

Computationally Sound Mechanized Proofs of Correspondence Assertions

Bruno Blanchet

CNRS, École Normale Supérieure, Paris

blanchet@di.ens.fr

Abstract

We present a new mechanized prover for showing correspondence assertions for cryptographic protocols in the computational model. Correspondence assertions are useful in particular for establishing authentication. Our technique produces proofs by sequences of games, as standard in cryptography. These proofs are valid for a number of sessions polynomial in the security parameter, in the presence of an active adversary. Our technique can handle a wide variety of cryptographic primitives, including shared- and public-key encryption, signatures, message authentication codes, and hash functions. It has been implemented in the tool CryptoVerif and successfully tested on examples from the literature.

1. Introduction

Correspondence assertions on cryptographic protocols are properties of the form “if some events have been executed, then some other events have been executed”, where each event corresponds to a certain point in the protocol, possibly with arguments. An event can be formalized by a special instruction event $e(M_1, \dots, M_m)$, which simply records that the event $e(M_1, \dots, M_m)$ has been executed. Woo and Lam [61] introduced correspondence assertions to express the authentication properties of cryptographic protocols, such as “if B terminates a run of the protocol, apparently with A , then A has started a run of the protocol, apparently with B .” This property can be written more formally “if event $B_{\text{terminates}}(A)$ has been executed, then event $A_{\text{starts}}(B)$ has been executed”, where event $B_{\text{terminates}}(X)$ occurs at the point where B terminates a run and he thinks he talks to X , and event $A_{\text{starts}}(Y)$ occurs at the point where A starts a run with Y . Correspondence assertions have become a standard tool for reasoning on cryptographic protocols.

The main novelty of our work lies in the model in which we prove correspondence assertions. Indeed, there are two main models for cryptographic protocols. In the compu-

tational model, cryptographic primitives are functions on bitstrings and the adversary is a polynomial-time probabilistic Turing machine. In this realistic model, proofs are usually manual. In the formal, Dolev-Yao model, cryptographic primitives are considered as perfect blackboxes represented by function symbols, and the adversary is restricted to compute with these blackboxes. There already exist several techniques for proving correspondence assertions automatically in this abstract model, e.g. [17, 36]. However, in general, these proofs are not sound with respect to the computational model.

Since the seminal paper by Abadi and Rogaway [6], there has been much interest in relating both models [4, 11, 14, 30, 31, 38, 39, 48, 49], to show the soundness of the Dolev-Yao model with respect to the computational model, and thus obtain automatic proofs of protocols in the computational model. However, this approach has limitations: since the computational and Dolev-Yao models do not correspond exactly, additional hypotheses are necessary in order to guarantee soundness. (For example, for symmetric encryption, key cycles have to be excluded, or a specific security definition of encryption is needed [8].)

In this paper, we adopt a different approach: our tool proves correspondences directly in the computational model. In order to achieve such proofs, we extend our previous approach for secrecy [19, 20]. We produce proofs by sequences of games, as used by cryptographers [16, 55–57]: the initial game represents the protocol, for which we want to prove that the probability of breaking a certain correspondence is negligible; intermediate games are obtained each from the previous one by transformations such that the difference of probability between consecutive games is negligible; the final game is such that the desired probability can directly be shown to be negligible from the form of the game. The desired probability is then negligible in the initial game.

In order to extend our approach to correspondence assertions, we slightly extend the calculus that we use to represent games, so that it can specify events. The game transformations that we used for secrecy can also be used for correspondences, without change. However, we still need

to check that the correspondence holds on the final game. So, we introduce a rich language of correspondence assertions, and show how to check them automatically. This language allows one to specify both injective correspondences (if some event has been executed n times, then some other events have been executed at least n times) and non-injective correspondences (if some events have been executed, then some other events have been executed at least once), as well as properties of the form “if some events have been executed, then some formula holds”.

Moreover, we also show how to use correspondences in order to prove mutual authentication and authenticated key exchange. Mutual authentication is an immediate consequence of correspondences. The situation is more subtle for authenticated key exchange: intuitively, we need to prove the secrecy of the key. Since the key is shared between two participants of the protocol, the secrecy of the key is not simply the secrecy of a single variable, as we could prove in [19, 20]. However, we show that by combining correspondences with the secrecy of the variable that contains the key for one of the participants of the protocol, we can prove the standard notion of authenticated key exchange.

The prover succeeds in a fully automatic way for many examples. For delicate cases, our prover allows the user to indicate the main game transformations to perform, such as applying the security of a certain cryptographic primitive for a certain secret key. Importantly, the prover is always sound, whatever indications the user gives.

The verification of correspondences has been implemented in our prover `CryptoVerif` (19200 lines of Ocaml for version 1.06 of `CryptoVerif`), available at <http://www.di.ens.fr/~blanchet/cryptoc-eng.html>.

Related Work Results that show the soundness of the Dolev-Yao model with respect to the computational model, e.g. [31, 39, 49], make it possible to use Dolev-Yao provers in order to prove correspondences in the computational model. In particular, a tool [29] has been built based on [31] in order to make computational proofs using the Dolev-Yao prover AVISPA, for protocols that use public-key encryption and signatures. However, computational soundness 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).

Several frameworks exist for formalizing proofs of protocols in the computational model. Backes, Pfitzmann, and Waidner [10–12] have designed an abstract cryptographic library including symmetric and public-key encryption, message authentication codes, signatures, and nonces and shown its soundness with respect to computational primitives, under arbitrary active attacks. This framework

shares some limitations with the computational soundness results, for instance the exclusion of key cycles and the fact that symmetric encryption has to be authenticated. It relates the computational model to a non-standard version of the Dolev-Yao model, in which the length of messages is still present. It has been used for a computationally-sound machine-checked proof of the Needham-Schroeder-Lowe protocol [58].

Canetti and Herzog [26] show how a Dolev-Yao-style symbolic analysis can be used to prove security properties of protocols (including authentication) within the framework of universal composability [24], for a restricted class of protocols using public-key encryption as only cryptographic primitive. Then, they use the automatic Dolev-Yao verification tool `ProVerif` [18] for verifying protocols in this framework.

Canetti et al. [25] use the framework of time-bounded task-PIOAs (Probabilistic Input/Output Automata) for proving cryptographic protocols in the computational model. This framework allows them to combine probabilistic and non-deterministic behaviors.

Lincoln et al. [44, 45, 47, 50, 54] developed a probabilistic polynomial-time calculus for the analysis of security protocols. This calculus comes with a notion of process equivalence, used in particular to prove authentication properties in [45]. This calculus resembles ours in that both are probabilistic polynomial-time variants of the pi calculus. (The restriction chooses a fresh random number. The replication is polynomially bounded.) However, it differs from our calculus since it uses an explicit probabilistic scheduler while, in our calculus, the adversary schedules the processes. Our calculus also adds arrays in order to store all values of variables, which is key to our proofs, as we shall see below.

Datta et al. [32, 33] have designed a computationally sound logic that enables them to prove computational security properties using a logical deduction system.

Corin and Hartog [28] use a probabilistic Hoare-style logic for formalizing game-based cryptographic proofs.

All these frameworks can be used to prove security properties of protocols in the computational sense, but except for [26] which relies on a Dolev-Yao prover and for the machine-checked proofs of [58], they have not been mechanized up to now, as far as we know.

Other works provide proofs in the computational model, but only for secrecy. Laud [41] designed an automatic analysis for protocols using shared-key encryption, with passive adversaries. He extended it to active adversaries, but with only one session of the protocol [42]. The type system of [9, 43] handles shared-key and public-key encryption, with an unbounded number of sessions. This system relies on the Backes-Pfitzmann-Waidner library.

Barthe, Cederquist, and Tarento [13, 59] have formal-

ized the generic model and the random oracle model in the interactive theorem prover Coq, and proved signature schemes in this framework. In contrast to our specialized prover, proofs in generic interactive theorem provers require a lot of human effort, to build a detailed enough proof for the theorem prover to check it.

Halevi [37] explains that implementing an automatic prover based on sequences of games would be useful, and suggests ideas in this direction, but does not actually implement one.

Outline The next section recalls the process calculus that we use to represent games and extends it with events. Section 3 defines the correspondence assertions that we prove. Section 4 recalls the definition of observational equivalence and extends it with events. Section 5 illustrates on an example the game transformations used in our proofs. Section 6 details how we prove correspondences. Section 7 shows how to prove standard notions of authentication and authenticated key exchange using correspondences. Finally, Section 8 summarizes our experimental results and Section 9 concludes. The long version of the paper [21] contains details on the semantics of the calculus, the proof engine we use for reasoning on games, the proofs of our results, and our experiments.

2. A Calculus for Games

In this section, we review the process calculus defined in [19, 20] in order to model games used in computational security proofs. This calculus has been carefully designed to make the automatic proof of cryptographic protocols easier. We extend this calculus with parametric events, which serve in the definition of correspondences.

