



HAL
open science

The Squirrel Prover and its Logic

David Baelde, Stéphanie Delaune, Charlie Jacomme, Adrien Koutsos, Joseph Lallemand

► **To cite this version:**

David Baelde, Stéphanie Delaune, Charlie Jacomme, Adrien Koutsos, Joseph Lallemand. The Squirrel Prover and its Logic. ACM SIGLOG News, 2024, 11 (2), 10.1145/3665453.3665461 . hal-04579038

HAL Id: hal-04579038

<https://inria.hal.science/hal-04579038>

Submitted on 17 May 2024

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

The Squirrel Prover and its Logic

David Baelde, Univ Rennes, CNRS, IRISA

Stéphanie Delaune, CNRS, Univ Rennes, IRISA

Charlie Jacomme, Université de Lorraine, LORIA, Inria Nancy Grand-Est

Adrien Koutsos, Inria Paris

Joseph Lallemand, CNRS, Univ Rennes, IRISA

The Squirrel system [Baelde et al. 2021] is an interactive prover for the verification of cryptographic protocols. It relies on a dedicated higher-order logic and provides security guarantees akin to those classically obtained by cryptographers, against attackers modeled as arbitrary probabilistic polynomial-time Turing machines. In this paper, we provide a high-level introduction to the logic behind Squirrel, and briefly describe some of the interesting technical challenges encountered during its construction.

1. INTRODUCTION

Security protocols are widely used today to secure transactions that take place over public channels like the Internet. Common uses include the secure transfer of sensitive information such as credit card numbers, or user authentication on a system. Because of their presence in many widely used applications (*e.g.* electronic commerce, government-issued ID), developing methods and tools to verify security protocols has become an important research challenge. Such tools help increase our trust in protocols, and hence on the applications that rely on them.

Formal methods have brought various approaches to prove that cryptographic protocols indeed guarantee the expected security properties. An effective approach in this area of research consists in modeling cryptographic messages as first-order terms, together with an equational theory that represents attacker capabilities. This idea, originally proposed in [Dolev and Yao 1981], has been refined over the years, resulting in a variety of so-called *symbolic* models. These models encompass broad categories of attackers and facilitate the automated verification of protocols. They have led to the development of successful tools such as ProVerif [Blanchet 2001] and Tamarin [Meier et al. 2013]. However, it is important to note that security in a symbolic model does not necessarily imply security in the cryptographers' standard model, called the *computational* model. In that model, attackers are represented by probabilistic polynomial-time Turing machines (PPTMs), and one proves that a protocol is indistinguishable from an idealized, obviously secure version of it. Verification techniques for the computational model, though crucially needed, often exhibit less flexibility or automation compared to ones for symbolic models. As an illustration, secret keys are faithfully modeled in the computational model as long bitstrings that are drawn uniformly at random, whereas they are modeled using abstract names in symbolic models. In symbolic models, two distinct secret keys are represented by different names, which cannot be equal. However, in the computational model, as in reality, it is possible (although unlikely) that the sampled bitstrings are equal. In this column, we present a recent logic-based method for verifying cryptographic protocols in the computational model, and some practical aspects of its implementation in the Squirrel tool [Baelde et al. 2021; Baelde et al. 2023]. This system is built on the Computationally Complete Symbolic Attacker (CCSA) approach of [Bana and Comon-Lundh 2012; Bana and Comon-Lundh 2014], which relies on the symbolic setting of logic, but avoids the limitations of the symbolic models mentioned above. Instead of modeling attacker capabilities by rules stating what the adversary can do, the CCSA method relies on the specification of what the attacker *cannot* do. Starting from the security properties of cryptographic primitives, one derives rules that articulate which pairs of message sequences are indistinguishable. These

rules are proved sound w.r.t. the interpretation of terms as PPTMs. Therefore, a proof of a security property using these rules implies security in the computational model under the initial cryptographic assumptions. The CCSA logic was later extended into a meta-logic, which served as basis for the first version of Squirrel [Baelde et al. 2021], before being generalized to a fully-fledged higher-order logic in [Baelde et al. 2023]. The Squirrel tool is a proof assistant developed in a collaborative effort, which has been successfully used on a variety of case studies [Baelde et al. 2022; Comon et al. 2020; Cremers et al. 2022]. Instructions for installing the system are available on its website, together with a user manual, tutorials, and an in-browser interface for playing with the tool without installing it:

<https://squirrel-prover.github.io>



We first provide, in Section 2, an introduction to the computational model and the CCSA approach, showing in particular how cryptographic assumptions translate into logical rules. We elaborate on this in Section 3 to show how protocols can be modeled in CCSA logic, discussing a subtle issue w.r.t. the intended notion of security. Finally, Section 4 formally defines the higher-order CCSA logic, on which Squirrel is based, and discusses some of its interesting technical features: the distinction between local and global formulas, the subtleties tied to reasoning about probabilistic objects, and the key notion of bi-deduction for reasoning about computational indistinguishability.

Related work. Squirrel is only a recent addition to the list of available systems for verifying cryptographic protocols. We have mentioned above some tools that provide guarantees in symbolic models; we now briefly present the tools providing guarantees in the computational model. Several such systems exist, based on different approaches. The earliest one is CryptoVerif [Blanchet 2008], which mechanizes proofs based on high-level game transformations, following the style of pen-and-paper proofs. The most prominent system today is probably EasyCrypt [Barthe et al. 2011], which is a proof assistant also featuring higher-order logic. It notably embeds a domain-specific *probabilistic relational Hoare logic* [Barthe et al. 2009], which can capture cryptographic game transformations. This design makes it possible to carry out most cryptographic arguments within EasyCrypt. Other systems are currently being developed, with various goals: notably CryptHOL [Basin et al. 2020], which explores an alternative modeling technique in Isabelle/HOL, and F^* [Swamy et al. 2016], which is a general-purpose program verification framework based on refinement types that can be used (via external arguments) to provide computational security guarantees. Those approaches can be compared on several criteria [Barbosa et al. 2021].

Here, we only compare with the closest two tools, and only with respect to three criteria, to highlight differences: modeling, automation, and proof methodology. Regarding modeling, CryptoVerif and Squirrel have a similar level of detail and expressivity, although CryptoVerif supports a larger range of cryptographic assumptions. EasyCrypt is more expressive, with a higher level of detail, at the cost of modeling overhead, and thus better suited for proving properties of primitives. Protocol specifications in CryptoVerif and Squirrel are given in a process algebra, whereas protocols are encoded in EasyCrypt as APIs, which is inconvenient for interactive protocols. Finally, CryptoVerif does not support stateful protocols, while Squirrel and EasyCrypt do.

Overall, the current level of automation of Squirrel sits somewhere between CryptoVerif, which can apply cryptographic arguments automatically, although it often requires hints about which game transformations to use, and EasyCrypt, which does not automatically apply cryptographic arguments.

Finally, the most important difference between Squirrel and the other tools is the associated proof methodology: CryptoVerif relies on game transformations and EasyCrypt performs Hoare-style proofs of programs, while Squirrel reasons over execution traces of protocols.

2. A BASIC INDISTINGUISHABILITY LOGIC

We introduce in this section the CCSA approach, which is designed to reason on cryptographic protocols, *i.e.* concurrent programs relying on cryptographic primitives to achieve some functionality while running in a malicious environment. To simplify our exposition, we delay the treatment of general protocols to Section 3, and restrict our presentation to simple sequential cryptographic games for the time being. These games are used by cryptographers to express security properties of cryptographic primitives against attackers modeled as arbitrary polynomial-time computations. A key concept about games is *computational indistinguishability*: roughly, two games are indistinguishable if any polynomial-time adversary has at-best a negligible probability in distinguishing them. In the CCSA approach, we use first-order logic to establish computational indistinguishability: we first model the messages derived during a game’s execution as first-order terms, and then reason about computational indistinguishability in first-order logic, using inference rules derived from cryptographic assumptions.

A key aspect of the logic is that it abstracts cryptographic arguments as much as possible. Notably, its syntax does not mention the length of keys or the value of probabilities. Further, reduction-based arguments are implicit: they justify the soundness of the proof rules, but the code underlying reductions is never explicitly stated.

2.1. The core of the logic

Security in the real world is often conditional, as there is usually a small but non-zero probability that the adversary manages to break the security of a system. Typically, one cannot rule out that a lucky attacker guesses a secret key. To discard such unlikely cases, security is considered up to a *negligible* probability of attack. A function $f : \mathbb{N} \rightarrow [0, 1]$ is *negligible* if it is asymptotically lower than the inverse of any polynomial, *i.e.* if $\forall k \in \mathbb{N}. \exists n_0. \forall n \geq n_0. f(n) \leq \frac{1}{n^k}$. Conversely, f is *overwhelming* if $1 - f$ is negligible. In what follows, we will thus use a *security parameter* $\eta \in \mathbb{N}$ which can be, *e.g.*, the length of the secret keys, and security must hold with *overwhelming* probability w.r.t. η .

In the cryptographic literature, security properties are commonly specified using so-called *games*, in which an adversary tries to mount an attack: the game might represent an adversary attempting to guess a secret value, to forge a signature, . . . Games are usually written in pseudo-code using an imperative style, as shown in Figure 1. The statement $x \stackrel{\$}{\leftarrow} D$ stores the result of a random sampling following distribution D into variable x : *e.g.* $\text{sk} \stackrel{\$}{\leftarrow} \{0, 1\}^\eta$ uniformly samples a secret key sk of η bits. Games are parameterized by an abstract *interactive attacker*, denoted by \mathcal{A} . The statement $o \leftarrow \mathcal{A}(i)$ calls the adversary \mathcal{A} on input i , and stores its answer in variable o . The adversary is stateful, retaining information across its invocations. Standard assignments are denoted in the same way: $x \leftarrow e$ stores the result of evaluating e into variable x .

Describing cryptographic games as probabilistic imperative programs lets cryptographers rely on the reader’s intuitive understanding of their semantics, and avoids introducing convoluted execution models. However, this comes at a cost: formally reasoning over such programs can be difficult (one needs to deal with statefulness, loop invariants, probabilistic dependencies, . . .); moreover, encoding complex protocols as imperative programs is possible but not natural, and complicates security proofs.

