel{R}{\leftarrow} T_r; \\ &\quad \text{let } x_m : \text{bitstring} = \text{enc}(\text{k2b}(x'_k), x_k, x_r) \text{ in} \\ &\quad \text{return } (x_m, \text{mac}'(x_m, x_{mk})) \\ Q'_B &= \text{foreach } i' \leq n \text{ do } O_B[i'](x'_m, x_{ma}) := \\ &\quad \text{find } u \leq n \text{ suchthat defined}(x_m[u]) \wedge x'_m = x_m[u] \wedge \text{verify}'(x'_m, x_{mk}, x_{ma}) \text{ then} \\ &\quad (\\ &\quad \quad \text{if true then let } \text{i}_\perp(\text{k2b}(x''_k)) = \text{dec}(x'_m, x_k) \text{ in return } () \\ &\quad) \\ &\quad \text{else} \\ &\quad (\\ &\quad \quad \text{if false then let } \text{i}_\perp(\text{k2b}(x''_k)) = \text{dec}(x'_m, x_k) \text{ in return } () \\ &\quad) \end{aligned}$$

The initial definition of x_{mr} is removed and replaced with a new definition, which we still call x_{mr} . The term $\text{mac}(x_m, x_{mk})$ is replaced with $\text{mac}'(x_m, x_{mk})$. The term $\text{verify}(x'_m, x_{mk}, x_{ma})$ becomes $\text{find } u \leq n \text{ suchthat defined}(x_m[u]) \wedge x'_m = x_m[u] \wedge \text{verify}'(x'_m, x_{mk}, x_{ma}) \text{ then true else}$

false, which yields Q'_B after the call to **expand**, which transforms the **find** term into a process. The process looks up the message x'_m in the array x_m , which contains the messages whose MAC has been computed with key x_{mk} . If the MAC of x'_m has never been computed, the check always fails (it returns false) by the definition of security of the MAC. Otherwise, it returns true when $\text{verify}'(x'_m, x_{mk}, x_{ma})$.

After applying **simplify**, Q'_A is unchanged and Q'_B becomes

$$Q''_B = \text{foreach } i' \leq n \text{ do } O_B[i'](x'_m, x_{ma}) := \\ \text{find } u \leq n \text{ suchthat defined}(x_m[u], x'_k[u]) \wedge x'_m = x_m[u] \wedge \text{verify}'(x'_m, x_{mk}, x_{ma}) \text{ then} \\ \text{let } x''_k : T_k = x'_k[u] \text{ in return } ()$$

First, the tests if true then ... and if false then ... are simplified. The term $\text{dec}(x'_m, x_k)$ is simplified knowing $x'_m = x_m[u]$ by the **find** condition, $x_m[u] = \text{enc}(\text{k2b}(x'_k[u]), x_k, x_r[u])$ by the assignment that defines x_m , and $\text{dec}(\text{enc}(m, k, r), k) = i_\perp(m)$ by (enc). So we have $\text{dec}(x'_m, x_k) = i_\perp(\text{k2b}(x'_k[u]))$. By injectivity of i_\perp and k2b , the assignment to x''_k simply becomes $x''_k = x'_k[u]$, using the equations $\forall x : \text{bitstring}, i_\perp^{-1}(i_\perp(x)) = x$ and $\forall x : T_k, \text{k2b}^{-1}(\text{k2b}(x)) = x$.

Then we apply the security of encryption: $\text{enc}(\text{k2b}(x'_k), x_k, x_r)$ becomes $\text{enc}'(\text{Z}(\text{k2b}(x'_k)), x_k, x_r)$. After **simplify**, it becomes $\text{enc}'(\text{Z}_k, x_k, x_r)$, using $\forall x : T_k, \text{Z}(\text{k2b}(x)) = \text{Z}_k$ (which expresses that all keys have the same length).