We illustrate this calculus on the following example, inspired by the corrected Woo-Lam public key protocol [62]:

$$\begin{aligned} B &\rightarrow A : (N, B) \\ A &\rightarrow B : \{pk_A, B, N\}_{sk_A} \end{aligned}$$

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; the prover can obviously also handle the version with A 's identity). The signatures are assumed to be (existentially) unforgeable under chosen message attacks (UF-CMA) [35], 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 process G_0 , explained below:

$$\begin{aligned} G_0 &= c_0(); \text{ new } rk_A : \textit{keyseed}; \text{ let } pk_A = \text{pkgen}(rk_A) \text{ in} \\ &\quad \text{let } sk_A = \text{skgen}(rk_A) \text{ in } \overline{c_1}\langle pk_A \rangle; (Q_A \mid Q_B) \\ Q_A &= !^{i_A \leq n} c_2[i_A](x_N : \textit{nonce}, x_B : \textit{host}); \\ &\quad \text{event } e_A(pk_A, x_B, x_N); \text{ new } r : \textit{seed}; \\ &\quad \overline{c_3}[i_A]\langle \text{sign}(\text{concat}(pk_A, x_B, x_N), sk_A, r) \rangle \\ Q_B &= !^{i_B \leq n} c_4[i_B](x_{pk_A} : \textit{pkey}); \text{ new } N : \textit{nonce}; \\ &\quad \overline{c_5}[i_B]\langle N, B \rangle; c_6[i_B](s : \textit{signature}); \\ &\quad \text{if } \text{verify}(\text{concat}(x_{pk_A}, B, N), x_{pk_A}, s) \text{ then} \\ &\quad \text{if } x_{pk_A} = pk_A \text{ then event } e_B(x_{pk_A}, B, N) \end{aligned}$$

The process G_0 is assumed to run in interaction with an adversary, which also models the network. G_0 first receives an empty message on channel c_0 , sent by the adversary. Then, it chooses randomly with uniform probability a bitstring rk_A in the type $\textit{keyseed}$, by the construct $\text{new } rk_A : \textit{keyseed}$. A type T , such as $\textit{keyseed}$, aims at denoting a set of bitstrings. However, the considered set of bitstrings depends on the security parameter η , which determines the length of keys. So, more precisely, a type T corresponds for each value of η to a set of bitstrings denoted by $I_\eta(T)$. 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 outputs the public key pk_A on channel c_1 , so that the adversary has this public key.

After outputting this message, the control passes to the receiving process, which is part of the adversary. Several processes are then made available, which represent the roles of A and B in the protocol: the process $Q_A \mid Q_B$ is the parallel composition of Q_A and Q_B ; it makes simultaneously available the processes defined in Q_A and Q_B . Let Q'_A and Q'_B be such that $Q_A = !^{i_A \leq n} Q'_A$ and $Q_B = !^{i_B \leq n} Q'_B$. The replication $!^{i_A \leq n} Q'_A$ represents n copies of the process Q'_A , indexed by the replication index i_A . (The symbol n corresponds to an integer $I_\eta(n)$ for each value of the security parameter η ; $I_\eta(n)$ is required to be a polynomially bounded function of η .) The process Q'_A begins with an input on channel $c_2[i_A]$; the channel is indexed with i_A so that the adversary can choose which copy of the process Q'_A receives the message by sending it on channel $c_2[i_A]$ for the appropriate value of i_A . The situation is similar for Q'_B , which expects a message on channel $c_4[i_B]$. The adversary can then run each copy of Q'_A or Q'_B simply by sending a message on the appropriate channel $c_2[i_A]$ or $c_4[i_B]$.

The process Q'_B first expects on channel $c_4[i_B]$ a message x_{pk_A} in the type \textit{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*: a type T is *large* when the inverse of its cardinal $\frac{1}{|T_\eta(T)|}$ is negligible, so that collisions between independent random numbers chosen uniformly in a large type have negligible probability. (The probability $f(\eta)$ is *negligible* when for all polynomials q , there exists $\eta_0 \in \mathbb{N}$ such that for all $\eta > \eta_0$, $f(\eta) \leq \frac{1}{q(\eta)}$. The probability $f(\eta)$ is *overwhelming* when $1 - f(\eta)$ is negligible.) Q'_B sends the message (N, B) on channel $c_5[i_B]$. The control then passes to the receiving process, included in the adversary. This process is expected to forward this message (N, B) on channel $c_2[i_A]$, but may proceed differently in order to mount an attack against the protocol.

Upon receiving a message (x_N, x_B) on channel $c_2[i_A]$, 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 outputs 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 receiving process, which should forward this message on channel $c_6[i_B]$ if it wishes to run the protocol correctly.

Upon receiving a message s on $c_6[i_B]$, 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 .

In our calculus, all variables defined under a replication are implicitly arrays. For example, the variable x_N defined under $!^{i_A \leq n}$ is implicitly an array indexed by the replication index i_A : x_N is an abbreviation for $x_N[i_A]$. Similarly, x_B is an abbreviation for $x_B[i_A]$, r for $r[i_A]$, x_{pk_A} for $x_{pk_A}[i_B]$, N for $N[i_B]$, and s for $s[i_B]$. Using arrays allows us to remember the values of the variables in each copy of the processes, so that the whole state of the system is available. In our calculus, arrays replace lists often used by cryptographers in their proofs. For example, during the proof, all messages signed under sk_A would be stored in a list, and by the unforgeability of signatures, when the verification of the signature of a message succeeds, we would be sure that this message occurs in the list. In our calculus, we will store messages in arrays instead. Arrays come with a lookup construct: `find $u_1 \leq n_1, \dots, u_m \leq n_m$ suchthat defined(M_1, \dots, M_l) \wedge M then P else P'` looks for indices u_1, \dots, u_m such that M_1, \dots, M_l are defined and M is true. When

such indices are found, it executes P ; otherwise, it executes P' . When several values of indices are possible, each possible value is chosen with the same probability. For example, `find $u \leq n$ suchthat defined($x_N[u]$) \wedge $x_N[u] = N$` then P looks for an index u such that $x_N[u]$ is defined and equal to N . Here, the `find` construct does not occur in the initial game, but will be introduced by game transformations.

As detailed in [19, 20], we require some *well-formedness invariants* to guarantee that bitstrings are of their expected type and that arrays are used properly (that each cell of an array is assigned at most once during execution and that variables are accessed only after being initialized).

All processes of our calculus run in probabilistic polynomial time. The semantics of the calculus is defined by a probabilistic reduction relation on semantic configurations \mathbb{C} . We denote by `initConfig(Q)` the initial configuration associated to process Q . We refer the reader to [19] for additional details on this calculus and its semantics. Given a mapping ρ from variable names to bitstrings, we write $\rho, M \Downarrow a$ when the term M (built from function symbols and variables, without array accesses) evaluates to bitstring a . We denote by `Pr[$Q \rightsquigarrow \bar{c}(a)$]` the probability that Q outputs the bitstring a on channel c after some reductions. We denote by \mathcal{E} a sequence of events of the form $e(a_1, \dots, a_n)$, where e is an event symbol and a_1, \dots, a_n are bitstrings. We denote by `Pr[$\exists(\mathbb{C}, \mathcal{E}), \text{initConfig}(Q) \xrightarrow{\mathcal{E}} \mathbb{C} \wedge \phi(\mathbb{C}, \mathcal{E})$]` the probability that there exists a sequence of events \mathcal{E} and a semantic configuration \mathbb{C} such that Q reduces to \mathbb{C} , executing events \mathcal{E} on the trace, and the logical formula $\phi(\mathbb{C}, \mathcal{E})$ holds. We denote by `Pr[$Q \rightsquigarrow \mathcal{E}$]` = `Pr[$\exists \mathbb{C}, \text{initConfig}(Q) \xrightarrow{\mathcal{E}} \mathbb{C} \wedge \mathbb{C}$ does not reduce]` the probability that the process Q executes exactly the sequence of events \mathcal{E} , in the order of \mathcal{E} . These probabilities depend on the security parameter η ; we omit it to lighten notations.

We use an *evaluation context* C to represent the adversary. An evaluation context is a process with a hole, of one of the following forms: a hole $[\]$, a process in parallel with an evaluation context $Q \mid C$, or a restriction `newChannel $c; C$` , which limits the scope of the channel c to the context C . We denote by $C[Q]$ the process obtained by replacing the hole of C with Q . When V is a set of variables defined in Q , an evaluation context C is said to be *acceptable* for (Q, V) if and only if C does not contain events, the common variables of C and Q are in V , and $C[Q]$ satisfies the well-formedness invariants. The set V contains the variables the context is allowed to access (using `find`).