Game \mathcal{G} :

1 : $sk \xleftarrow{\$} \{0, 1\}^\eta$;	A secret key sk of length η is randomly sampled. The
2 : $m \leftarrow \mathcal{A}()$;	interactive attacker \mathcal{A} is asked to provide a message m
3 : $t \leftarrow \text{enc}(m, sk)$;	which is then encrypted (using sk) and sent back to \mathcal{A} .
4 : $x \leftarrow \mathcal{A}(t)$;	Afterwards, \mathcal{A} returns a second message x which is the
5 : return x	final output of the game.

Fig. 1. An example cryptographic game.

Cryptographic games in CCSA. In the CCSA approach, we do not explicitly represent games. Instead, we shall only represent the messages computed at different points in a game, using first-order terms. Those are pure, unlike the stateful games, and thus easier to reason about. Our terms are built using:

- *honest function symbols* which represent the various primitives used to compute messages (e.g. pairing, encryption);
- *attacker function symbols*, noted att_i for $i \in \mathbb{N}$, modeling arbitrary computations performed by adversaries;
- *name symbols*, representing the sources of randomness: essentially, a name is a pointer to a memory cell holding a value sampled at random before the game starts.

Example 2.1. Let us illustrate this approach by modeling the final result of the game of Figure 1 as a single term. The sampling of the secret key is modeled by a name symbol sk . The first call to the attacker $\mathcal{A}()$ is modeled by the term $\text{att}_0()$. In the imperative game modeling style, internal functions (e.g. the encryption enc) can be probabilistic, but this is not so in our approach, where we precisely track probabilistic dependencies. Thus, we explicitly specify the randomness used by the encryption function enc , by introducing a name symbol r and using the term $\text{enc}(\text{att}_0(), r, sk)$ to model t in the game. Then, the second call to the adversary can be modeled using a different adversarial function symbol, as: $\text{att}_1(\text{enc}(\text{att}_0(), r, sk))$. Importantly, we have modeled two successive calls to the stateful attacker \mathcal{A} in the game by two pure functions (att_0 and att_1). This is without loss of generality, as any state computed during the first call to \mathcal{A} and used in the second call can be recomputed when modeling the second call as att_1 .

Models of the logic. Defining the logic’s semantics in order to have a faithful translation from games to terms requires a bit of care. The interpretation of a term is parameterized by the security parameter $\eta \in \mathbb{N}$, as well as two explicit sources of randomness provided by a pair $\rho = (\rho_h, \rho_a)$ of tapes filled with random bits. A model \mathbb{M} of our logic must:

- provide for each honest function symbol f its interpretation as a PPTM \mathcal{M}_f ;
- associate to each name n a unique sub-sequence of η bits in ρ_h using a PPTM \mathcal{M}_n with access to ρ_h , such that distinct names use disjoint parts of ρ_h ;
- interpret any adversarial function symbol att as a PPTM \mathcal{M}_{att} which can only access the random tape ρ_a (but not ρ_h).

For the sake of simplicity, we forced names to be interpreted as uniform and independent random samplings in $\{0, 1\}^\eta$.

Our models are first-order models with a tailored interpretation domain. Thus, terms can be interpreted in \mathbb{M} using the standard semantics of first-order logic. In our specific setting, a term t is interpreted as a function associating, to each value of the security parameter η and random tapes ρ , the interpretation $\llbracket t \rrbracket_{\mathbb{M}}^{\eta, \rho}$. More precisely, the function $\eta, \rho \mapsto \llbracket t \rrbracket_{\mathbb{M}}^{\eta, \rho}$ is computed by a PPTM.

For example, the interpretation $\llbracket n \rrbracket_{\mathbb{M}}^{\eta, \rho}$ of a name symbol n is computed by $\mathcal{M}_n(\eta, \rho_h)$, and an adversarial computation is computed by $\llbracket \text{att}_i(t) \rrbracket_{\mathbb{M}}^{\eta, \rho} = \mathcal{M}_{\text{att}_i}(\eta, \llbracket t \rrbracket_{\mathbb{M}}^{\eta, \rho}, \rho_a)$. Importantly, this interpretation is compositional and makes all probability dependencies explicit, which will facilitate reasoning over cryptographic messages.

Example 2.2. Coming back to \mathcal{G} from Figure 1, the value in the variable x at the end of \mathcal{G} 's execution follows the same probability distribution as

$$\llbracket \text{att}_1(\text{enc}(\text{att}_0(), r, \text{sk})) \rrbracket_{\mathbb{M}}^{\eta, \rho}$$

where the tapes in ρ are sampled at random, and where \mathbb{M} interprets the encryption and attacker functions as does \mathcal{G} .

Using Turing machines to interpret terms gives us a high level of expressivity. For instance, conditional branching in games can be internalized in the first-order terms using an (if b then t else e) function symbol whose interpretation is forced to be the natural one in all models.

Example 2.3. Consider the following game:

Game \mathcal{G}_{sk}

```

1:  $\text{sk} \xleftarrow{\$} \{0, 1\}^\eta$ ;
2:  $m \leftarrow \mathcal{A}()$ ;
3: if  $m = \text{sk}$  then
4:   return 0
5: else
6:   return 1
```

This game samples a fresh secret key and queries the attacker to obtain a guess m for the key. The game outputs 0 if the attacker correctly guessed the key, and 1 otherwise. The message returned at the end of the game can be represented by the following term:

$$\phi_{\text{sk}} \stackrel{\text{def}}{=} \text{if } (\text{att}_0() = \text{sk}) \text{ then } 0 \text{ else } 1.$$

2.2. The computational indistinguishability predicate \sim

One of the most commonly used notions to define cryptographic properties is *computational indistinguishability*. Roughly, indistinguishability is expressed as a guessing game in which an adversary must figure out which one of two scenarios \mathcal{G}_I or \mathcal{G}_R it is interacting with. More precisely, two games \mathcal{G}_I and \mathcal{G}_R are indistinguishable, which we write $\mathcal{G}_I \approx \mathcal{G}_R$, when any PPTM adversary has at-best a negligible probability of guessing whether it is interacting with \mathcal{G}_I or \mathcal{G}_R . Formally, we define the *advantage* of the attacker \mathcal{A} as the probability that it makes the two games behave differently. We then require this advantage to be negligible in the security parameter η :

$$\forall \mathcal{A} \in \text{PPTM}. \quad \eta \mapsto |\Pr(\mathcal{G}_I(\mathcal{A}) = 1) - \Pr(\mathcal{G}_R(\mathcal{A}) = 1)| \text{ is negligible.}$$

Typically, the game \mathcal{G}_R will correspond to an attack against the real primitives, while \mathcal{G}_I will represent an attack against an idealized implementation of the primitives for which security is obvious.

Example 2.4. Coming back to Example 2.3, the indistinguishability $\mathcal{G}_{\text{sk}} \approx (\text{return } 1)$ states that the probability that the adversary \mathcal{A} guesses the secret key sk is negligible. In that case, a simple probabilistic independence argument can be used to show that the probability that \mathcal{A} computes sk is $\frac{1}{2^\eta}$. Thus, the indistinguishability holds.

Squirrel's logic relies on a predicate \sim to represent computational indistinguishability. Formally, for two terms u, v , the predicate $u \sim v$ holds in a model \mathbb{M} if $\mathcal{G}_u^{\mathbb{M}} \approx \mathcal{G}_v^{\mathbb{M}}$, where $\mathcal{G}_t^{\mathbb{M}}$ is the game where the adversary is provided with $\llbracket t \rrbracket_{\mathbb{M}}^{\eta, \rho}$ and must produce a bit b , which the game returns. In other words, $u \sim v$ in \mathbb{M} when:

$$\forall \mathcal{A} \in \text{PPTM}. \quad \eta \mapsto |\Pr_{\rho}(\mathcal{A}(\llbracket u \rrbracket_{\mathbb{M}}^{\eta, \rho}, \rho_a) = 1) - \Pr_{\rho}(\mathcal{A}(\llbracket v \rrbracket_{\mathbb{M}}^{\eta, \rho}, \rho_a) = 1)| \text{ is negligible.}$$

$$\begin{array}{c}
\text{REFL} \\
\hline
u \sim u
\end{array}
\quad
\begin{array}{c}
\text{REWRITE} \\
\frac{(u = v) \sim \text{true} \quad C[v] \sim w}{C[u] \sim w}
\end{array}
\quad
\begin{array}{c}
\text{FRESH} \\
\hline
(t \neq n) \sim \text{true}
\end{array}
\quad
\begin{array}{l}
\text{when } t \text{ is a ground term} \\
\text{in which } n \text{ does not occur}
\end{array}$$

Fig. 2. Basic inference rules over \sim .

Example 2.5. Reusing \mathcal{G}_{sk} and ϕ_{sk} from Example 2.3, checking $\phi_{\text{sk}} \sim 1$ amounts to verifying that $\mathcal{G}_{\text{sk}} \approx (\text{return } 1)$.

In the example above, we have used the indistinguishability predicate on booleans. In that case, the semantics of \sim can be restated in a simpler way, without a quantification over all distinguishers \mathcal{A} : it simply means that the two booleans have the same probability of being 1. The general semantics of \sim becomes useful, though, to reason on it by decomposing terms: for instance, it allows us to derive $f(u) \sim f(v)$ from $u \sim v$ for any u and v of arbitrary types. Indeed, the existence of a PPTM distinguishing between $f(u)$ and $f(v)$ implies the existence of a PPTM distinguishing between u and v : the latter distinguisher is obtained by composing the former with the PPTM computing f .

Logical rules. We now present some rules that can be used to reason over \sim . We intuitively describe some of those rules here, starting with the simpler rules that do not rely on any security assumptions over the primitives.