So we obtain the following game:

$$Q''_0 = \text{Start}() := x_k \stackrel{R}{\leftarrow} T_k; x_{mk} \stackrel{R}{\leftarrow} T_{mk}; \text{return } (); (Q''_A \mid Q''_B) \\ Q''_A = \text{foreach } i \leq n \text{ do } O_A[i]() := x'_k \stackrel{R}{\leftarrow} T_k; x_r \stackrel{R}{\leftarrow} T_r; \\ \text{let } x_m : \text{bitstring} = \text{enc}'(\text{Z}_k, x_k, x_r) \text{ in} \\ \text{return } (x_m, \text{mac}'(x_m, x_{mk}))$$

where Q''_B remains as above.

6 Proof Strategy

Extension: This description is based on the initial proof strategy. Additional advice has been added since then.

Up to now, we have described the available game transformations. Next, we explain how we organize these transformations in order to prove protocols.

At the beginning of the proof and after each successful cryptographic transformation (that is, a transformation of Section 5.2), the prover executes **simplify** and tests whether the desired security properties are proved, as described in Section 4. If so, it stops.

In order to perform the cryptographic transformations and the other syntactic transformations, our proof strategy relies on the idea of advice. Precisely, the prover tries to execute each available cryptographic transformation in turn. When such a cryptographic transformation fails, it returns some syntactic transformations that could make the desired transformation work. (These are the advised transformations.) Then the prover tries to perform these syntactic transformations. If they fail, they may also suggest other advised transformations, which are then executed. When the syntactic transformations finally succeed, we retry the desired cryptographic transformation, which may succeed or fail, perhaps with new advised transformations, and so on.

The prover determines the advised transformations as follows:

- Assume that we try to execute a cryptographic transformation, and need to recognize a certain term M of L , but we find in Q_0 only part of M , the other parts being variable accesses $x[\dots]$ while we expect function applications. In this case, we advise **remove_assign binder** x . For example, if Q_0 contains $\text{sign}(M', x_{sk}, x_r)$ and we look for $\text{sign}(x_m, \text{skgen}(x_{rk}), x_r)$, we advise **remove_assign binder** x_{sk} . (The transformation of Example 8 is advised for this reason.)
- When we try to execute **remove_assign binder** x , x has several definitions, and there are accesses to variable x guarded by **find** in Q_0 , we advise **SArename** x .
- When we check whether x is secret or one-session secret, we have an assignment $\text{let } x[\tilde{i}] : T = y[\tilde{M}]$ in P , and there is at least one assignment defining y , we advise **remove_assign binder** y .

When we check whether x is secret or one-session secret, we have an assignment $\text{let } x[\tilde{i}] : T = y[\tilde{M}]$ in P , y is defined by random choices, y has several definitions, and some variable accesses to y are not of the form $\text{let } y'[\tilde{i}'] : T = y[\tilde{M}']$ in P' , we advise **SArename** y .

7 Conclusion

The tool CryptoVerif produces proofs by sequences of games like those manually written by cryptographers. It generates the games, using an automatic proof strategy or guidance from the user, who specifies the transformations to perform. It supports a wide variety of cryptographic primitives specified by indistinguishability axioms. Many of these primitives are included in a library so that the user does not have to redefine them. It can prove secrecy, correspondence, and indistinguishability properties. It has been applied to substantial case studies, including Signal [52], TLS 1.3 [23], and WireGuard [58].

CryptoVerif still has limitations. In particular, the size of games tends to grow too fast, which limits its ability to deal with large examples, especially because some game transformations require the game to be expanded first by the **expand** transformation, which duplicates the code from each test until the end of protocol. Planned improvements include allowing more game transformations to work without previous application of **expand**; allowing internal oracle calls in games, in order to share code between different parts of the game; using composition results in order to make proofs more modular. Moreover, some game transformations could be generalized. For instance, the transformation **merge_branches** merges branches of a test when they execute the same code; the detection that several branches execute equivalent code could be made more flexible, by allowing reorderings of instructions for instance. CryptoVerif only considers blackbox adversaries: it does not support proofs that manipulate the code of the adversary, such as the forking lemma [65].