When P is under replications $!^{i_1 \leq n_1} \dots !^{i_m \leq n_m}$, we say that the *replication indices at P* are i_1, \dots, i_m . We denote by \tilde{i} a sequence of replication indices i_1, \dots, i_m , and by \tilde{M} a sequence of terms M_1, \dots, M_m . We denote by `fc(P)` the set of free channels of P , and by `var(P)` the set of variables that occur in P . We also use the notation `var(\cdot)` for contexts, terms, and formulas.

3. Definition of Correspondences

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

3.1. Non-injective Correspondences

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

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

Terms M, M_1, \dots, M_m in formulae must not contain array accesses, and their variables are assumed to be distinct from variables of processes. The formula M holds when M evaluates to true. The formula $\text{event}(e(M_1, \dots, M_m))$ holds when the event $e(M_1, \dots, M_m)$ has been executed. The conjunction and disjunction are defined as usual. More formally, we write $\rho, \mathcal{E} \vdash \phi$ when the sequence of events \mathcal{E} satisfies the formula ϕ , in the environment ρ that maps variables to bitstrings. We define $\rho, \mathcal{E} \vdash \phi$ as follows:

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

Formulae denoted by ψ are conjunctions of events.

Definition 1 The sequence of events \mathcal{E} satisfies the correspondence $\psi \Rightarrow \phi$, written $\mathcal{E} \vdash \psi \Rightarrow \phi$, if and only if for all ρ defined on $\text{var}(\psi)$ such that $\rho, \mathcal{E} \vdash \psi$, there exists an extension ρ' of ρ to $\text{var}(\phi)$ such that $\rho', \mathcal{E} \vdash \phi$.

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

Definition 2 The process Q satisfies the correspondence $\psi \Rightarrow \phi$ with public variables V if and only if for all evaluation contexts C acceptable for (Q, V) , $\text{Pr}[\exists(\mathbb{C}, \mathcal{E}), \text{initConfig}(C[Q]) \xrightarrow{\mathcal{E}} \mathbb{C} \wedge \mathcal{E} \not\vdash \psi \Rightarrow \phi]$ is negligible.

A process satisfies $\psi \Rightarrow \phi$ when the probability that it generates a sequence of events \mathcal{E} that does not satisfy $\psi \Rightarrow \phi$ is negligible, in the presence of an adversary represented by the context C .

Example 1 Referring to the example G_0 of Section 2, the correspondence

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

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

The correspondence

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

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.

3.2. Injective Correspondences

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

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

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

This definition differs from the case of non-injective correspondences in that we propagate the pseudo-formula ϕ^τ

and, in the case of injective events, we make sure that the event has been executed at step τ by requiring that $\tau \neq \perp$ and $e(a_1, \dots, a_m) = \mathcal{E}(\tau)$.

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

Definition 3 The sequence of events \mathcal{E} satisfies the *correspondence* $\psi \Rightarrow \phi$, written $\mathcal{E} \vdash \psi \Rightarrow \phi$, if and only if there exists a component-wise injective \mathbb{F} such that for all ρ defined on $\text{var}(\psi)$, for all ψ^τ such that $\rho, \mathcal{E} \vdash^{\psi^\tau} \psi$, there exists an extension ρ' of ρ to $\text{var}(\phi)$ such that $\rho', \mathcal{E} \vdash^{\mathbb{F}(\psi^\tau)} \phi$.

Intuitively, a sequence of events \mathcal{E} satisfies $\psi \Rightarrow \phi$ when, if \mathcal{E} satisfies ψ with execution steps defined by ψ^τ , then \mathcal{E} satisfies ϕ with execution steps defined by $\mathbb{F}(\psi^\tau)$. The injectivity is guaranteed because \mathbb{F} is component-wise injective. Definition 2 is unchanged for injective correspondences.

Example 2 Referring to the example G_0 of Section 2, the correspondence

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

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

The correspondence

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

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

When no injective event occurs in $\psi \Rightarrow \phi$, Definition 3 reduces to the definition of non-injective correspondences.

3.3. Property

The next lemma is straightforward. It shows that correspondences are preserved by adding a context.

Lemma 1 *If Q satisfies a correspondence c with public variables V and C is an evaluation context acceptable for (Q, V) , then for all $V' \subseteq V \cup (\text{var}(C) \setminus \text{var}(Q))$, $C[Q]$ satisfies c with public variables V' .*

4. Observational Equivalence

The notion of observational equivalence is key to proofs by sequences of games. It can be seen as an adaptation to the computational model of the notion of observational equivalence used in the spi calculus [3] in the Dolev-Yao model. We review the definition observational equivalence and its properties, adapting them to the presence of events.

In the next definition, we use an evaluation context C to represent an algorithm that tries to distinguish Q from Q' .

Definition 4 (Observational equivalence) Let Q and Q' be two processes that satisfy the well-formedness invariants. Let V be a set of variables defined in Q and Q' , with the same types.

We say that Q and Q' are *observationally equivalent* with public variables V , written $Q \approx^V Q'$, when for all evaluation contexts C acceptable for (Q, V) and (Q', V) , for all channels c and bitstrings a , $|\Pr[C[Q] \rightsquigarrow \bar{c}(a)] - \Pr[C[Q'] \rightsquigarrow \bar{c}(a)]|$ is negligible and $\sum_{\mathcal{E}} |\Pr[C[Q] \rightsquigarrow \mathcal{E}] - \Pr[C[Q'] \rightsquigarrow \mathcal{E}]|$ is negligible.

This definition formalizes that the probability that an algorithm C distinguishes the games Q and Q' is negligible. The context C is allowed to access directly the variables in V (using `find`). When V is empty, we write $Q \approx Q'$.

This definition makes events observable, so that observationally equivalent processes execute the same events with overwhelming probability.

The following lemma is straightforward:

Lemma 2 1. \approx^V is reflexive, symmetric, and transitive.

2. If $Q \approx^V Q'$ and C is an evaluation context acceptable for (Q, V) and (Q', V) , then for all $V' \subseteq V \cup (\text{var}(C) \setminus (\text{var}(Q) \cup \text{var}(Q')))$, $C[Q] \approx^{V'} C[Q']$.
3. If $Q \approx^V Q'$ and Q satisfies a correspondence c with public variables V , then so does Q' .

The transitivity of \approx^V and Property 3 of Lemma 2 are key to performing proofs by sequences of games. Indeed, our prover starts from a game G_0 corresponding to the real protocol, and builds a sequence of observationally equivalent

games $G_0 \approx^V G_1 \approx^V \dots \approx^V G_m$. By transitivity, we conclude that $G_0 \approx^V G_m$. By Property 3, if G_m satisfies a certain correspondence with public variables V , then so does G_0 . The sequence $G_0 \approx^V G_1 \approx^V \dots \approx^V G_m$ is built by game transformations. Some of these transformations rely on security assumptions of cryptographic primitives; others are syntactic transformations used to simplify games. Since these transformations are the same for correspondences as for secrecy, we do not detail them here, and refer the reader to [19, 20]. (These transformations leave events unchanged.) Next, we illustrate them on an example.

5. A Proof by a Sequence of Games

In this section, we explain the transformations performed on the process G_0 of Section 2. 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:

$$\begin{aligned}
G_1 &= c_0(); \text{ new } rk_A : \text{keyseed}; \\
&\quad \text{let } pk_A = \text{pkgen}'(rk_A) \text{ in } \overline{c_1}(pk_A); (Q_{1A} \mid Q_{1B}) \\
Q_{1A} &= !^{i_A \leq n} c_2[i_A](x_N : \text{nonce}, x_B : \text{host}); \\
&\quad \text{event } e_A(pk_A, x_B, x_N); \\
&\quad \text{let } m = \text{concat}(pk_A, x_B, x_N) \text{ in} \\
&\quad \text{new } r : \text{seed}; \overline{c_3}[i_A](\text{sign}'(m, \text{skgen}'(rk_A), r)) \\
Q_{1B} &= !^{i_B \leq n} c_4[i_B](x_{pk_A} : \text{pkey}); \text{ new } N : \text{nonce}; \\
&\quad \overline{c_5}[i_B](N, B); c_6[i_B](s : \text{signature}); \\
&\quad \text{find } u \leq n \text{ such that defined}(m[u], x_B[u], x_N[u]) \\
&\quad \quad \wedge (x_{pk_A} = pk_A) \wedge (B = x_B[u]) \wedge (N = x_N[u]) \\
&\quad \quad \wedge \text{verify}'(\text{concat}(x_{pk_A}, B, N), x_{pk_A}, s) \text{ then} \\
&\quad \text{event } e_B(x_{pk_A}, B, N)
\end{aligned}$$

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 (1) and (2) with any public variables V , in particular with $V = \emptyset$. By Lemma 2, Property 3, G_0 also satisfies these correspondences with public variables $V = \emptyset$. Let us sketch how the proof of correspondence (1) for the game G_1 will proceed. Let Q'_{1A} and Q'_{1B} such that $Q_{1A} = !^{i_A \leq n} Q'_{1A}$ and $Q_{1B} = !^{i_B \leq n} Q'_{1B}$. 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 (1). 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. In the next section, we detail and formalize this reasoning, both for non-injective and injective correspondences.