Example 2.6. The following formulas are valid:

- (1) $t \sim t$ for any term t ; indistinguishability is reflexive.
- (2) $n \sim n'$ for any names n, n' ; two uniform random samplings are indistinguishable.
- (3) $(n, n) \not\sim (n, n')$; the attacker sees, in one case, the same value twice, and in the other, two distinct values (with overwhelming probability). Notice that this implies that \sim does not lift to an arbitrary context: we have seen that $n \sim n'$, but this example shows that we do not have for all context C that $C[n] \sim C[n']$.
- (4) $(\text{if true then } u \text{ else } v = u) \sim \text{true}$ for any terms u, v ; more generally, any term occurring in an indistinguishability can be rewritten into an equal term.
- (5) $(n \neq n') \sim \text{true}$; names, *i.e.* random samplings, collide with negligible probability.

We provide a first set of inference rules in Figure 2. **REFL** corresponds to the reflexivity of \sim , and **REWRITE** allows replacing two terms that are equal with overwhelming probability in any context. Finally, **FRESH** exploits the fact that a term syntactically contains all its probabilistic dependencies: if a name n does not occur in a term t , then n is a uniform random sampling independent of t , and thus $t = n$ can only be true with negligible probability.

Example 2.7. Let us go back to Example 2.5 and the formula $\phi_{\text{sk}} \sim 1$. We assume an additional proof rule **SIMPL₌** that allows to replace a term with another one equal to it with probability one, which we note \equiv , and where for instance, $v \equiv (\text{if false then } u \text{ else } v)$. We can prove our goal with the following derivation:

$$\begin{array}{c}
\text{REFL} \\
\hline
1 \sim 1
\end{array}
\quad
\begin{array}{c}
\text{SIMPL}_{=} \\
\hline
((\text{att}_0() = \text{sk}) = \text{false}) \sim \text{true}
\end{array}
\quad
\begin{array}{c}
\text{SIMPL}_{=} \\
\hline
(\text{if false then } 0 \text{ else } 1) \sim 1
\end{array}
\quad
\begin{array}{c}
\text{REWRITE} \\
\hline
(\text{if } \text{att}_0() = \text{sk} \text{ then } 0 \text{ else } 1) \sim 1
\end{array}$$

Here, we use twice **SIMPL₌**: once with the previously mentioned equality over a false conditional, and once using the fact that $((\text{att}_0() = \text{sk}) = \text{false}) \equiv (\text{att}_0() \neq \text{sk})$.

2.3. Unforgeability of hash functions: EUF-CMA

To reason over cryptographic games, a final piece of the puzzle is missing: assumptions over cryptographic primitives. For each primitive (a cipher, a signature, a hash, ...), there is a usual set of cryptographic assumptions made over it. Such assumptions are expressed as an indistinguishability between two games, and cryptographic proofs take the form of reduction-based arguments. To prove that the security of the primitive, for instance expressed as $P_R \approx P_I$, implies the security of some game $\mathcal{G}_R \approx \mathcal{G}_I$, we prove the contrapositive: assuming that there is a distinguisher against $\mathcal{G}_R \approx \mathcal{G}_I$, we build a distinguisher against $P_R \approx P_I$. In our logic, we hide those reduction-based arguments behind dedicated rules, one for each cryptographic assumption.

Consider the example of a keyed hash function $h(x, sk)$ and of the EUF-CMA assumption. Intuitively, a keyed hash function is a function that takes a message and a key, and produces a short unpredictable value. More precisely, if a key sk is secret and randomly sampled, then an attacker has a negligible probability of guessing the hash $h(m, sk)$ of any m , unless the attacker was directly given this value. Such a hash function is called *unforgeable*, and this is formalized in the EUF-CMA_h cryptographic game:

EUF-CMA _h	Hashing Oracle $O(x)$
1: $sk \xleftarrow{\$} \{0, 1\}^n$;	1: $\mathcal{L} \leftarrow \{x\} \cup \mathcal{L}$;
2: $\mathcal{L} \leftarrow \emptyset$;	2: return $h(x, sk)$
3: $t, m \leftarrow \mathcal{A}^O()$;	
4: return $(t = h(m, sk) \wedge m \notin \mathcal{L})$	

This game asks the attacker to return a pair of values t, m such that $t = h(m, sk)$. The attacker is allowed to interact with the hashing oracle O , but it must return a m that was never queried to O . This is enforced by storing in a set \mathcal{L} all oracle inputs.

We can express the security of the hash function by assuming that for all attackers \mathcal{A} , $\text{EUF-CMA}_h \approx (\text{return false})$. Our goal is now to derive an inference rule which is sound under this assumption.

As a warm-up, we restrict ourselves to models where the interpretation of h satisfies the EUF-CMA assumption, and analyze the validity of a few statements:

- $(h(0, sk) = \text{att}_0()) \sim \text{false}$ is valid. Otherwise, we would have a model (and thus an attacker) able to output with non-negligible probability the value of $h(0, sk)$ without having access to any other hash values. This attacker would trivially break EUF-CMA_h, by directly outputting $0, h(0, sk)$.
- $(h(0, sk) = \text{att}_1(sk)) \sim \text{false}$ is not valid, because the secret key is leaked to the attacker. As such, the computation of $\text{att}_1(sk)$ cannot be seen as a computation made by an adversary \mathcal{A}^O in the EUF-CMA_h game, and the assumption does not apply.
- $(h(0, sk) = \text{att}_1(h(1, sk))) \sim \text{false}$ is valid. Otherwise, we would have a model and thus an attacker breaking EUF-CMA_h: the attacker would compute $h(1, sk)$ with a call to oracle $O(1)$, after which $\mathcal{L} = \{1\}$, and return $\text{att}_1(h(1, sk))$, which is the hash of $0 \notin \mathcal{L}$.

To generalize this, we need to answer the following question: under which conditions over t and m is $(h(m, sk) = t) \sim \text{false}$ a valid formula? We answer this by providing sufficient conditions under which the terms t and m can be produced by some attacker interacting with the EUF-CMA_h game in such a way that m is never queried to O .

The first main condition is that all syntactic occurrences of sk in t and m must be as a key to h (*i.e.* all occurrences are of the form $h(u, sk)$ for some u). Otherwise, the terms cannot be simulated, as the adversary \mathcal{A} against EUF-CMA_h does not know the key.

Once we know that sk is only used in key position, it is easy to syntactically track in t and m what will be computed using queries to O : it is precisely the set of subterms

Game $\mathcal{G}(\mathcal{R})$	Tag \mathcal{T}	Reader (real) $\mathcal{R}_r(x)$	Reader (ideal) $\mathcal{R}_i(x)$
$\text{sk} \xleftarrow{\$} \{0, 1\}^\eta;$ $\mathcal{L} \leftarrow \emptyset;$ $b \leftarrow \mathcal{A}^{\mathcal{T}, \mathcal{R}}();$ return b	$\bar{n} \xleftarrow{\$} \{0, 1\}^\eta;$ $x \leftarrow \langle n, h(n, \text{sk}) \rangle;$ $\mathcal{L} \leftarrow \{x\} \cup \mathcal{L};$ return x	$x_n \leftarrow \text{fst}(x);$ $x_h \leftarrow \text{snd}(x);$ return $(h(x_n, \text{sk}) = x_h)$	return $x \in \mathcal{L}$

Fig. 3. Authentication game for the BASIC HASH protocol.

of the form $h(x, \text{sk})$, *i.e.* $S_{\text{sk}}^h(t, m) = \{h(u, \text{sk}) \mid h(u, \text{sk}) \text{ subterm of } t \text{ or } m\}$. Using this set, we can then define a formula, inside our logic, expressing that all previously seen hash values are distinct from m . Those two conditions are sufficient to obtain a rule enabling us to use EUF-CMA to establish indistinguishabilities in the logic:

$$\text{EUF-CMA} \frac{\bigwedge_{h(u, \text{sk}) \in S_{\text{sk}}^h(t, m)} ((u = m) \sim \text{false})}{(h(m, \text{sk}) = t) \sim \text{false}} \text{ sk only occurs as } h(x, \text{sk}) \text{ in } t \text{ and } m$$

This rule has a logical premise, which becomes a proof obligation when the rule is applied to a proof goal, and a syntactic side condition automatically checked by Squirrel.

3. MODELING INTERACTIVE PROTOCOLS WITH RECURSIVE FUNCTIONS

Now that we have introduced the core CCSA framework, let us turn to how cryptographic protocols are modeled. Compared to the simple games of the previous section, protocols are concurrent systems, involving several communicating agents. An agent may for instance be a key server, a bank, a user's terminal... When analyzing the security of a protocol, it is desirable to make worst-case assumptions. A typical one is that the adversary has full control over the network: it receives all messages output by protocol agents, and is tasked with feeding inputs to them, with messages resulting from arbitrary computations. We also assume that the adversary schedules the actions of the protocol's agents: it decides when to spawn a new session, when to advance in a session, ... This can actually be modeled using cryptographic games, representing the agents as oracles to which the adversary has access. Adversarial computations then induce a sequence of interactions with the protocol's agents, *i.e.* an execution trace of the protocol. Following this style, quantifying over all adversaries implies quantifying over all execution traces. We present this approach in more details in what follows, illustrating it on a simplified version of the BASIC HASH protocol [Brusò et al. 2010].

3.1. An example: BASIC HASH protocol

The BASIC HASH protocol is a simple RFID protocol in which a tag \mathcal{T} tries to authenticate itself to a reader \mathcal{R} . This is an access control protocol: *e.g.* the tag may be embedded in an RFID fob, and the reader could guard access to some building. The tag and the reader both share a secret key sk . Whenever the tag wants to authenticate itself to the reader, it samples a random value n , hashes it with sk using an unforgeable keyed hash function as presented in Section 2.3, and sends the pair $\langle n, h(n, \text{sk}) \rangle$ of the name and its hash digest to the reader. The reader can check that a message x it received from the network was generated by the tag by extracting the first component $\text{fst}(x)$ of x , re-computing its hash using the secret key sk and checking that this yields the second component $\text{snd}(x)$ of its input.