Acknowledgments We warmly thank David Pointcheval for his advice and explanations of the computational proofs of protocols. This project would not have been possible without him. We also thank Jacques Stern for initiating this work. The design and implementation of CryptoVerif was partly done while Bruno Blanchet was at CNRS and at Ecole Normale Supérieure. This work was partly supported by the French National Research Agency (ANR) under the projects FormaCrypt (ARA SSIA 2005), ProSe (VERSO 2010, decision number 2010-VERS-004), TECAP (decision number ANR-17-CE39-0004-03) and received funding from the France 2030 program managed by the ANR under the references ANR-22-PECY-0006 (PEPR Cybersecurity SVP) and ANR-22-PETQ-0008 (PEPR Quantic PQ-TLS).

References

- [1] M. Abadi, B. Blanchet, and C. Fournet. The applied pi calculus: Mobile values, new names, and secure communication. *Journal of the ACM*, 65(1):1:1–1:41, Oct. 2017.
- [2] M. Abadi and C. Fournet. Mobile values, new names, and secure communication. In *28th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'01)*, pages 104–115, New York, NY, Jan. 2001. ACM Press.
- [3] M. Abadi and P. Rogaway. Reconciling two views of cryptography (the computational soundness of formal encryption). *Journal of Cryptology*, 15(2):103–127, 2002.
- [4] M. Abdalla, P.-A. Fouque, and D. Pointcheval. Password-based authenticated key exchange in the three-party setting. *IEEE Proceedings Information Security*, 153(1):27–39, Mar. 2006.
- [5] R. Affeldt, D. Nowak, and K. Yamada. Certifying assembly with formal cryptographic proofs: the case of BBS. In *9th International Workshop on Automated Verification of Critical Systems (AVoCS'09)*, volume 23 of *Electronic Communications of the EASST*. EASST, Sept. 2009.
- [6] J. B. Almeida, M. Barbosa, G. Barthe, M. Campagna, E. Cohen, B. Grégoire, V. Pereira, B. Portela, P.-Y. Strub, and S. Tasiran. A machine-checked proof of security for AWS key management service. In *ACM SIGSAC Conference on Computer and Communications Security (CCS'19)*, pages 63–78, New York, NY, Nov. 2019. ACM Press.
- [7] M. Backes, D. Hofheinz, and D. Unruh. CoSP: A general framework for computational soundness proofs. In *ACM Conference on Computer and Communications Security (CCS'09)*, pages 66–78, New York, NY, Nov. 2009. ACM Press.
- [8] M. Backes and P. Laud. Computationally sound secrecy proofs by mechanized flow analysis. In *13th ACM Conference on Computer and Communications Security (CCS'06)*, pages 370–379, New York, NY, Nov. 2006. ACM Press.
- [9] M. Backes and B. Pfitzmann. Symmetric encryption in a simulatable Dolev-Yao style cryptographic library. In *17th IEEE Computer Security Foundations Workshop*, pages 204–218, Los Alamitos, CA, June 2004. IEEE Computer Society Press.
- [10] M. Backes, B. Pfitzmann, and M. Waidner. A composable cryptographic library with nested operations. In *10th ACM conference on Computer and communication security (CCS'03)*, pages 220–230, New York, NY, Oct. 2003. ACM Press.
- [11] D. Baelde, S. Delaune, A. Koutsos, C. Jacomme, and S. Moreau. An interactive prover for protocol verification in the computational model. In *42nd IEEE Symposium on Security and Privacy (S&P'21)*, pages 537–554, Los Alamitos, CA, May 2021. IEEE Computer Society Press.
- [12] G. Barthe, J. M. Crespo, Y. Lakhnech, and B. Schmidt. Mind the gap: Modular machine-checked proofs of one-round key exchange protocols. In E. Oswald and M. Fischlin, editors, *Advances in Cryptology – EUROCRYPT 2015*, volume 9057 of *Lecture Notes in Computer Science*, pages 689–718, Berlin, Heidelberg, Apr. 2015. Springer.
- [13] G. Barthe, M. Daubignard, B. Kapron, and Y. Lakhnech. Computational indistinguishability logic. In *17th ACM Conference on Computer and Communications Security (CCS'10)*, pages 375–386, New York, NY, Oct. 2010. ACM Press.