6. Proving Correspondences

In this section, we explain how our prover shows that a game satisfies a correspondence. We first sketch the technique we use for collecting properties of games, then we handle the simpler case of non-injective correspondences, before generalizing to injective correspondences.

6.1. Reasoning on Games

The proof of correspondences relies on two techniques for reasoning on games. These techniques were already

used for simplifying games, so we summarize them briefly and refer the reader to [19] for details.

First, we collect facts that hold at each program point in the game. We use the following facts: the term M means that M is true, $\text{defined}(M)$ means that M is defined, and $\text{event}(e(M_1, \dots, M_m))$ means that the event $e(M_1, \dots, M_m)$ has been executed. The set of true facts collected at program point P is denoted by \mathcal{F}_P . We collect these facts as follows:

- We take into account facts that come from assignments and tests above P . For example, in the process if M then P , we have $M \in \mathcal{F}_P$, since M is true when P is executed.

In our running example G_1 , at the program point P just after the event e_B , \mathcal{F}_P 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 when P is executed. (\mathcal{F}_P also contains other facts, which are useless for proving the desired correspondences, so we do not list them.)

- When we already know that $x[\widetilde{M}]$ is defined at P (that is, $\text{defined}(M) \in \mathcal{F}_P$ and $x[\widetilde{M}]$ is a subterm of M), some definition of $x[\widetilde{i}]$ must have been executed, with $\widetilde{i} = \widetilde{M}$, so the facts F that hold at all definitions of x also hold at P , for $\widetilde{i} = \widetilde{M}$: $F\{\widetilde{M}/\widetilde{i}\} \in \mathcal{F}_P$.

In the example G_1 , we have $\text{defined}(m[u[i_B]]) \in \mathcal{F}_P$, and, when $m[i_A]$ is defined, $\text{event}(e_A(pk_A, x_B[i_A], x_N[i_A]))$ holds, so $\text{event}(e_A(pk_A, x_B[i_A], x_N[i_A]))\{u[i_B]/i_A\} \in \mathcal{F}_P$, that is, $\text{event}(e_A(pk_A, x_B[u[i_B]], x_N[u[i_B]])) \in \mathcal{F}_P$. In order words, since m is defined at index $u[i_B]$, event e_A has been executed in the copy of Q'_{1A} of index $u[i_B]$.

Second, we use an equational prover, inspired by the Knuth-Bendix completion algorithm [40]. From a set of facts \mathcal{F} , it generates rewrite rules by orienting equalities of \mathcal{F} , and uses these rewrite rules to infer new facts from the elements of \mathcal{F} . It also takes into account that collisions between uniformly distributed random elements of a large type have negligible probability, so it transforms an equality $x[\widetilde{M}] = x[\widetilde{M}']$ into $\widetilde{M} = \widetilde{M}'$ when x is defined only by restrictions new $x : T$ and T is a large type. (If the indices were different, the considered cells of x would contain independent random numbers chosen uniformly in the large type T , so the probability of equality would be negligible.)

We say that \mathcal{F} yields a contradiction when the equational prover can derive false from \mathcal{F} (for example, when \mathcal{F} contains an inequality $M_1 \neq M_2$, rewritten by the rewrite rules into $M \neq M$, which is then rewritten into false).

6.2. Non-injective Correspondences

Intuitively, in order to prove that a process Q_0 satisfies a non-injective correspondence $\psi \Rightarrow \phi$, we collect all facts that hold at events in ψ and show that these facts imply ϕ using the equational prover.

We collect facts that hold when the event F has been executed, as follows.

Definition 5 (P follows F , $\mathcal{F}_{F,P}$) When $F = \text{event}(e(M_1, \dots, M_m))$ and P is such that event $e(M'_1, \dots, M'_m); P$ occurs in Q_0 , we say that P follows F , and we define $\mathcal{F}_{F,P} = \theta' \mathcal{F}_P \cup \{\theta' M'_j = M_j \mid j \leq m\}$ where the substitution θ' is a renaming of the replication indices at P to distinct fresh replication indices.

Intuitively, when the event F has been executed, it has been executed by some subprocess of Q_0 , so there exists a subprocess event $e(M'_1, \dots, M'_m); P$ in Q_0 such that, for some replication indices defined by θ' , the event $e(M'_1, \dots, M'_m)$ has been executed and it is equal to the event F , hence $\theta' M'_j = M_j$ holds for $j \leq m$. Moreover, since the program point P , which follows F , has been reached, $\theta' \mathcal{F}_P$ holds. Hence $\mathcal{F}_{F,P} = \theta' \mathcal{F}_P \cup \{\theta' M'_j = M_j \mid j \leq m\}$ holds.

Let θ be a substitution equal to the identity on the variables of ψ . This substitution gives values to existentially quantified variables 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'_1, \dots, M'_m such that $\text{event}(e(M'_1, \dots, M'_m)) \in \mathcal{F}$ 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.

Non-injective correspondences are proved as follows.

Proposition 1 Let $\psi \Rightarrow \phi$ be a non-injective correspondence, with $\psi = F_1 \wedge \dots \wedge F_m$. If for every P_1 that follows F_1, \dots , for every P_m that follows F_m , there exists a substitution θ equal to the identity on the variables of ψ and such that $\mathcal{F}_{F_1, P_1} \cup \dots \cup \mathcal{F}_{F_m, P_m} \Vdash_\theta \phi$, then Q_0 satisfies $\psi \Rightarrow \phi$ with any public variables V .

Intuitively, when $\psi = F_1 \wedge \dots \wedge F_m$ holds, $\mathcal{F}_{F_1, P_1} \cup \dots \cup \mathcal{F}_{F_m, P_m}$ hold. For some θ equal to the identity on ψ , $\mathcal{F}_{F_1, P_1} \cup \dots \cup \mathcal{F}_{F_m, P_m}$ implies $\theta\phi$, so $\theta\phi$ holds. Hence the correspondence is satisfied.

Example 3 Let us prove that the example G_1 satisfies (1). For $\psi = F = \text{event}(e_B(x, y, z))$, the only process P that follows F is the process after event $e_B(x_{pk_A}, B, N)$, 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 $\mathcal{F}_{F,P} = \mathcal{F}_P\{i'_B/i_B\} \cup \{x_{pk_A}[i'_B] = x, B = y, N[i'_B] = z\}$ hold for some value of i'_B , where \mathcal{F}_P has been studied in Section 6.1 and $\theta' = \{i'_B/i_B\}$.

Furthermore, the substitution θ is the identity since all variables of ϕ also occur in ψ . Then we just have to show that $\mathcal{F}_{F,P}$ implies $\phi = \text{event}(e_A(x, y, z))$, that is, $\mathcal{F}_{F,P} \models_{\theta} \text{event}(e_A(x, y, z))$. Since $\text{event}(e_A(pk_A, x_B[u[i_B]], x_N[u[i_B]])) \in \mathcal{F}_P$, we have $\text{event}(e_A(pk_A, x_B[u[i'_B]], x_N[u[i'_B]])) \in \mathcal{F}_{F,P}$, 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 $\mathcal{F}_{F,P}$: $x_{pk_A}[i'_B] = x, B = y, N[i'_B] = z, x_{pk_A}[i_B] = pk_A, B = x_B[u[i'_B]]$, and $N[i'_B] = x_N[u[i'_B]]$.

Hence, G_1 satisfies (1) 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 $\text{var}(\psi)$. It is defined on other variables when checking $\mathcal{F} \models_{\theta} M$ by trying to find θ such that $\theta M \in \mathcal{F}$, and when checking $\mathcal{F} \models_{\theta} \text{event}(e(M_1, \dots, M_m))$ 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.

6.3. 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 as follows.

We require that in the initial game of the sequence, which represents the real protocol, if the event e is used as injective event in a correspondence, then two occurrences of e always occur in different branches of find or if. This property is preserved by the game transformations, so the game Q_0 on which we test the correspondences satisfies this property. This property guarantees that for each value of the replication indices, each injective event is executed at most once.

We add as first argument of every event in Q_0 the tuple (i_1, \dots, i_m) of replication indices at the program point at which the event is executed. We add as first argument of every event in $\psi \Rightarrow \phi$ a fresh variable. Then the initial process satisfies the initial correspondence if and only if the modified process satisfies the modified correspondence.

The addition of replication indices to events allows us to distinguish executions of the same injective event: these executions always have distinct replication indices by the requirement of the previous paragraph.