Authentication game. We can express the fact that this protocol provides some form of authentication using the cryptographic game presented in Figure 3. The definitions of \mathcal{T} and \mathcal{R}_r correspond to the tag and reader agents described above, with a single addition that will be useful to express the security property: before sending its output

$x = \langle n, h(n, sk) \rangle$ to the reader, the tag logs x into a set \mathcal{L} . Further, \mathcal{R}_i corresponds to an idealized reader, which checks that its input x originates from \mathcal{T} by directly verifying that x appears in the log \mathcal{L} . Intuitively, the idealized reader looks, across space and time, into the past internal memories of the tag to check whether it indeed generated x . Obviously, \mathcal{R}_i is only useful as a trick to model the expected authentication property, and cannot be physically implemented.

The cryptographic game $\mathcal{G}(\mathcal{R})$, parameterized by the reader $\mathcal{R} \in \{\mathcal{R}_r, \mathcal{R}_i\}$, samples the secret key sk , initializes the log \mathcal{L} , and runs the adversary $\mathcal{A}^{\mathcal{T}, \mathcal{R}}$ by letting it interact with the tag and reader oracles. The adversary must try to guess a bit b indicating whether it is interacting with the real or ideal reader \mathcal{R} : its guess b is the final result of the game. The protocol is secure if \mathcal{A} has a negligible advantage in guessing the correct bit b , *i.e.* if the computational indistinguishability $\mathcal{G}(\mathcal{R}_r) \approx \mathcal{G}(\mathcal{R}_i)$ holds. Indeed, if an attacker can make \mathcal{R}_r accept while breaking authentication (*i.e.* no corresponding tag produced the message), trying to do the same execution with \mathcal{R}_i would fail by construction, making it trivial to distinguish the two worlds.

3.2. Modeling protocol executions

During its execution, \mathcal{A} can call the tag and reader oracles any number of times and according to any interleaving. We are going to model the interaction of \mathcal{A} with the protocol agents \mathcal{T}, \mathcal{R} along an execution trace tr representing a fixed but arbitrary interleaving. Thus, an *execution trace* tr is a finite sequence τ_1, \dots, τ_n of *timestamps*, where the j -th timestamp τ_j in the trace represents the agent the adversary interacted with at this step. To distinguish multiple interactions with the same agents, actions are indexed by a unique replication index i in some set of indices \mathcal{I} . For BASIC HASH, we use $\tau = T(i)$ and $\tau = R(i)$ for an interaction with respectively the tag and the reader. For convenience, we also force τ_1 to a special value `init`.

$$\mathcal{T} ::= \tau_1, \dots, \tau_n \quad \tau_i ::= \text{init} \mid T(i) \mid R(i) \quad (n \in \mathbb{N}, i \in \mathcal{I})$$

Modeling execution traces. We assume a typed logic, to enable reasoning over different kinds of objects. Models of the logic provide an interpretation domain for each type, where for instance `message` is interpreted as the set of bitstrings and `bool` as $\{0, 1\}$. Intuitively, the type of all terms used in Section 2 is `message` or `bool`, with *e.g.* the equality function symbol `=` typed as `message` \rightarrow `message` \rightarrow `bool`. The types `index` and `timestamp` respectively represent the sets of indices \mathcal{I} and timestamps \mathcal{T} , and we use function symbols `T` and `R`, both of type `index` \rightarrow `timestamp`, to build timestamps from indices, and `init` : `timestamp` for the initial timestamp. The execution trace tr is implicitly encoded using two function symbols:

$$\text{happens} : \text{timestamp} \rightarrow \text{bool} \quad \cdot < \cdot : \text{timestamp} \rightarrow \text{timestamp} \rightarrow \text{bool}$$

where `happens(τ)` states that τ has been scheduled in the execution trace, and $\tau < \tau'$ that τ occurred before τ' in the trace. Each model \mathbb{M} of the logic fixes the execution trace by interpreting the types and function symbols above. We require, through an ad hoc restriction on models of the logic, that the interpretation of `index` in any model is finite and independent of the security parameter η . Further, we assume that any scheduled timestamp uniquely corresponds to either `init`, some `T(i)` or some `R(i)`.

Modeling the game's execution. In the logic, the interaction of the adversary \mathcal{A} with the oracles \mathcal{T} and \mathcal{R} along an execution trace is modeled using several mutually recursive functions called *macros*. A macro is identified by its name `m` and is evaluated at a particular timestamp τ : informally, `m@ τ` refers to the value of `m` at point τ in the execution. We will use the following macros: `input@ τ` and `output@ τ` denote the messages input and output by the oracle scheduled at time τ in the execution trace;

and $\text{frame}_{@T}$ is roughly the sequence of all messages output by any agent during the execution up to timestamp T . The frame is crucial to express security properties, as it represents the full list of messages seen by the attacker. We assume a function symbol $\text{pred} : \text{timestamp} \rightarrow \text{timestamp}$ that maps any scheduled timestamp different from init to its predecessor in the execution trace, and all other timestamps to a fixed constant undef , distinct from all scheduled timestamps. The macros can be defined as follows for the game $\mathcal{G}(\mathcal{R}_r)$ using the real reader:

$$\begin{aligned}
\text{input}_r@T &= \text{match } T \text{ with} & \text{frame}_r@T &= \text{match } T \text{ with} \\
& \quad | \text{init} \rightarrow \text{empty} & & \quad | \text{init} \rightarrow \text{empty} \\
& \quad | _ \rightarrow \text{att}_r(\text{frame}_r@\text{pred}(T)) & & \quad | _ \rightarrow \langle \text{frame}_r@\text{pred}(T), \text{output}_r@T \rangle \\
\text{output}_r@T &= \text{match } T \text{ with} \\
& \quad | \text{init} \rightarrow \text{empty} \\
& \quad | T(i) \rightarrow \langle n(i), h(n(i), \text{sk}) \rangle \\
& \quad | R(i) \rightarrow h(\text{fst}(\text{input}_r@T), \text{sk}) = \text{snd}(\text{input}_r@T)
\end{aligned}$$

where empty is a symbol of type message representing the empty bitstring, and where we generalized names to be of type $\text{index} \rightarrow \text{message}$, so that each $n(i)$ is an independent fresh name. The r subscript indicates that these macros are for the game using the real reader. The macros definitions for the game with the ideal reader $\mathcal{G}(\mathcal{R}_i)$, subscripted by i , are identical except for the output macro: we only need to introduce a macro $\mathcal{L}@T$ for the value of the log \mathcal{L} after the execution of the oracle at T , and to modify the $R(i)$ case of output :

$$\begin{aligned}
\text{output}_i@T &= \text{match } T \text{ with} & \mathcal{L}@T &= \text{match } T \text{ with} \\
& \quad | R(i) \rightarrow \text{input}_i@T \in \mathcal{L}@T & & \quad | \text{init} \rightarrow \emptyset \\
& \quad | \dots \rightarrow \dots \text{ (as before)} & & \quad | T(i) \rightarrow \{\text{output}_i@T\} \cup \mathcal{L}@T \\
& & & \quad | R(i) \rightarrow \mathcal{L}@T
\end{aligned}$$

Then, the formula $\text{frame}_r@T \sim \text{frame}_i@T$ expresses the fact that, for any model \mathbb{M} defining an execution trace tr and an adversary \mathcal{A} , the adversary has a negligible probability of distinguishing the real and ideal versions of the protocol, *i.e.* $\mathcal{G}(\mathcal{R}_r) \approx \mathcal{G}(\mathcal{R}_i)$ holds *along this particular execution trace* (see Section 3.3 below for a discussion of the security guarantees provided by this formula).

Cryptographic reasoning with recursive functions. The EUF-CMA reasoning rule presented in Section 2.3 does not directly apply to terms containing recursive definitions: we need to be able to reason over the occurrences of the hash function inside terms, and thus possibly inside an arbitrary recursive unfolding. A first way to fix the EUF-CMA rule is to forbid any occurrence of a hash in all the recursive functions occurring inside a term, but it is of course too limiting: in BASIC HASH, we have for instance a set of hash occurrences of the form $h(n(i), \text{sk})$ in output_r . We need to be able to reason over such a set, for instance using some universal quantification over indices or timestamps. We formally introduce such constructions in Section 4.2, and extend the cryptographic rules in Section 4.4.

Implementation in Squirrel. The encoding of protocols using macros as described above can be generalized to support a large class of security protocols. Given a protocol \mathcal{P} , we can use the same definitions for frame and input , and we only need to adapt the definition of output according to the protocol specification, and, if necessary, introduce additional macros for the protocol's mutable variables (such as the macro \mathcal{L} for the log \mathcal{L}). This method is expressive enough to encode complex protocol behaviors, and allows all relevant information to be expressed as messages that our logic can handle.

However, it is arguably not very intuitive to manipulate for users. To help with that, Squirrel allows users to describe protocols as processes in a process algebra (a variant of the applied π -calculus), and automatically translates them as macros (the actual definition of Squirrel macros is slightly more involved than the one presented here, we refer readers to [Baelde et al. 2022] for details).

3.3. The issue of asymptotic security

With the modeling style described here for protocols, we always reason on an arbitrary but fixed trace. When proving a property, we show that it holds (up to negligible probability) in all models of the logic, and the model includes the interpretation of the order on timestamps, *i.e.* the interleaving of actions, as well as the interpretation of all attacker function symbols. In particular, this means that in any given model, the trace and the interpretation of adversarial functions is fixed, and only then do the security parameter η and random tape ρ vary. In other words, when proving a property ϕ , what we establish can informally be described by the formula:

$$\forall \text{ trace tr. } \forall \text{ adversary } \mathcal{A}. \Pr_{\rho}(\phi \text{ holds in tr against } \mathcal{A}) \text{ is overwhelming in } \eta.$$

While such a statement gives good security guarantees, it is weaker than what is commonly expected. In cryptographic games, the adversary is indeed given access to the protocol through oracles it may call at will any (polynomial) number of times. That is, the adversary can *choose* the trace, rather than having it imposed on him. In particular, the length of the trace may be chosen by the adversary and thus depend on η , while it is fixed in our model. That would correspond to the stronger (informal) formula:

$$\forall \text{ adversary } \mathcal{A}. \Pr_{\rho}(\phi \text{ holds in tr chosen by } \mathcal{A}) \text{ is overwhelming in } \eta.$$

It is possible to overcome this limitation by changing the way we model protocols, to let the adversary adaptively choose the trace during the protocol execution. However, such alternatives have not yet been explored in practice.