- [14] G. Barthe, B. Grégoire, S. Z. Béguelin, and Y. Lakhnech. Beyond provable security. Verifiable IND-CCA security of OAEP. In A. Kiayias, editor, *Topics in Cryptology - CT-RSA 2011*, volume 6558 of *Lecture Notes in Computer Science*, pages 180–196, Berlin, Heidelberg, Feb. 2011. Springer.
- [15] G. Barthe, B. Grégoire, S. Héraud, and S. Z. Béguelin. Formal certification of ElGamal encryption. A gentle introduction to CertiCrypt. In P. Degano, J. Guttman, and F. Martinelli, editors, *5th International Workshop on Formal Aspects in Security and Trust, FAST 2008*, volume 5491 of *Lecture Notes in Computer Science*, pages 1–19, Berlin, Heidelberg, 2009. Springer.
- [16] G. Barthe, B. Grégoire, S. Héraud, and S. Z. Béguelin. Computer-aided security proofs for the working cryptographer. In P. Rogaway, editor, *Advances in Cryptology - CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 71–90, Berlin, Heidelberg, Aug. 2011. Springer.
- [17] G. Barthe, B. Grégoire, and S. Zanella. Formal certification of code-based cryptographic proofs. In *36th ACM SIGPLAN - SIGACT Symposium on Principles of Programming Languages (POPL'09)*, pages 90–101, New York, NY, Jan. 2009. ACM Press.
- [18] D. Basin, A. Lochbihler, and S. R. Sefidgar. CryptHOL: Game-based proofs in higher-order logic. *Journal of Cryptology*, 33:494–566, 2020.
- [19] S. Z. Béguelin, G. Barthe, S. Héraud, B. Grégoire, and D. Hedin. A machine-checked formalization of sigma-protocols. In *23rd Computer Security Foundations Symposium (CSF'10)*, pages 246–260, Los Alamitos, CA, July 2010. IEEE Computer Society Press.
- [20] S. Z. Béguelin, B. Grégoire, G. Barthe, and F. Olmedo. Formally certifying the security of digital signature schemes. In *30th IEEE Symposium on Security and Privacy, S&P 2009*, pages 237–250, Los Alamitos, CA, May 2009. IEEE Computer Society Press.
- [21] M. Bellare and C. Namprempre. 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.
- [22] 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>.
- [23] K. Bhargavan, B. Blanchet, and N. Kobeissi. Verified models and reference implementations for the TLS 1.3 standard candidate. In *IEEE Symposium on Security and Privacy (S&P'17)*, pages 483–503, Los Alamitos, CA, May 2017. IEEE Computer Society Press.
- [24] B. Blanchet. Automatic proof of strong secrecy for security protocols. In *IEEE Symposium on Security and Privacy*, pages 86–100, May 2004.
- [25] 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>.