We extend Definition 5 to injective events, with exactly the same definition as for non-injective events. We let I_P be the image by θ' of the tuple of replication indices at P , where θ' is the renaming defined in Definition 5.

The proof of injective correspondences extends that for non-injective correspondences: for a correspondence $\psi \Rightarrow \phi$, 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 replication indices, then each injective event of ϕ has different replication indices. In order to achieve this proof, we collect information on the replication indices of events, 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 replication indices of the considered injective event of ϕ , stored in a tuple M_0 ; these indices are computed when we prove that this event is executed;
- the replication indices of the injective events of ψ , stored as a mapping $\mathcal{I} = \{j \mapsto I_{P_j} \mid F_j \text{ is an injective event}\}$, where $\psi = F_1 \wedge \dots \wedge F_m$ and P_j is the process that executes F_j , for $j \leq m$;
- the set \mathcal{V} containing the replication indices in \mathcal{F} and the variables 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 replication indices of two executions of the injective events of ψ are different, then the replication indices of the corresponding executions of the considered injective event of ϕ are also different. 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 $\mathcal{F} \mapsto_{\theta} \phi$ used for non-injective correspondences to $\mathcal{F} \mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi$, 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:

$$\begin{aligned} \mathcal{F} &\mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \perp} M \text{ if and only if} \\ &\mathcal{F} \cup \{\neg\theta M\} \text{ yields a contradiction} \\ \mathcal{F} &\mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \perp} \text{event}(e(i, M_1, \dots, M_m)) \text{ if and only if} \\ &\text{there exist } M'_0, M'_1, \dots, M'_m \text{ such that} \\ &\text{event}(e(M'_0, M'_1, \dots, M'_m)) \in \mathcal{F} \text{ and } \mathcal{F} \cup \\ &\{\theta i \neq M'_0 \vee \bigvee_{j=1}^m \theta M_j \neq M'_j\} \text{ yields a contradiction} \\ \mathcal{F} &\mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{S}} \text{inj-event}(e(i, M_1, \dots, M_m)) \text{ if and only if} \\ &\text{there exist } M'_0, M'_1, \dots, M'_m \text{ such that} \\ &\text{event}(e(M'_0, M'_1, \dots, M'_m)) \in \mathcal{F}, \\ &\mathcal{F} \cup \{\theta i \neq M'_0 \vee \bigvee_{j=1}^m \theta M_j \neq M'_j\} \text{ yields a} \\ &\text{contradiction, and } (\mathcal{F}, M'_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}. \\ \mathcal{F} &\mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1 \wedge \mathcal{C}_2} \phi_1 \wedge \phi_2 \text{ if and only if} \\ &\mathcal{F} \mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1 \text{ and } \mathcal{F} \mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_2} \phi_2 \\ \mathcal{F} &\mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1 \vee \mathcal{C}_2} \phi_1 \vee \phi_2 \text{ if and only if} \\ &\mathcal{F} \mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_1} \phi_1 \text{ or } \mathcal{F} \mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}_2} \phi_2 \end{aligned}$$

These formulae 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 $(\mathcal{F}, M'_0, \mathcal{I}, \mathcal{V}) \in \mathcal{S}$.

Injective correspondences are proved as follows.

Proposition 2 *Let $\psi \Rightarrow \phi$ be a correspondence, with $\psi = F_1 \wedge \dots \wedge F_m$.*

Assume that, for all events e used as injective events in $\psi \Rightarrow \phi$, two occurrences of the event e always occur in different branches of find or if in Q_0 .

Assume that there exists \mathcal{C} such that $\vdash \mathcal{C}$ and for every P_1 that follows F_1, \dots , for every P_m that follows F_m , there exists a substitution θ equal to the identity on the variables of ψ and such that $\mathcal{F}_{F_1, P_1} \cup \dots \cup \mathcal{F}_{F_m, P_m} \mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi$ where $\mathcal{I} = \{j \mapsto I_{P_j} \mid F_j \text{ is an injective event}\}$ and $\mathcal{V} = \text{var}(I_{P_1}) \cup \dots \cup \text{var}(I_{P_m}) \cup \text{var}(\psi)$.

Then Q_0 satisfies $\psi \Rightarrow \phi$ with any public variables 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 $\mathcal{F}_{F_1, P_1} \cup \dots \cup \mathcal{F}_{F_m, P_m} \mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \phi$.

Example 4 Let us prove that the example G_1 satisfies (2). After adding replication indices to events, the process contains events $e_A(i_A, pk_A, x_B, x_N)$ and $e_B(i_B, x_{pk_A}, B, N)$, and we prove the correspondence $\psi \Rightarrow \phi = \text{inj-event}(e_B(i, x, y, z)) \Rightarrow \text{inj-event}(e_A(i', x, y, z))$. As in Section 6.1, we compute the set \mathcal{F}_P of facts that hold at the program point P just after event e_B . However, m is defined at index $i_A = u[i_B]$ now implies that $\text{event}(e_A(u[i_B], pk_A,$

$x_B[u[i_B]], x_N[u[i_B]]) \in \mathcal{F}_P$. The process P follows $F = \text{event}(e_B(i, x, y, z))$ and $\mathcal{F} = \mathcal{F}_{F, P} = \mathcal{F}_P\{i'_B/i_B\} \cup \{i'_B = i, x_{pk_A}[i'_B] = x, B = y, N[i'_B] = z\}$.

Similarly to the proof of $\mathcal{F} \mapsto_{\theta} \text{event}(e_A(x, y, z))$ in Example 3, we can show that $\mathcal{F} \mapsto_{\theta}^{\mathcal{I}, \mathcal{V}, \mathcal{C}} \text{event}(e_A(i', x, y, z))$ where $\mathcal{I} = \{1 \mapsto i'_B\}$ encodes the replication indices of the events of ψ , $\mathcal{V} = \{i'_B, i, x, y, z\}$ contains the replication indices of \mathcal{F} and the variables of ψ , $\mathcal{C} = \mathcal{S} = \{(\mathcal{F}, u[i'_B], \mathcal{I}, \mathcal{V})\}$. ($\mathcal{C} = \mathcal{S}$ because the formula ψ is reduced to a single event; $M'_0 = u[i_B]$ contains the replication indices of the event e_A contained in \mathcal{F} : $\text{event}(e_A(u[i'_B], pk_A, x_B[u[i'_B]], x_N[u[i'_B]])) \in \mathcal{F}$.)

In order to prove injectivity, it remains to show that $\vdash \mathcal{C}$. Let $\theta'' = \{i''_B/i'_B, i''/i, x''/x, y''/y, z''/z\}$. We need to show that $\mathcal{F} \cup \theta''\mathcal{F} \cup \{i'_B \neq i''_B, u[i'_B] = u[i''_B]\}$ yields a contradiction, that is, if the replication indices of the event e_B in ψ are distinct ($i'_B \neq i''_B$), then the replication indices of the event e_A in ϕ are also distinct ($u[i'_B] \neq u[i''_B]$).

\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 restrictions of the large type *nonce*, $N[i'_B] = N[i''_B]$ implies $i'_B = i''_B$ with overwhelming probability, by eliminating collisions. This equality contradicts $i'_B \neq i''_B$, so we obtain the desired injectivity and G_1 satisfies (2) with any public variables V .

7. Authentication and Key Exchange

In this section, we show how correspondences can be used to prove mutual authentication and authenticated key exchange, as formalized in cryptography following the seminal paper by Bellare and Rogaway [15] and more recent formalizations [7, 27].

7.1. Mutual Authentication

For simplicity, we consider a protocol that includes two roles, initiator and responder, played by two participants A and B , respectively. Other participants are included in the adversary. The protocol consists of a sequence of messages exchanged alternatively from the initiator to the responder and from the responder to the initiator. Such a configuration can be represented by a process of the form

$$Q_0 = \text{Init}; (!^{i_A \leq n} Q_A \mid !^{i_B \leq n} Q_B \mid Q_S)$$

where *Init* is an initialization process (creating keys of A and B for instance), Q_A and Q_B represent respectively the initiator A and the responder B , and Q_S represents a process that allows the adversary to register keys of other (possibly dishonest) participants, so that they can take part in sessions of the protocol with A and B . The processes Q_A and Q_B do not contain replications.