4. THE GENERAL LOGIC & PROOF SYSTEM

We have considered so far a typed language with first-order terms and a single indistinguishability predicate, with inference rules that can be expressed as axioms in first-order logic. Except for the addition of recursively defined functions, all this fits within the original CCSA logic of [Bana and Comon-Lundh 2014], which is plain first-order logic. Thus, proving statements in that logic can be done using any proof system for classical first-order logic; only the CCSA axiom schemes would reflect the intended probabilistic semantics of terms and the cryptographic assumptions on primitives.

We have seen, however, that syntactic side conditions, which are needed for cryptographic rules, cannot be easily lifted to terms containing recursive definitions. In addition, the logic appears ill-suited for some security properties. Recall the BASIC HASH example, whose security we defined by using a real and an ideal system. Yet, the intuitive security property should only be about a high-level fact verified by all executions of the real system: whenever a reader accepts, the value it received was produced by an honest tag.

To improve this aspect of the logic, a core idea is to introduce a more specific logic, where special status is given to terms of type `bool`. We have seen axioms involving indistinguishability, of the form $t \sim \text{true}$ or $t \sim \text{false}$. For such formulas, using computational indistinguishability for all PPTM distinguishers is overkill, and we can take a much simpler approach: $t \sim \text{true}$ simply means that t is true with overwhelming probability. Further, it is convenient to view t itself as a specific kind of formula: to this effect, we can use function symbols representing propositional connectives, as we did for the if-then-else construction; less obviously, the same can be done for quantifiers. In that

context, t will be called a *local* formula, while the real first-order formulas are called *global*. It turns out that CCSA proofs tend to intensively reason on local formulas. This calls for a proof system that conveniently handles these like proper formulas. Squirrel comes with one such proof system, which we present in this section.

4.1. Higher-order terms with a relaxed semantics

Before properly introducing local and global formulas and their proof systems, we must come back to the semantics of terms. Indeed, since [Baelde et al. 2023], the logic behind Squirrel features higher-order terms. There have been several motivations for this significant generalization of the original CCSA logic. First, Squirrel developments predominantly consist of proofs of local formulas, and higher-order features at this level help structure these proofs in a modular way, and re-use them. Second, the constraints of the original CCSA logic, namely that all terms are interpreted as PPTMs, make sense for modeling cryptographic protocols and adversaries but are sometimes too restrictive when writing proofs. For instance, it can be useful to talk about the discrete logarithm over finite groups, even though this operation is (hopefully) not PTIME. It can also be useful to quantify over infinite types in local formulas, *e.g.* to specify properties of primitives over messages: anticipating on what follows, we will be able to write local formulas such as $\forall x, y : \text{message}. \text{fst}(\langle x, y \rangle) = x$.

The terms of our logic are thus simply-typed λ -calculus terms, additionally featuring recursive definitions. In the tool, some limited forms of polymorphism are even available, but we leave them unaccounted for in the theory. Each type τ is interpreted in a model \mathbb{M} as a family of sets $\llbracket \tau \rrbracket_{\mathbb{M}}^{\eta}$ indexed by the security parameter η , with $\llbracket \tau_1 \rightarrow \tau_2 \rrbracket_{\mathbb{M}}^{\eta} = \llbracket \tau_1 \rrbracket_{\mathbb{M}}^{\eta} \rightarrow \llbracket \tau_2 \rrbracket_{\mathbb{M}}^{\eta}$. We require that $\llbracket \text{bool} \rrbracket_{\mathbb{M}}^{\eta} = \{0, 1\}$ and $\llbracket \text{message} \rrbracket_{\mathbb{M}}^{\eta} = \{0, 1\}^*$ for all \mathbb{M} and η . The potential dependency in η is however crucial *e.g.* to faithfully model the finite groups used in Diffie-Hellman key exchanges.

Terms are no longer interpreted as PPTMs, but more generally as discrete random variables, *i.e.* functions from random tapes to the desired domain. More precisely, a model fixes a set of *finite tapes* $\mathbb{T}_{\mathbb{M}, \eta}$ for each η , and we let $\mathbb{RV}_{\mathbb{M}}(\tau)$ be the set of all η -indexed sequences of functions from $\mathbb{T}_{\mathbb{M}, \eta}$ to $\llbracket \tau \rrbracket_{\mathbb{M}}^{\eta}$. Then, given \mathbb{M} , η and ρ , any term t of type τ is interpreted as $(\rho \in \mathbb{T}_{\mathbb{M}, \eta} \mapsto \llbracket t \rrbracket_{\mathbb{M}}^{\eta, \rho})_{\eta \in \mathbb{N}} \in \mathbb{RV}_{\mathbb{M}}(\tau)$ as follows:

$$\begin{aligned} \llbracket x \rrbracket_{\mathbb{M}}^{\eta, \rho} &= \mathbb{M}(x)(\eta, \rho) \\ \llbracket t \ t' \rrbracket_{\mathbb{M}}^{\eta, \rho} &= \llbracket t \rrbracket_{\mathbb{M}}^{\eta, \rho}(\llbracket t' \rrbracket_{\mathbb{M}}^{\eta, \rho}) \\ \llbracket \lambda(x : \tau'). t \rrbracket_{\mathbb{M}}^{\eta, \rho} &= (a \in \llbracket \tau' \rrbracket_{\mathbb{M}}^{\eta} \mapsto \llbracket t \rrbracket_{\mathbb{M}[x/a]}^{\eta, \rho}) \end{aligned}$$

where $\mathbb{M}(x)$ is the interpretation of x in \mathbb{M} , and $\mathbb{M}[x/a]$ is the model modified to interpret x as a random variable X such that $X(\eta, \rho) = a$. Note that the interpretation of a term only depends on the interpretation of its subterms for the same values of η and ρ , and thus the value of $\llbracket x/a \rrbracket_{\mathbb{M}}^{\eta, \rho}(x)$ is only relevant on η, ρ . We refer the reader to [Baelde et al. 2023] for the full details, including the interpretation of recursive definitions when they satisfy a suitable well-foundedness condition, as well as how to check that some terms do correspond to PPTM computations in order to apply cryptographic assumptions.

While it could seem more natural to have infinite tapes, rather than finite but arbitrarily long ones, this has technical motivations. It ensures that all functions computed on tapes actually correspond to random variables, *i.e.* that they are measurable. This restriction of our logic (and of the original CCSA logic) is not limiting when it comes to modeling cryptographic protocols and attackers, which all run in PTIME. It is however crucial that the length of tapes can grow with η . The restriction to finite tapes is standard in formal cryptographic reasoning: it is present, *e.g.*, in approaches based on Hoare logics [Barthe et al. 2011; Petcher and Morrisett 2015; Basin et al. 2020].

4.2. Local and global formulas

Squirrel's logic is structured around two kinds of formulas: we call *global* formulas the formulas of first-order logic built around the \sim predicate, and *local* formulas the terms of type `bool`. To clearly distinguish the two, we add tildes to logical connectives and quantifiers at the global level. We also note formulas with different letters: ϕ and ψ for local formulas and F and G for global ones.

$$\begin{aligned} F &:= F \tilde{\wedge} F' \mid F \tilde{\vee} F' \mid F \tilde{\Rightarrow} F' \mid \tilde{\sim} F \mid \tilde{\forall} x.F \mid \tilde{\exists} x.F \mid \tilde{u} \sim \tilde{v} \mid [\phi] \mid \dots \\ \phi &:= \phi \wedge \phi' \mid \phi \vee \phi' \mid \phi \Rightarrow \phi' \mid \neg \phi \mid \forall x.\phi \mid \exists x.\phi \mid u = v \mid \dots \end{aligned}$$

Atomic global formulas include indistinguishabilities, but also atoms of the form $[\phi]$, which can be understood as $\phi \sim \text{true}$. The syntax above is open-ended, as more predicates will be used: in addition to custom predicates introduced by the user, *e.g.* to model verifications on messages, we will make use of several extra predicates later in this section. Note, however, that we will not use equality at the global level.

Indistinguishability atoms $\tilde{u} \sim \tilde{v}$ can only be formed when u_i and v_i have the same type for all i . Moreover, this type must be of order at most 1, *i.e.* of the form $\tau_1 \rightarrow \dots \rightarrow \tau_n$ where all τ_k are base atomic types. Further, we assume that elements of base types can be encoded as bitstrings, which allows interpreting these atoms as in Section 2.2, with a generalization: we now consider a distinguisher that can access the (semantics of) terms as oracles, which can be called on arbitrary inputs to obtain new data.

As said before, local formulas are just boolean terms. As such, the syntax above should be understood as syntactic sugar: for instance, $\phi \wedge \psi$ stands for $(\wedge) \phi \psi$ where (\wedge) is a constant function of type `bool` \rightarrow `bool` \rightarrow `bool`. We also view local quantifiers in this way: $\forall(x : \tau). \phi$ is a notation for $\forall_\tau (\lambda(x : \tau). \phi)$ where \forall_τ is a constant of type $(\tau \rightarrow \text{bool}) \rightarrow \text{bool}$. We require that all models interpret these logical constants as expected. In particular, $\llbracket \forall x. \phi \rrbracket_{\mathbb{M}}^{\eta, \rho} = 1$ iff $\llbracket \phi \rrbracket_{\mathbb{M}[x/a]}^{\eta, \rho} = 1$ for all $a \in \llbracket \tau \rrbracket_{\mathbb{M}}^{\eta}$.