- [26] B. Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Transactions on Dependable and Secure Computing*, 5(4):193–207, Oct.–Dec. 2008.
- [27] B. Blanchet. Automatically verified mechanized proof of one-encryption key exchange. Cryptology ePrint Archive, Report 2012/173, Apr. 2012. Available at <http://eprint.iacr.org/2012/173>.
- [28] B. Blanchet. Modeling and verifying security protocols with the applied pi calculus and ProVerif. *Foundations and Trends in Privacy and Security*, 1(1–2):1–135, Oct. 2016.
- [29] B. Blanchet. Cryptoverif: a computationally-sound security protocol verifier (initial version with communication on channels). Research report RR-9525, Inria, Oct. 2023. Available at <https://inria.hal.science/hal-04246199>.
- [30] B. Blanchet, M. Abadi, and C. Fournet. Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming*, 75(1):3–51, Feb.–Mar. 2008.
- [31] I. Boureau, C. C. Drăgan, F. Dupressoir, D. Gérard, and P. Lafourcade. Mechanised models and proofs for distance-bounding. In *34th IEEE Computer Security Foundations Symposium (CSF'21)*, Los Alamitos, CA, 2021. IEEE Computer Society Press.
- [32] 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.
- [33] D. Cadé and B. Blanchet. From computationally-proved protocol specifications to implementations and application to SSH. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, 4(1):4–31, Mar. 2013.
- [34] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Symposium on Foundations of Computer Science (FOCS)*, pages 136–145, Los Alamitos, CA, Oct. 2001. IEEE Computer Society Press. An updated version is available at Cryptology ePrint Archive, <http://eprint.iacr.org/2000/067>.
- [35] R. Canetti, L. Cheung, D. Kaynar, M. Liskov, N. Linch, O. Pereira, and R. Segala. Time-bounded task-PIOAs: A framework for analyzing security protocols. In S. Dolev, editor, *20th Symposium on Distributed Computing (DISC)*, volume 4167 of *Lecture Notes in Computer Science*, pages 238–253, Berlin, Heidelberg, Sept. 2006. Springer.
- [36] R. Canetti and J. Herzog. Universally composable symbolic analysis of cryptographic protocols (the case of encryption-based mutual authentication and key exchange). Cryptology ePrint Archive, Report 2004/334, 2004. Available at <http://eprint.iacr.org/2004/334>.
- [37] H. Comon-Lundh and V. Cortier. Computational soundness of observational equivalence. In *15th ACM conference on Computer and communications security (CCS'08)*, pages 109–118, New York, NY, Oct. 2008. ACM Press.
- [38] V. Cortier, C. C. Drăgan, F. Dupressoir, B. Schmidt, P.-Y. Strub, and B. Warinschi. Machine-checked proofs of privacy for electronic voting protocols. In *IEEE Symposium on Security and Privacy (SP'17)*, pages 993–1008, Los Alamitos, CA, 2017. IEEE Computer Society Press.

- [39] V. Cortier, H. Hördegen, and B. Warinschi. Explicit randomness is not necessary when modeling probabilistic encryption. In C. Dima, M. Minea, and F. Tiplea, editors, *Workshop on Information and Computer Security (ICS 2006)*, volume 186 of *Electronic Notes in Theoretical Computer Science*, pages 49–65. Elsevier, Sept. 2006.
- [40] V. Cortier, S. Kremer, and B. Warinschi. A survey of symbolic methods in computational analysis of cryptographic systems. *Journal of Automated Reasoning*, 46(3-4):225–259, Apr. 2011.
- [41] V. Cortier and B. Warinschi. Computationally sound, automated proofs for security protocols. In M. Sagiv, editor, *Proc. 14th European Symposium on Programming (ESOP'05)*, volume 3444 of *Lecture Notes in Computer Science*, pages 157–171, Berlin, Heidelberg, Apr. 2005. Springer.
- [42] V. Cortier and B. Warinschi. A composable computational soundness notion. In *18th ACM Conference on Computer and Communications Security (CCS'11)*, pages 63–74, New York, NY, Oct. 2011. ACM Press.
- [43] J. Courant, M. Daubignard, C. Ene, P. Lafourcade, and Y. Lakhnech. Towards automated proofs for asymmetric encryption schemes in the random oracle model. In *15th ACM conference on Computer and communications security (CCS'08)*, pages 371–380, New York, NY, Oct. 2008. ACM Press.
- [44] J. Courant, M. Daubignard, C. Ene, P. Lafourcade, and Y. Lakhnech. Automated proofs for asymmetric encryption. In D. Dams, U. Hannemann, and M. Steffen, editors, *Concurrency, Compositionality, and Correctness*, volume 5930 of *Lecture Notes in Computer Science*, pages 300–321, Berlin, Heidelberg, 2010. Springer.
- [45] J. Courant, C. Ene, and Y. Lakhnech. Computationally sound typing for non-interference: The case of deterministic encryption. In V. Arvind and S. Prasad, editors, *27th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'07)*, volume 4855 of *Lecture Notes in Computer Science*, pages 364–375, Berlin, Heidelberg, Dec. 2007. Springer.
- [46] A. Datta, A. Derek, J. C. Mitchell, V. Shmatikov, and M. Turuani. Probabilistic polynomial-time semantics for a protocol security logic. In L. Caires and L. Monteiro, editors, *ICALP 2005: the 32nd International Colloquium on Automata, Languages and Programming*, volume 3580 of *Lecture Notes in Computer Science*, pages 16–29, Berlin, Heidelberg, July 2005. Springer.
- [47] A. Datta, A. Derek, J. C. Mitchell, and B. Warinschi. Computationally sound compositional logic for key exchange protocols. In *19th IEEE Computer Security Foundations Workshop (CSFW'06)*, pages 321–334, Los Alamitos, CA, July 2006. IEEE Computer Society Press.
- [48] J. Gancher, S. Gibson, P. Singh, S. Dharanikota, and B. Parno. OWL: Compositional verification of security protocols via an information-flow type system. In *2023 IEEE Symposium on Security and Privacy (S&P)*, pages 1114–1131, Los Alamitos, CA, May 2023. IEEE Computer Society Press.
- [49] S. Goldwasser, S. Micali, and R. Rivest. A digital signature scheme secure against adaptive chosen-message attacks. *SIAM Journal of Computing*, 17(2):281–308, Apr. 1988.