We assume that the protocol contains an odd number of rounds r , so that the first and last messages of the protocol are both from the initiator to the responder. (The other case can be handled similarly.) We assume that the process Q_A stores the messages of the protocol in variables x_1, \dots, x_r and that Q_B stores them in variables y_1, \dots, y_r . The initiator process Q_A starts by receiving a message that is not really part of the protocol, and which contains the identity Y of the responder with which A is supposed to run a session. The last (r -th) message sent by process Q_A is assumed to be a pair containing, in addition to the last message of the protocol x_r , either $\text{accept}_A(Y)$, when the protocol ran as expected, or reject , when the protocol failed. The process Q_B is assumed to send a $(r+1)$ -th message containing either $\text{accept}_B(X)$ or reject just after B received and checked the last message of the protocol, where X is the identity of its expected partner (inferred by B from the protocol messages). We designate by Q_A^i the copy of Q_A of index $i_A = i$ and by Q_B^i the copy of Q_B of index $i_B = i$. We say that Q_A^i *accepts with B* when it sends $\text{accept}_A(B)$ as second component of its last message; Q_B^i *accepts with A* when it sends $\text{accept}_B(A)$ as $(r+1)$ -th message.

A *session identifier* is a function sid of the protocol messages; $\text{sid}(x_1, \dots, x_r)$ is typically a subsequence of the messages x_1, \dots, x_r , often the whole sequence. We also define a partial session identifier $\text{sid}'(x_1, \dots, x_{r-1})$, useful since the r -th message is not available to B when A accepts. We require that $\text{sid}(x_1, \dots, x_r) = \text{sid}(y_1, \dots, y_r)$ implies $\text{sid}'(x_1, \dots, x_{r-1}) = \text{sid}'(y_1, \dots, y_{r-1})$. We say that Q_A^i and Q_B^i are (real) partners when they have the same session identifier: $\text{sid}(x_1[i], \dots, x_r[i]) = \text{sid}(y_1[i'], \dots, y_r[i'])$.

Definition 6 We say that Q_0 is a *secure mutual authentication protocol* with session identifiers sid and sid' if:

1. if the adversary just sends B to Q_A^i as first message and relays messages faithfully between Q_A^i and $Q_B^{i'}$, then Q_A^i accepts with B and $Q_B^{i'}$ accepts with A ;
2. with overwhelming probability, there exists an injective function that maps each index i of a process Q_A^i that accepts with B to the index i' of a process $Q_B^{i'}$ with expected partner A such that $\text{sid}'(x_1[i], \dots, x_{r-1}[i]) = \text{sid}'(y_1[i'], \dots, y_{r-1}[i'])$;
3. with overwhelming probability, there exists an injective function that maps each index i' of a process $Q_B^{i'}$ that accepts with A to the index i of a process Q_A^i that accepts with B such that $\text{sid}(x_1[i], \dots, x_r[i]) = \text{sid}(y_1[i'], \dots, y_r[i'])$.

In item 2, $Q_B^{i'}$ has not accepted yet when Q_A^i accepts, so we cannot require that $Q_B^{i'}$ accepts with A ; we only require that $Q_B^{i'}$ has expected partner A (so that, if it accepts later, it

accepts with A). The first condition is easy to check manually, as already noticed in [15]: it expresses that the protocol works when A and B interact without adversary. The last two conditions mean that each session of A corresponds to a distinct session of B , and conversely, with overwhelming probability. They can be verified using correspondences, as shown by the following proposition.

Proposition 3 Let Q'_0 be obtained from Q_0 by adding

- event $\text{part}_A(Y, \text{sid}'(x_1, \dots, x_{r-1}))$; event $\text{full}_A(Y, \text{sid}(x_1, \dots, x_r))$ just before A sends x_r , $\text{accept}_A(Y)$;
- event $\text{full}_B(X, \text{sid}(y_1, \dots, y_r))$ just before B sends $\text{accept}_B(X)$;
- event $\text{part}_B(X, \text{sid}'(y_1, \dots, y_{r-1}))$ just before B sends y_{r-1} .

If Q_0 satisfies the first condition of Definition 6 and Q'_0 satisfies the correspondences

$$\text{inj-event}(\text{part}_A(B, x)) \Rightarrow \text{inj-event}(\text{part}_B(A, x)) \quad (3)$$

$$\text{inj-event}(\text{full}_B(A, x)) \Rightarrow \text{inj-event}(\text{full}_A(B, x)) \quad (4)$$

with public variables $V = \emptyset$, then Q_0 is a secure mutual authentication protocol with session identifiers sid and sid' .

The proof of this proposition is straightforward from the definitions. Obviously, many other versions of authentication can be verified using correspondences, for example by requiring non-injective properties instead of injective ones or by requiring authentication in one direction only instead of mutual authentication.

7.2. Authenticated Key Exchange

We adopt the same hypotheses as for mutual authentication. Furthermore, we assume that Q_A sends or receives the j -th message of the protocol on channel $c_{Aj}[i_A]$, and similarly Q_B on channel $c_{Bj}[i_B]$. The channels $c_{Aj}[i_A]$ and $c_{Bj}[i_B]$ are not used for other purposes. We assume that, just before Q_A ends accepting, it stores the established key in variable k_A of type T , and sends $x_r, \text{accept}_A(Y)$ on channel $c_{Ar}[i_A]$. We assume that, just before Q_B ends accepting, it stores the established key in variable k_B of type T and sends $\text{accept}_B(X)$ on channel $c_{Br+1}[i_B]$.

We consider here the Real-Or-Random model [7]: the adversary is allowed to ask several test queries, which either all return the session key (real) or all return a random key (random). Our goal is to show that the adversary has a negligible probability of distinguishing these two situations. As shown in [7], the Real-Or-Random model is stronger than the Find-Then-Guess model of [15]. When the test queries

return the real session key, they are defined by the process $Q_T = Q_{TA} \mid Q_{TB}$, where

$$Q_{TA} = !^{i \leq n_T} \text{test}_A[i](u_A);$$

if defined($k_A[u_A]$) then $\overline{\text{test}_A[i]} \langle k_A[u_A] \rangle$

and Q_{TB} is defined symmetrically. When the test queries return a random key, they are defined by the process $Q'_T = Q'_{TA} \mid Q'_{TB}$, where

$$Q'_{TA} = !^{i \leq n_T} \text{test}_A[i](u_A);$$

if defined($k_A[u_A], Y[u_A]$) then

if $Y[u_A] \neq B$ then $\overline{\text{test}_A[i]} \langle k_A[u_A] \rangle$ else

find $u \leq n_T$ suchthat defined($u_A[u], r_A[u]$) \wedge

$u_A[u] = u_A$ then $\overline{\text{test}_A[i]} \langle r_A[u] \rangle$ else

find $u \leq n_T$ suchthat defined($u_B[u], r_B[u],$

$x_1[u_A], \dots, x_r[u_A], y_1[u_B[u]], \dots, y_r[u_B[u]]$) \wedge

$\text{sid}(x_1[u_A], \dots, x_r[u_A]) = \text{sid}(y_1[u_B[u]], \dots,$

$y_r[u_B[u]])$ then $\overline{\text{test}_A[i]} \langle r_B[u] \rangle$ else

new $r_A : T; \overline{\text{test}_A[i]} \langle r_A \rangle$

and Q'_{TB} is defined symmetrically. When the expected partner of A is not B , the session is executed with a dishonest participant; then, the test query Q'_{TA} returns the real key. When the test query Q'_{TA} has already been asked to the same copy of Q_A (of index $u_A[u] = u_A$), or to a copy of Q_B with the same session identifier (of index $u_B[u]$ such that $\text{sid}(x_1[u_A], \dots, x_r[u_A]) = \text{sid}(y_1[u_B[u]], \dots, y_r[u_B[u]])$), Q'_{TA} returns the same result as in the previous test query. Otherwise, Q'_{TA} returns a fresh random key r_A .

Definition 7 We say that Q_0 is a *secure authenticated key exchange* over T with session identifiers sid and sid' if Q_0 is a secure mutual authentication protocol with session identifiers sid and sid' and the following are true:

1. if the adversary just sends B to Q_A^i as first message and relays messages faithfully between Q_A^i and Q_B^i , then Q_A^i accepts with B , Q_B^i accepts with A , $k_A[i] = k_B[i]$, and this random variable is uniformly distributed in T ;
2. $Q_0 \mid Q_T \approx Q_0 \mid Q'_T$.

The first point of this definition means that the protocol works correctly when A and B interact without adversary. The second point expresses the indistinguishability when the real key (returned by Q_T) and a random key (returned by Q'_T).

As shown in [19, 20], our prover can prove the secrecy of a variable x , defined as follows:

Definition 8 (Secrecy) Assume x of type T is defined in Q under a single replication $!^{i \leq n}$. Let Q' be obtained from Q by removing events. The process Q *preserves the secrecy* of x when $Q' \mid R_x \approx Q' \mid R'_x$, where

$$R_x = !^{i \leq n'} c[i](u : [1, n]); \text{ if defined}(x[u]) \text{ then } \overline{c[i]} \langle x[u] \rangle$$

$$R'_x = !^{i \leq n'} c[i](u : [1, n]); \text{ if defined}(x[u]) \text{ then}$$

find $u' \leq n'$ suchthat defined($y[u'], u[u']$) $\wedge u[u'] = u$

then $\overline{c[i]} \langle y[u'] \rangle$ else new $y : T; \overline{c[i]} \langle y \rangle$

$c \notin \text{fc}(Q')$, and $u, u', y \notin \text{var}(Q')$.