Example 4.1. Consider once again the BASIC HASH protocol of Figure 3. Instead of expressing its security by using an equivalence, we can also state it with a local formula over the previously defined macros `inputr` and `outputr`:

$$\begin{aligned} \forall i : \text{index}. \text{happens}(\text{R}(i)) \wedge \text{output}_r @ \text{R}(i) = \text{true} \\ \Rightarrow \exists j : \text{index}. \text{T}(j) < \text{R}(i) \wedge \text{input}_r @ \text{R}(i) = \text{output}_r @ \text{T}(j) \end{aligned}$$

Asking that this formula holds with overwhelming probability properly expresses authentication for any trace: whenever some tag $\text{R}(i)$ accepts, its input value must have been produced in the past by an honest reader $\text{T}(j)$.

As usual in first-order logic, a global formula can be satisfied (or not) in a model: we write $\mathbb{M} \models F$ when this is the case. A global formula is valid when it is satisfied in all models. We write $F \models G$ if, for all \mathbb{M} , $\mathbb{M} \models F$ implies $\mathbb{M} \models G$. A local formula, however, is valid if it is overwhelmingly true in all models: ϕ is valid iff $[\phi]$ is valid, *i.e.* $\eta \mapsto \Pr_{\rho \in \mathbb{T}_{\mathbb{M}, \eta}} (\llbracket \phi \rrbracket_{\mathbb{M}}^{\eta, \rho} = 1)$ is overwhelming.

Example 4.2. We have $[\phi \Rightarrow \psi] \models [\phi] \tilde{\Rightarrow} [\psi]$: if both $\phi \Rightarrow \psi$ and ϕ are overwhelmingly true, then so must be ψ . However, $[\phi] \tilde{\Rightarrow} [\psi] \not\models [\phi \Rightarrow \psi]$: indeed, $[\phi] \tilde{\Rightarrow} [\psi]$ might be satisfied in a model where ϕ is not overwhelmingly true, but true for half of the tapes; hence ψ could be always false, and $[\phi \Rightarrow \psi]$ does not hold.

Example 4.3. We have $[\phi] \tilde{\vee} [\psi] \models [\phi \vee \psi]$ but not the converse: there might be models where only ϕ holds for half of the random tapes and only ψ holds for the other half, hence $\phi \vee \psi$ is overwhelmingly true (even exactly true) while neither ϕ nor ψ is.

In each model, global formulas are either fixed to true or false, but their quantifiers range over random variables, *i.e.* elements of $\mathbb{R}\mathbb{V}_{\mathbb{M}}(\tau)$ for some τ . In contrast, local formulas are probabilistic and interpreted in $\mathbb{R}\mathbb{V}_{\mathbb{M}}(\text{bool})$, but their quantifiers range over individuals, *i.e.* elements $a \in \llbracket \tau \rrbracket_{\mathbb{M}}^{\eta, \rho}$. Despite this essential difference, we can relate quantifiers at the two levels, as shown in [Baelde et al. 2023, Proposition 2].

PROPOSITION 4.4. *For all \mathbb{M} , we have $\mathbb{M} \models [\forall(x : \tau).\phi]$ iff $\mathbb{M} \models \tilde{\forall}(x : \tau). [\phi]$, and similarly for the existential quantifiers.*

4.3. Proof system for global and local reasoning

Squirrel’s proof system is given in a natural deduction style (though the tool’s tactics are sometimes closer to sequent calculus rules) and it deals with two kinds of sequents. Global sequents are of the form $\mathcal{E}; \Theta \vdash F$ where F is a global formula, Θ is a set of global formulas, and \mathcal{E} is a list of typed variable declarations containing at least all free variables of Θ, F . Such a sequent can be read as the global formula $\tilde{\forall}\mathcal{E}.(\tilde{\wedge}\Theta \Rightarrow F)$, and it is valid when the associated formula is valid. Local sequents are of the form $\mathcal{E}; \Theta; \Gamma \vdash \phi$ where ϕ is a local formula, Γ is a set of local formulas, and the other components are as before, with \mathcal{E} declaring all free variables of Θ, Γ and ϕ . The formula associated to a local sequent is $\tilde{\forall}\mathcal{E}.(\tilde{\wedge}\Theta \Rightarrow [\wedge\Gamma \Rightarrow \phi])$.

It is important to note, in the meaning of local sequents, that the second implication happens at the local level. As a result, the *global* hypotheses of Θ and the *local* hypotheses of Γ take a different meaning, illustrated in the following selected rules:

$$\frac{}{\mathcal{E}; \Theta, F \vdash F} \quad \frac{}{\mathcal{E}; \Theta; \Gamma, \phi \vdash \phi} \quad \frac{}{\mathcal{E}; \Theta, [\phi]; \Gamma \vdash \phi} \quad \frac{\mathcal{E}; \Theta; \emptyset \vdash \phi}{\mathcal{E}; \Theta \vdash [\phi]}$$

The first rule is the usual axiom rule for the global logic. The second and third rules are local and global versions of the axiom rule for the local logic. The last rule allows proving a global sequent whose conclusion is $[\phi]$ from the corresponding local sequent.

The local equality predicate gives rise to the ability to rewrite at the local level, which is in fact a generalization of the **REWRITE** rule from Figure 2:

$$\frac{\mathcal{E}; \Theta; \Gamma \vdash u = v \quad \mathcal{E}; \Theta; \Gamma \vdash \phi\{x \mapsto u\}}{\mathcal{E}; \Theta; \Gamma \vdash \phi\{x \mapsto v\}}$$

In the tool, a powerful **rewrite** tactic is provided, which builds on the previous rule as well as more basic rules, to perform various kinds of rewriting. The tactic can be invoked on a goal $\mathcal{E}; \Theta; \Gamma \vdash \phi$ to selectively replace occurrences of u into v inside Θ, Γ and ϕ when $u = v$ is part of the hypotheses (or can be simply derived from them). If $u = v$ is a local hypothesis, however, rewriting will only be possible in Γ and ϕ ; rewriting in Θ requires a global hypothesis $[u = v]$, and is only possible at specific occurrences in Θ .

The previous result relating global and local quantifiers justifies that our sequents only feature a top-level environment \mathcal{E} corresponding to a global quantification. The rules for universal quantifiers in our proof system are as follows, assuming that x does not occur in \mathcal{E} , and that t is a well-typed term of type τ in environment \mathcal{E} :

$$\frac{\mathcal{E}, x : \tau; \Theta \vdash F}{\mathcal{E}; \Theta \vdash \tilde{\forall}(x : \tau).F} \quad \frac{\mathcal{E}; \Theta \vdash \tilde{\forall}(x : \tau).F}{\mathcal{E}; \Theta \vdash F\{x \mapsto t\}} \quad \frac{\mathcal{E}, x : \tau; \Theta; \Gamma \vdash \phi}{\mathcal{E}; \Theta; \Gamma \vdash \forall(x : \tau).\phi} \quad \frac{\mathcal{E}; \Theta; \Gamma \vdash \forall(x : \tau).\phi}{\mathcal{E}; \Theta; \Gamma \vdash \phi\{x \mapsto t\}}$$

Let us mention one last articulation between the global and local logics:

$$\frac{\mathcal{E}; \Theta; \Gamma_1 \vdash \phi_1 \quad \mathcal{E}; \Theta \vdash (\wedge\Gamma_0 \Rightarrow \phi_0) \sim (\wedge\Gamma_1 \Rightarrow \phi_1)}{\mathcal{E}; \Theta; \Gamma_0 \vdash \phi_0}$$

In the tool, this rule is made available through a convenient `rewrite equiv` tactic, typically used to establish local formulas when they essentially derive from indistinguishability assumptions.

Example 4.5. A common assumption over hash functions, stronger than EUF-CMA, states that they can be seen as pseudo-random functions (PRF): for a fresh secret key sk , an attacker cannot distinguish between $x \mapsto h(x, sk)$ and a fully random function. It notably yields in our logic that $h(0, sk), h(1, sk) \sim n, m$. Applying on both sides of the equivalence the function $\lambda x, y : \text{message}. x \neq y$, we can obtain the validity of the indistinguishability formula $h(0, sk) \neq h(1, sk) \sim n \neq m$. With this global formula and the `rewrite equiv` tactic, we can reduce the proof of the local sequent $\mathcal{E}; \Theta; \cdot \vdash h(0, sk) \neq h(1, sk)$ to the proof of $\mathcal{E}; \Theta; \cdot \vdash n \neq m$.

A more complete description of our proof system is given in [Baelde et al. 2023]. However, it is important to point out that the system presented in that paper (or any earlier one) is only a collection of correct rules: it does not come with any completeness or cut-elimination result. Ongoing investigations are starting to provide such proof-theoretical results, though in a propositional fragment that allows to view our two-level logic in the richer framework of modal logic.

4.4. Cryptographic reasoning with recursive functions

It is natural to lift cryptographic reasoning to either local or global reasoning depending on the nature of the cryptographic assumption. In the case of EUF-CMA from Section 2.3, it could become a local rule of the form:

$$\text{EUF-CMA} \frac{\mathcal{E}; \Theta; \Gamma \vdash \bigwedge_{h(u, sk) \in S_{sk}^h(t, m)} (u \neq m)}{\mathcal{E}; \Theta; \Gamma \vdash h(m, sk) \neq t} \text{ sk only occurs as } h(x, sk) \text{ in } t \text{ and } m$$

Yet, now that terms contain recursive functions, the meaning of forbidding occurrences of sk is unclear: forbidden computations (e.g. a leak of a secret key) might occur deep inside the recursive functions. Further, sk may be an indexed key $k(i)$, and occurrences with distinct indices $k(j)$ for $j \neq i$ should be ignored. Checking for such indirect occurrences requires a deeper analysis of the terms involved.

Example 4.6. Let $k(i)$ be a name symbol indexed by an integer i and representing a key, and consider the following recursive function:

$$\text{keys} = \lambda(i : \text{index}). \text{if } (i = 0) \text{ then empty else } \langle k(i), \text{keys}(i - 1) \rangle$$

where we assume to have functions symbols for the standard operations over integers. Basically, $\text{keys}(i)$ computes the list $[k(i), k(i - 1), \dots, k(1)]$ (encoded as nested pairs). Consider some index terms i_0, j_0 and assume that $j_0 < i_0$. To prove $h(\text{att}(\text{keys}(j_0)), k(i_0)) \neq t$ with EUF-CMA, we need to check, among other things, that $k(i_0)$ does not occur in $\text{keys}(j_0)$. The term $\text{keys}(j_0)$ cannot be evaluated without providing a concrete value for j_0 . Moreover, if we attempted to directly check the condition on $\text{keys}(j_0)$'s body, we would find that it contains $k(i)$ (for a bound i), which can *a priori* be equal to $k(i_0)$. To conclude here, we need to exploit the hypothesis that $j_0 < i_0$ and details on how the recursive function keys is defined. Thus, we need a more involved check exploiting additional information that cannot be directly obtained from a superficial scrutiny of the terms.