- [50] R. Janvier, Y. Lakhnech, and L. Mazaré. Completing the picture: Soundness of formal encryption in the presence of active adversaries. In M. Sagiv, editor, *Proc. 14th European Symposium on Programming (ESOP'05)*, volume 3444 of *Lecture Notes in Computer Science*, pages 172–185, Berlin, Heidelberg, Apr. 2005. Springer.
- [51] D. E. Knuth and P. B. Bendix. Simple word problems in universal algebras. In J. Leech, editor, *Computational Problems in Abstract Algebra*, pages 263–297. Pergamon Press, Oxford, UK, 1970.
- [52] N. Kobeissi, K. Bhargavan, and B. Blanchet. Automated verification for secure messaging protocols and their implementations: A symbolic and computational approach. In *2nd IEEE European Symposium on Security and Privacy (EuroS&P'17)*, pages 435–450, Los Alamitos, CA, Apr. 2017. IEEE Computer Society Press.
- [53] P. Laud. Handling encryption in an analysis for secure information flow. In P. Degano, editor, *Programming Languages and Systems, 12th European Symposium on Programming, ESOP'03*, volume 2618 of *Lecture Notes in Computer Science*, pages 159–173, Berlin, Heidelberg, Apr. 2003. Springer.
- [54] P. Laud. Symmetric encryption in automatic analyses for confidentiality against active adversaries. In *IEEE Symposium on Security and Privacy*, pages 71–85, May 2004.
- [55] 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.
- [56] P. Laud and I. Tšahhrov. A user interface for a game-based protocol verification tool. In P. Degano and J. Guttman, editors, *6th International Workshop on Formal Aspects in Security and Trust (FAST2009)*, volume 5983 of *Lecture Notes in Computer Science*, pages 263–278, Berlin, Heidelberg, Nov. 2009. Springer.
- [57] P. Laud and V. Vene. A type system for computationally secure information flow. In M. Liškiewicz and R. Reischuk, editors, *15th International Symposium on Fundamentals of Computation Theory (FCT'05)*, volume 3623 of *Lecture Notes in Computer Science*, pages 365–377, Berlin, Heidelberg, Aug. 2005. Springer.
- [58] B. Lipp, B. Blanchet, and K. Bhargavan. A mechanised cryptographic proof of the Wire-Guard virtual private network protocol. In *IEEE European Symposium on Security and Privacy (EuroS&P'19)*, pages 231–246, Stockholm, Sweden, June 2019. IEEE Computer Society.
- [59] 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.
- [60] D. Nowak. A framework for game-based security proofs. In S. Qing, H. Imai, and G. Wang, editors, *Information and Communications Security, 9th International Conference, ICICS 2007*, volume 4861 of *Lecture Notes in Computer Science*, pages 319–333, Berlin, Heidelberg, Dec. 2007. Springer.
- [61] D. Nowak. On formal verification of arithmetic-based cryptographic primitives. In P. J. Lee and J. H. Cheon, editors, *Information Security and Cryptology - ICISC 2008, 11th International Conference*, volume 5461 of *Lecture Notes in Computer Science*, pages 368–382, Berlin, Heidelberg, Dec. 2008. Springer.