Intuitively, this definition means that the adversary cannot distinguish the array x from an array of uniformly distributed random values by performing several test queries represented by R_x and R'_x , with non-negligible probability.

Proposition 4 Let Q'_0 be obtained from Q_0 by replacing $c_{Ar}[i_A] \langle x_r, \text{accept}_A(Y) \rangle$ with

event $\text{part}_A(Y, \text{sid}'(x_1, \dots, x_{r-1}))$;

event $\text{full}_A(Y, k_A, \text{sid}(x_1, \dots, x_r))$;

if $Y = B$ then

let $k'_A = k_A$ in $\overline{c_{Ar}[i_A]} \langle x_r, \text{accept}_A(Y) \rangle$

else

$\overline{c_{Ar}[i_A]} \langle x_r, \text{accept}_A(Y) \rangle; c_{AK}[i_A](); \overline{c_{AK}[i_A]} \langle k_A \rangle$

and $\overline{c_{Br+1}[i_B]} \langle \text{accept}_B(X) \rangle$ with

event $\text{full}_B(X, k_B, \text{sid}(y_1, \dots, y_r))$;

if $X = A$ then

$\overline{c_{Br+1}[i_B]} \langle \text{accept}_B(X) \rangle$

else

$\overline{c_{Br+1}[i_B]} \langle \text{accept}_B(X) \rangle; c_{BK}[i_B](); \overline{c_{BK}[i_B]} \langle k_B \rangle$

and adding event $\text{part}_B(X, \text{sid}'(y_1, \dots, y_{r-1}))$ just before Q_B sends y_{r-1} .

If Q_0 satisfies the first condition of Definition 7, Q'_0 preserves the secrecy of k'_A , and Q'_0 satisfies the correspondences

$$\text{inj-event}(\text{part}_A(B, x)) \Rightarrow \text{inj-event}(\text{part}_B(A, x)) \quad (5)$$

$$\text{inj-event}(\text{full}_B(A, k, x)) \Rightarrow \text{inj-event}(\text{full}_A(B, k, x)) \quad (6)$$

$$\text{event}(\text{full}_B(A, k, x)) \wedge \text{event}(\text{full}_A(B, k', x)) \Rightarrow k = k' \quad (7)$$

with public variables $\{k'_A\}$, then Q_0 is a *secure authenticated key exchange* with session identifiers sid and sid' .

The process Q'_0 adds events as for mutual authentication, except that the exchanged key is added to the events full_A

and $full_B$. Furthermore, when A runs a session with B , it stores the key in the variable k'_A . When A runs a session with $Y \neq B$, it allows the adversary to obtain the exchanged key, by sending a message on c_{AK} , and symmetrically when B runs a session with $X \neq A$. (The test queries also allow the adversary to get the key in this case.) As for Proposition 3, the first condition of Definition 7 is easy to check manually. The first two correspondences imply mutual authentication. The equivalence $Q_0 \mid Q_T \approx Q_0 \mid Q'_T$ is obtained by combining the last two correspondences with the secrecy of k'_A . Intuitively, the correspondences allow us to show that each element of k_B in a session with A is in fact also an element of k'_A (which we can find by looking for the same session identifier), so showing that k'_A cannot be distinguished from an array of independent random numbers is sufficient to show the secrecy of the key. The correspondences must be true with public variables $\{k'_A\}$, so that the context is allowed to access k'_A : in the proof, the process Q'_0 is put in a context that implements the test queries by calling the processes $R_{k'_A}$ or $R'_{k'_A}$ of Definition 8, which directly access k'_A .

8. Experimental Results

We have successfully tested our prover on examples of protocols of the literature: Yahalom [23] with and without key confirmation, Otway-Rees [53], and the original and corrected versions of Woo-Lam shared-key [36] and public-key [60, 62], Needham-Schroeder public-key [46, 51], Denning-Sacco public-key [5, 34], and Needham-Schroeder shared-key [51, 52] with and without key confirmation. For each protocol, we have tried to prove one-way or mutual authentication or authenticated key exchange, depending on the goal of the protocol. Our prover obviously does not prove properties that do not hold. It succeeds in proving properties that hold, in all cases except one: it cannot show (4) for the original version of the Needham-Schroeder shared-key protocol, because it fails to prove that $N_B[i] \neq N_B[i'] - 1$ with overwhelming probability, where N_B is a nonce.

Our prover can make subtle distinctions, which are typically not made by Dolev-Yao provers. For instance, it can model two notions of security for signatures: one in which the adversary is allowed to forge a new signature for an already signed message; the other in which the adversary cannot forge any signature. With the latter definition, for the corrected Woo-Lam public key protocol [62], it can show that the signature is authenticated (both participants have exactly the same signature), while it cannot with the former definition, because the two participants may have different signatures for the same message.

The total runtime for all these tests is 29 s on a Pentium M 1.8 GHz.

9. Conclusion

We have presented the first tool for proving correspondences by sequences of games, in the computational model. This tool works with no or very little help from the user, handles a wide variety of cryptographic primitives, and produces proofs valid for a polynomial number of sessions in the presence of an active adversary.

Although this tool can prove complex correspondences, with conjunctions and disjunctions, our examples use rather simple ones. Complex correspondences proved useful in case studies [1, 2] in the Dolev-Yao model; we plan to use them in similar situations in the computational model. Our tool can also be used to analyze protocols or combinations of primitives that are outside the scope of the Dolev-Yao model. For example, in [22], in collaboration with David Pointcheval, we have used it to prove the Full Domain Hash signature scheme. We plan to consider other such examples in the future.

Acknowledgments We thank David Pointcheval for very helpful discussions on this paper. This work was partly supported by the ANR project ARA SSIA Formacrypt.

References

- [1] M. Abadi and B. Blanchet. Computer-assisted verification of a protocol for certified email. *Science of Computer Programming*, 58(1–2):3–27, Oct. 2005. Special issue SAS'03.
- [2] M. Abadi, B. Blanchet, and C. Fournet. Just fast keying in the pi calculus. *ACM Transactions on Information and System Security*. To appear. An extended abstract appears in ESOP'04.
- [3] M. Abadi and A. D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148(1):1–70, Jan. 1999. An extended version appeared as Digital Equipment Corporation Systems Research Center report No. 149, January 1998.
- [4] M. Abadi and J. Jürjens. Formal eavesdropping and its computational interpretation. In *TACS'01*, volume 2215 of *LNCS*, pages 82–94. Springer, Oct. 2001.
- [5] M. Abadi and R. Needham. Prudent engineering practice for cryptographic protocols. *IEEE Transactions on Software Engineering*, 22(1):6–15, Jan. 1996.
- [6] M. Abadi and P. Rogaway. Reconciling two views of cryptography (the computational soundness of formal encryption). *Journal of Cryptology*, 15(2):103–127, 2002.
- [7] M. Abdalla, P.-A. Fouque, and D. Pointcheval. Password-based authenticated key exchange in the three-party setting. *IEE Proceedings Information Security*, 153(1):27–39, Mar. 2006.
- [8] P. Adão, G. Bana, J. Herzog, and A. Scedrov. Soundness of formal encryption in the presence of key-cycles. In *ESORICS'05*, volume 3679 of *LNCS*, pages 374–396. Springer, Sept. 2005.