Generalized subterms. We therefore introduce a generalized notion of subterm, to compute an over-approximation of all the subterms that may occur during the evaluation of a recursive function, for any model and random tapes. In addition, to allow exploiting semantic information (e.g. the fact that for any $k(i)$ in $\text{keys}(j)$, we have $i \neq 0$), we gather

together with each subterm a local formula expressing the condition under which that subterm may be computed. A generalized subterm, or *occurrence*, σ in a term t is a tuple (\mathcal{V}, ϕ, s) : s is, roughly, a subterm of t (expanding recursive definitions), \mathcal{V} is a set of typed variables bound above s in t , and ϕ is the condition under which s is evaluated in t , gathered from all if conditions. We use this notion, denoted $\text{ST}(t)$, in cryptographic rules to replace the previous conditions on syntactic subterms. For instance, a condition expressing that $k(i_0)$ does not occur in t now instead requires that for each $(\mathcal{V}, \phi, k(i)) \in \text{ST}(t)$, the formula $\forall \mathcal{V}. \phi \Rightarrow i_0 \neq i$ holds. Taking all this into account, the **EUF-CMA** rule finally becomes:

$$\text{EUF-CMA} \frac{\mathcal{E}; \Theta; \Gamma \vdash \left(\bigwedge_{(\mathcal{V}, \phi, k(i)) \in \text{ST}_h(t, m)} \forall \mathcal{V}. \phi \Rightarrow i_0 \neq i \right) \wedge \left(\bigwedge_{(\mathcal{V}, \phi, h(u, k(i))) \in \text{ST}(t, m)} \forall \mathcal{V}. (\phi \wedge i_0 = i) \Rightarrow u \neq m \right)}{\mathcal{E}; \Theta; \Gamma \vdash h(m, k(i_0)) \neq t}$$

where $\text{ST}_h(t, m)$ is $\text{ST}(t, m)$, where we excluded occurrences of k in subterms of the form $h(\cdot, k(\cdot))$. An additional premise (omitted here) checks that t and m are computable in polynomial time.

Making the verification effective. As presented here, the set of occurrences in a term may be infinite: thus, we cannot generate one premise for each occurrence. To alleviate this issue, we use heuristics that help Squirrel construct a proof obligation for the user that over-approximates this infinite set in a way that makes it expressible as a single formula. We illustrate how this can be done on an example.

Example 4.7. Continuing Example 4.6, consider arbitrary indices i_0, j_0 such that $j_0 < i_0$. Recall that $\text{keys}(j_0)$ computes the list of keys $[k(j_0), \dots, k(1)]$. In the style described above, a condition for a rule may require us to consider all elements $(\mathcal{V}, \phi, k(i))$ of $\text{ST}(\text{att}(\text{keys}(j_0)))$ (for any i) to show that i cannot be equal to i_0 . There are countably many such elements, all of the form $(\emptyset, j_0 \neq 0 \wedge j_0 - 1 \neq 0 \wedge \dots \wedge j_0 - \ell \neq 0, k(j_0 - \ell))$ for $\ell \in \mathbb{N}$. They can all be subsumed by using instead the occurrence $(i, 0 < i \wedge i \leq j_0, k(i))$. The corresponding proof obligation is $\forall i : \text{index}. 0 < i \wedge i \leq j_0 \Rightarrow i \neq i_0$, which can indeed be proved under the hypothesis that $j_0 < i_0$.

4.5. Reasoning over traces: induction and case analysis

We now clarify the interpretation of traces in this logic, and what kind of reasoning we can perform over them. Following the modeling approach outlined in Section 3, recall that we may use $A : \text{index} \rightarrow \text{timestamp}$ to represent an indexed action. Assuming for simplicity that we only have this action A , we restrict our attention to models \mathbb{M} where

$$\llbracket \text{timestamp} \rrbracket_{\mathbb{M}} = \{\text{init}, \text{undef}\} \cup \{A(n) \mid n \in N\}$$

where N is a finite subset of $\llbracket \text{index} \rrbracket_{\mathbb{M}}$ and undef is a special value used to interpret all timestamps that do not happen. Then, the constants init and A are interpreted naturally, independently of η, ρ :

$$\llbracket \text{init} \rrbracket_{\mathbb{M}}^{\eta, \rho} = \text{init} \quad \llbracket A \rrbracket_{\mathbb{M}}^{\eta, \rho}(n) = A(n) \text{ for } n \in N \quad \llbracket A \rrbracket_{\mathbb{M}}^{\eta, \rho}(n) = \text{undef} \text{ otherwise}$$

The following axiom is sound in this class of models:

$$\llbracket \forall x : \text{timestamp}. \text{happens}(x) \Rightarrow (x = \text{init} \vee \exists i : \text{index}. x = A(i)) \rrbracket$$

It can effectively be used to reason by case analysis on timestamps, but only at the level of local formulas. At the global level, it is tempting to postulate a stronger axiom:

$$\tilde{\forall} x : \text{timestamp}. \llbracket \text{happens}(x) \rrbracket \Rightarrow ([x = \text{init}] \tilde{\forall} \exists i : \text{index}. [x = A(i)])$$

However, this is too strong, and unsound in our class of models: consider for instance a random variable over timestamps that is init for some random tapes and some $A(n)$ for other tapes. To avoid this problem, we crucially need to quantify over random variables that are actually constant, *i.e.* which depend neither on η nor on ρ . Assuming a global predicate $\text{const}(\cdot)$ with exactly this meaning, we can now write the following, which is sound in our class of models:

$$\forall x : \text{timestamp}. \text{const}(x) \tilde{\wedge} [\text{happens}(x)] \Rightarrow ([x = \text{init}] \tilde{\vee} \exists i : \text{index}. [x = A(i)])$$

In practice, such axioms are used in Squirrel proofs through the **case** tactic. This tactic, when given a timestamp, will perform a case analysis for local or global goals. However, when the goal is a global sequent, it will generate an additional subgoal to verify that the timestamp is constant, corresponding to the hypothesis in our axiom above. In practice, the constancy assumptions come from the target security property. For instance, the equivalence of protocols modeled in Section 3.2 would more precisely be stated as $\tilde{\forall}\tau. \text{const}(\tau) \Rightarrow \text{frame}_i @\tau \sim \text{frame}_e @\tau$.

Case analysis over timestamps is not always sufficient to prove a goal, and induction is sometimes needed. At the local level, induction is supported by the following axiom, which is obviously valid in our intended class of models:

$$[\forall (p : \text{timestamp} \rightarrow \text{bool}). (\forall x. (\forall y. y < x \Rightarrow p y) \Rightarrow p x) \Rightarrow (\forall x. p x)]$$

At the global level, the lack of higher-order quantification forces us to resort to an axiom scheme and, more importantly, we need $\text{const}(\cdot)$ assumptions. For any global formula F , the following axiom is valid:

$$(\tilde{\forall}x. \text{const}(x) \Rightarrow (\tilde{\forall}y. \text{const}(y) \tilde{\wedge} [y < x] \Rightarrow F\{x \mapsto y\}) \Rightarrow F) \Rightarrow (\tilde{\forall}x. \text{const}(x) \Rightarrow F)$$

Without the $\text{const}(\cdot)$ restrictions, the axiom would be unsound: in a nutshell, the problem is that the well-founded relation $<$ over $\llbracket \text{timestamp} \rrbracket_{\mathbb{M}}$ does not lift to a well-founded relation $[\cdot < \cdot]$ over random variables. This can be generalized to arbitrary sets of actions, and to arbitrary well-founded types: one may for instance reason by induction over natural numbers or trees, along the same lines.

4.6. Bi-deduction & cryptographic reductions

Given the scope of the logic, the question of automating some proof steps in Squirrel was quickly raised. A key property of computational indistinguishability is that it is preserved by public computations: if $\vec{u} \sim \vec{v}$ holds then for any adversarially computable function f , we have $f(\vec{u}) \sim f(\vec{v})$. This property, introduced in [Baelde et al. 2022], is of particular interest for two reasons:

- it is the theoretical justification behind many cryptographic proof steps;
- it is simple, notably because it is compositional, as opposed to indistinguishability \sim which is not (see point 3 of Example 2.6).

For these reasons, this property turned-out to be a nice target for automation, reducing the proof-burden put on the user and allowing more high-level reasoning steps.

Bi-Deduction. Let $\#(\vec{v}_0; \vec{v}_1)$ be a pair of *sequences of terms*, representing the left and right scenarios of an equivalence. We say that it is *bi-deducible* from another pair $\#(\vec{u}_0; \vec{u}_1)$, which we write $\#(\vec{u}_0; \vec{u}_1) \triangleright \#(\vec{v}_0; \vec{v}_1)$, if there exists a *simulator* function S that can be efficiently computed by the adversary and such that $S(\llbracket \vec{u}_i \rrbracket) = \llbracket \vec{v}_i \rrbracket$ for any $i \in \{0; 1\}$. Crucially, the same simulator S is used for both $i = 0$ and 1 : S 's behavior cannot depend on i as it does not know which side it is interacting with. The link

between bi-deduction and indistinguishability is captured by the following rule:

$$\frac{\text{BiDEDUCE} \quad \#(\vec{u}_0; \vec{u}_1) \triangleright \#(\vec{v}_0; \vec{v}_1) \quad \vec{u}_0 \sim \vec{u}_1}{\vec{v}_0 \sim \vec{v}_1}$$

Indeed, any efficient adversary \mathcal{A} against the conclusion of this rule can be turned into an adversary \mathcal{B} against the premise by simply composing \mathcal{A} with the simulator S witnessing the bi-deduction premises. \mathcal{B} is an efficient adversary as both \mathcal{A} and S are, and it has the same advantage as \mathcal{A} 's.