- [62] D. Nowak and Y. Zhang. A calculus for game-based security proofs. In *Provable Security, Fourth International Conference, ProvSec 2010*, volume 6402 of *Lecture Notes in Computer Science*, pages 35–52, Berlin, Heidelberg, Oct. 2010. Springer.
- [63] T. Okamoto and D. Pointcheval. The gap-problems: a new class of problems for the security of cryptographic schemes. In K. Kim, editor, *International Workshop on Practice and Theory in Public Key Cryptography (PKC'2001)*, volume 1992 of *Lecture Notes in Computer Science*, pages 104–118, Berlin, Heidelberg, Feb. 2001. Springer.
- [64] A. Petcher and G. Morrisett. The foundational cryptography framework. In R. Focardi and A. C. Myers, editors, *4th International Conference on Principles of Security and Trust (POST'15)*, volume 9036 of *Lecture Notes in Computer Science*, pages 53–72, Berlin, Heidelberg, Apr. 2015. Springer.
- [65] D. Pointcheval and J. Stern. Security proofs for signature schemes. In U. Maurer, editor, *Advances in Cryptology — EUROCRYPT '96*, volume 1070 of *Lecture Notes in Computer Science*, pages 387–398, Berlin, Heidelberg, May 1996. Springer.
- [66] V. Shoup. A proposal for an ISO standard for public-key encryption, Dec. 2001. ISO/IEC JTC 1/SC27.
- [67] V. Shoup. OAEP reconsidered. *Journal of Cryptology*, 15(4):223–249, Sept. 2002.
- [68] 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>.
- [69] G. Smith and R. Alpízar. Secure information flow with random assignment and encryption. In *4th ACM Workshop on Formal Methods in Security Engineering (FMSE'06)*, pages 33–43, Nov. 2006.
- [70] C. Sprenger, M. Backes, D. Basin, B. Pfizmann, and M. Waidner. Cryptographically sound theorem proving. In *19th IEEE Computer Security Foundations Workshop (CSFW-19)*, pages 153–166, Los Alamitos, CA, July 2006. IEEE Computer Society Press.
- [71] C. Sprenger and D. Basin. Cryptographically-sound protocol-model abstractions. In *23rd Annual IEEE Symposium on Logic in Computer Science*, pages 3–17, Los Alamitos, CA, June 2008. IEEE Computer Society Press.
- [72] I. Tšahhirov and P. Laud. Application of dependency graphs to security protocol analysis. In G. Barthe and C. Fournet, editors, *3rd Symposium on Trustworthy Global Computing (TGC'07)*, volume 4912 of *Lecture Notes in Computer Science*, pages 294–311, Berlin, Heidelberg, Nov. 2007. Springer.
- [73] T. Y. C. Woo and S. S. Lam. Authentication for distributed systems. In D. Denning and P. Denning, editors, *Internet Besieged: Countering Cyberspace Scofflaws*, pages 319–355. ACM Press and Addison-Wesley, Oct. 1997.

Inria

**RESEARCH CENTRE
PARIS**

2 rue Simone Iff - CS 42112
75589 Paris Cedex 12

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399