- [9] M. Backes and P. Laud. Computationally sound secrecy proofs by mechanized flow analysis. In *CCS'06*, pages 370–379. ACM, Nov. 2006.
- [10] M. Backes and B. Pfitzmann. Symmetric encryption in a simulatable Dolev-Yao style cryptographic library. In *CSFW'04*, pages 204–218. IEEE, June 2004.
- [11] M. Backes, B. Pfitzmann, and M. Waidner. A composable cryptographic library with nested operations. In *CCS'03*, pages 220–230. ACM, Oct. 2003.
- [12] M. Backes, B. Pfitzmann, and M. Waidner. Symmetric authentication within a simulatable cryptographic library. In *ESORICS'03*, volume 2808 of *LNCS*, pages 271–290. Springer, Oct. 2003.
- [13] G. Barthe, J. Cederquist, and S. Tarento. A machine-checked formalization of the generic model and the random oracle model. In *IJCAR'04*, volume 3097 of *LNCS*, pages 385–399. Springer, July 2004.
- [14] M. Baudet, V. Cortier, and S. Kremer. Computationally sound implementations of equational theories against passive adversaries. In *ICALP'05*, volume 3580 of *LNCS*, pages 652–663. Springer, July 2005.
- [15] M. Bellare and P. Rogaway. Entity authentication and key distribution. In *CRYPTO'93*, volume 773 of *LNCS*, pages 232–249. Springer, Aug. 1993.
- [16] M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *EUROCRYPT'06*, volume 4004 of *LNCS*, pages 409–426. Springer, May 2006. Extended version available at <http://eprint.iacr.org/2004/331>.
- [17] B. Blanchet. From secrecy to authenticity in security protocols. In *SAS'02*, volume 2477 of *LNCS*, pages 342–359. Springer, Sept. 2002.
- [18] B. Blanchet. Automatic proof of strong secrecy for security protocols. In *IEEE Symposium on Security and Privacy*, pages 86–100, May 2004.
- [19] B. Blanchet. A computationally sound mechanized prover for security protocols. Cryptology ePrint Archive, Report 2005/401, Nov. 2005. Available at <http://eprint.iacr.org/2005/401>.
- [20] B. Blanchet. A computationally sound mechanized prover for security protocols. In *IEEE Symposium on Security and Privacy*, pages 140–154, May 2006.
- [21] B. Blanchet. Computationally sound mechanized proofs of correspondence assertions. Cryptology ePrint Archive, Report 2007/128, Apr. 2007. Available at <http://eprint.iacr.org/2007/128>.
- [22] B. Blanchet and D. Pointcheval. Automated security proofs with sequences of games. In *CRYPTO'06*, volume 4117 of *LNCS*, pages 537–554. Springer, Aug. 2006.
- [23] M. Burrows, M. Abadi, and R. Needham. A logic of authentication. *Proceedings of the Royal Society of London A*, 426:233–271, 1989.
- [24] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS'01*, pages 136–145. IEEE, Oct. 2001. An updated version is available at Cryptology ePrint Archive, <http://eprint.iacr.org/2000/067>.
- [25] 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 *DISC'06*, volume 4167 of *LNCS*, pages 238–253. Springer, Sept. 2006.
- [26] R. Canetti and J. Herzog. Universally composable symbolic analysis of mutual authentication and key exchange protocols. In *TCC'06*, volume 3876 of *LNCS*, pages 380–403. Springer, Mar. 2006. Extended version available at <http://eprint.iacr.org/2004/334>.
- [27] R. Canetti and H. Krawczyk. Analysis of key-exchange protocols and their use for building secure channels. In *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 453–474. Springer, May 2001.
- [28] R. Corin. A probabilistic Hoare-style logic for game-based cryptographic proofs. In *ICALP'06*, volume 4052 of *LNCS*, pages 252–263. Springer, July 2006.
- [29] V. Cortier, H. Hördegen, and B. Warinschi. Explicit randomness is not necessary when modeling probabilistic encryption. In *ICS 2006*, Sept. 2006. Proceedings to appear.
- [30] V. Cortier, S. Kremer, R. Küsters, and B. Warinschi. Computationally sound symbolic secrecy in the presence of hash functions. In *FSTTCS'06*, volume 4246 of *LNCS*, pages 176–187. Springer, Dec. 2006.
- [31] V. Cortier and B. Warinschi. Computationally sound, automated proofs for security protocols. In *ESOP'05*, volume 3444 of *LNCS*, pages 157–171. Springer, Apr. 2005.
- [32] A. Datta, A. Derek, J. C. Mitchell, V. Shmatikov, and M. Turani. Probabilistic polynomial-time semantics for a protocol security logic. In *ICALP'05*, volume 3580 of *LNCS*, pages 16–29. Springer, July 2005.
- [33] A. Datta, A. Derek, J. C. Mitchell, and B. Warinschi. Computationally sound compositional logic for key exchange protocols. In *CSFW'06*, pages 321–334. IEEE Computer Society, July 2006.
- [34] D. E. Denning and G. M. Sacco. Timestamps in key distribution protocols. *Commun. ACM*, 24(8):533–536, Aug. 1981.
- [35] S. Goldwasser, S. Micali, and R. Rivest. A digital signature scheme secure against adaptative chosen-message attacks. *SIAM Journal of Computing*, 17(2):281–308, Apr. 1988.
- [36] A. Gordon and A. Jeffrey. Authenticity by typing for security protocols. In *CSFW'01*, pages 145–159. IEEE Computer Society, June 2001.
- [37] S. Halevi. A plausible approach to computer-aided cryptographic proofs. Cryptology ePrint Archive, Report 2005/181, June 2005. Available at <http://eprint.iacr.org/2005/181>.
- [38] J. Herzog. A computational interpretation of Dolev-Yao adversaries. In *WITS'03*, pages 146–155, Apr. 2003.
- [39] R. Janvier, Y. Lakhnech, and L. Mazaré. Completing the picture: Soundness of formal encryption in the presence of active adversaries. In *ESOP'05*, volume 3444 of *LNCS*, pages 172–185. Springer, Apr. 2005.
- [40] D. E. Knuth and P. B. Bendix. Simple word problems in universal algebras. In *Computational Problems in Abstract Algebra*, pages 263–297. Pergamon Press, 1970.
- [41] P. Laud. Handling encryption in an analysis for secure information flow. In *ESOP'03*, volume 2618 of *LNCS*, pages 159–173. Springer, Apr. 2003.

- [42] P. Laud. Symmetric encryption in automatic analyses for confidentiality against active adversaries. In *IEEE Symposium on Security and Privacy*, pages 71–85, May 2004.
- [43] P. Laud. Secrecy types for a simulatable cryptographic library. In *CCS'05*, pages 26–35. ACM, Nov. 2005.
- [44] P. D. Lincoln, J. C. Mitchell, M. Mitchell, and A. Scedrov. A probabilistic poly-time framework for protocol analysis. In *CCS'98*, pages 112–121, Nov. 1998.
- [45] P. D. Lincoln, J. C. Mitchell, M. Mitchell, and A. Scedrov. Probabilistic polynomial-time equivalence and security protocols. In *FM'99*, volume 1708 of *LNCS*, pages 776–793. Springer, Sept. 1999.
- [46] G. Lowe. Breaking and fixing the Needham-Schroeder public-key protocol using FDR. In *TACAS'96*, volume 1055 of *LNCS*, pages 147–166. Springer, 1996.
- [47] P. Mateus, J. Mitchell, and A. Scedrov. Composition of cryptographic protocols in a probabilistic polynomial-time process calculus. In *CONCUR'03*, volume 2761 of *LNCS*, pages 327–349. Springer, Sept. 2003.
- [48] D. Micciancio and B. Warinschi. Completeness theorems for the Abadi-Rogaway logic of encrypted expressions. *Journal of Computer Security*, 12(1):99–129, 2004.
- [49] D. Micciancio and B. Warinschi. Soundness of formal encryption in the presence of active adversaries. In *TCC'04*, volume 2951 of *LNCS*, pages 133–151. Springer, Feb. 2004.
- [50] 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.
- [51] R. M. Needham and M. D. Schroeder. Using encryption for authentication in large networks of computers. *Commun. ACM*, 21(12):993–999, Dec. 1978.
- [52] R. M. Needham and M. D. Schroeder. Authentication revisited. *Operating Systems Review*, 21(1):7, 1987.
- [53] D. Otway and O. Rees. Efficient and timely mutual authentication. *Operating Systems Review*, 21(1):8–10, 1987.
- [54] A. Ramanathan, J. Mitchell, A. Scedrov, and V. Teague. Probabilistic bisimulation and equivalence for security analysis of network protocols. In *FOSSACS'04*, volume 2987 of *LNCS*, pages 468–483. Springer, Mar. 2004.
- [55] V. Shoup. A proposal for an ISO standard for public-key encryption, Dec. 2001. ISO/IEC JTC 1/SC27.
- [56] V. Shoup. OAEP reconsidered. *Journal of Cryptology*, 15(4):223–249, Sept. 2002.
- [57] 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>.
- [58] C. Sprenger, M. Backes, D. Basin, B. Pfitzmann, and M. Waidner. Cryptographically sound theorem proving. In *CSFW'06*, pages 153–166. IEEE, July 2006.
- [59] S. Tarento. Machine-checked security proofs of cryptographic signature schemes. In *ESORICS'05*, volume 3679 of *LNCS*, pages 140–158. Springer, Sept. 2005.
- [60] T. Y. C. Woo and S. S. Lam. Authentication for distributed systems. *Computer*, 25(1):39–52, Jan. 1992.
- [61] T. Y. C. Woo and S. S. Lam. A semantic model for authentication protocols. In *Proceedings IEEE Symposium on Research in Security and Privacy*, pages 178–194, May 1993.
- [62] T. Y. C. Woo and S. S. Lam. Authentication for distributed systems. In *Internet Besieged: Countering Cyberspace Scofflaws*, pages 319–355. ACM Press and Addison-Wesley, Oct. 1997.