The **BiDEDUCE** rule transforms the \forall quantification over adversary of an equivalence into an \exists quantification: instead of proving a property against any possible adversaries, of which there are infinitely many, it suffices to establish the existence of a program witnessing bi-deduction. Of course, this observation is not new, as this is the key idea behind all reduction-based arguments, whether it be in cryptography, computability or computational complexity. The main novelty is how we use bi-deduction as an automated entailment checker integrated in Squirrel tactics.

Rules. As already mentioned, bi-deduction is compositional, which allows decomposing bi-deductions as follows:

$$\frac{\text{BD.TUPLE} \quad \#(\vec{u}_0; \vec{u}_1) \triangleright \#(v_1; w_1) \quad \dots \quad \#(\vec{u}_0; \vec{u}_1) \triangleright \#(v_n; w_n)}{\#(\vec{u}_0; \vec{u}_1) \triangleright \#(v_1, \dots, v_n; w_1, \dots, w_n)}$$

Here are two selected (simple) bi-deduction rules:

$$\frac{\text{BD.FA} \quad \#(\vec{u}_0; \vec{u}_1) \triangleright \#(\vec{v}_0; \vec{v}_1)}{\#(\vec{u}_0; \vec{u}_1) \triangleright \#(f(\vec{v}_0; \vec{v}_1))} \quad \frac{\text{BD.TRANS} \quad \#(\vec{u}_0; \vec{u}_1) \triangleright \#(\vec{w}_0; \vec{w}_1) \quad \#(\vec{u}_0, \vec{w}_0; \vec{u}_1, \vec{w}_1) \triangleright \#(\vec{v}_0; \vec{v}_1)}{\#(\vec{u}_0; \vec{u}_1) \triangleright \#(\vec{v}_0; \vec{v}_1)}$$

where common behaviors between the left and right scenarios are factorized when possible: *e.g.* $f(\#(\vec{v}_0; \vec{v}_1))$ denotes $\#(f(\vec{v}_0); f(\vec{v}_1))$. The **BD.FA** rule lets the adversary simulate the computation f : indeed, if S computes $\#(\vec{v}_0; \vec{v}_1)$ from $\#(\vec{u}_0; \vec{u}_1)$, then $f \circ S$ computes $\#(f(\vec{v}_0); f(\vec{v}_1))$ from the same inputs. The transitivity rule **BD.TRANS** allows extending the input $\#(\vec{u}_0; \vec{u}_1)$ of a simulator trying to compute $\#(\vec{v}_0; \vec{v}_1)$ by any sequence $\#(\vec{w}_0; \vec{w}_1)$ that is bi-deducible from $\#(\vec{u}_0; \vec{u}_1)$.

The more interesting bi-deduction rules deal with functions (seen as oracles) and conditioned terms, for which we now present two simplified rules. Conditioned terms are written $(t \mid \phi)$, which denotes the pair of terms $(\phi, \text{if } \phi \text{ then } t \text{ else witness})$, where witness is a function symbol of the same type as t .

$$\frac{\text{BD.ORACLE} \quad \begin{array}{c} \triangleright \#(t_0; t_1) \\ \vdash \left[\begin{array}{l} w_0\{y \mapsto t_0\} = v_0 \wedge \\ w_1\{y \mapsto t_1\} = v_1 \end{array} \right] \end{array}}{\lambda y. \#(w_0; w_1) \triangleright \#(v_0; v_1)} \quad \frac{\text{BD.COND} \quad \begin{array}{c} \triangleright \#(\psi_0; \psi_1) \\ \vdash \left[\begin{array}{l} \psi_0 \Rightarrow (\phi_0 \wedge v_0 = w_0) \wedge \\ \psi_1 \Rightarrow (\phi_1 \wedge v_1 = w_1) \end{array} \right] \end{array}}{\#((w_0 \mid \phi_0); (w_1 \mid \phi_1)) \triangleright \#((v_0 \mid \psi_0); (v_1 \mid \psi_1))}$$

The **BD.ORACLE** allows the simulator to call an oracle $\lambda y. \#(w_0; w_1)$ it received as input on any arguments $\#(t_0; t_1)$ that it can compute; the obtained terms $\#(w_0\{y \mapsto t_0\}; w_1\{y \mapsto t_1\})$ must be established, in the second proof obligation, to be equal to the wanted terms $\#(v_0; v_1)$. In **BD.COND**, the simulator retrieves from its inputs a conditioned term: the bi-deduction premise $\triangleright \#(\psi_0; \psi_1)$ guarantees that the simulator knows when to extract the terms from its input; and the logical premise ensures that when it does (*i.e.* when

ψ_0 or ψ_1 holds), the extracted value $\sharp(w_0; w_1)$ is given as inputs (*i.e.* when ϕ_0 or ϕ_1 holds) and is equal to the wanted value $\sharp(v_0; v_1)$.

Implementation. We implemented a fully automated heuristic procedure for bi-deduction and integrated it in the `apply` tactic of Squirrel: concretely, it is used to automatically check the entailment between equivalences, *e.g.* $\vec{u}_0 \sim \vec{u}_1 \Rightarrow \vec{v}_0 \sim \vec{v}_1$. While this procedure is heuristic, and thus incomplete, it remains a useful tool in practice.

Advanced bi-deduction. Bi-deduction gives a formal framework to establish the existence of a simulator S reducing an equivalence $\vec{u}_0 \sim \vec{u}_1$ to another equivalence $\vec{v}_0 \sim \vec{v}_1$. This looks very similar to the problem of applying a cryptographic assumption expressed as a pair of games $\mathcal{G}_0 \approx \mathcal{G}_1$ (in the style of Section 2): for example, given a target equivalence $\vec{u}_0 \sim \vec{u}_1$, does there exist S such that $S^{\mathcal{G}_i}$ computes \vec{u}_i for all $i \in \{0; 1\}$? Extending bi-deduction to be able to handle cryptographic games would allow supporting new cryptographic assumptions in Squirrel’s logic and implementation without having to manually design rules as we did in Section 2.3 and Section 4.4. This would drastically reduce the theoretical and programming work needed to support a new rule while limiting the risk of mistakes. Furthermore, it would allow non-expert users to add new cryptographic assumptions in Squirrel simply by writing, in user-space, the pair of games $\mathcal{G}_0 \approx \mathcal{G}_1$ associated to the assumptions.

Unfortunately, the simulators covered by our previous notion of bi-deduction are too restricted for this. First, they can only interact with *stateless* oracles that are provided as first-order terms of the logic. Second, they are deterministic programs: while they may receive randomly sampled values from their inputs, they cannot do random samplings themselves. This is too limiting: oracles of cryptographic games are often stateful (*e.g.* EUF-CMA relies on a stateful log \mathcal{L}); and the simulators justifying Squirrel’s cryptographic rules must sample all the randomness that does not come from the cryptographic game. Recently, we designed and implemented in Squirrel an extended notion of bi-deduction supporting stateful oracles and probabilistic simulators [Baelde et al. 2024]. This introduces a number of interesting aspects: we equipped the bi-deduction predicate with Hoare-style pre- and post-conditions to handle game statefulness; and random sampling usage is tracked through *randomness constraints* that ensure that the constructed simulator does not directly access the game’s random samplings. This approach, which showed good initial results, may become the standard way of handling cryptographic assumptions in Squirrel.

5. CONCLUSION

We have developed a logic based on the CCSA logic described in [Bana and Comon-Lundh 2014], along with proof systems for verifying reachability and equivalence properties within that framework. Our work demonstrates that this approach offers a straightforward, high-level strategy for conducting computer-assisted proofs of cryptographic protocols, offering asymptotic security guarantees within the computational model. This is reinforced by the development of the interactive prover Squirrel and its application across different case studies [Baelde et al. 2021; Baelde et al. 2022; Cremers et al. 2022; Baelde et al. 2024], summarized in Figure 4. The number of LoC mentioned includes both the model and the proof script. The cryptographic assumptions on which the proof relies are also indicated, as well as the security properties under study. We analyze reachability properties, namely various forms of authentication, and some equivalence-based properties, *e.g.* unlinkability on several RFID protocols, an anonymity property, as well as a strong form of secrecy. The case studies are divided into two blocks, with non-stateful protocols first and then stateful ones. The Squirrel prover was extended to post-quantum security proofs in [Cremers et al. 2022], the corresponding protocols with security guarantees against a quantum attacker are indicated with a \star .

Protocol name	LoC	Assumptions	Security Properties
Basic Hash *	100	PRF, EUF	auth. & unlinkability
Hash Lock *	130	PRF, EUF	auth. & unlinkability
LAK (with pairs) *	250	PRF, EUF	auth. & unlinkability
MW *	300	PRF, EUF, XOR	auth. & unlinkability
Feldhofer *	270	ENC-KP, INT-CTXT	auth. & unlinkability
Private authentication *	100	CCA ₁ , ENC-KP	anonymity
Private authentication	100	CCA\$	anonymity
Signed DDH [ISO 9798-3]	240	EUF, DDH	auth. & strong secrecy
IKEV1 _{PSK} *	850	PRF	auth. & strong secrecy
IKEV2 ^{SIGN} _{PSK} *	300	PRF, EUF	auth. & strong secrecy
BCGNP *	355	PRF, CCA ₁ , XOR	strong secrecy
FSXY *	620	PRF, CCA ₁ , XOR	strong secrecy
SC-AKE *	745	PRF, EUF, CCA ₁ , XOR	auth. & strong secrecy
CANAuth	450	EUF	auth.
SLK06	80	EUF	auth.
YPLRK05	160	EUF	auth.
OSK (without reader)	320	ROM	strong secrecy
OSK	370	EUF	auth.
YubiKey	270	INT-CTXT	auth.
YubiHSM	720	CCA ₁ , INT-CTXT	auth.

Fig. 4. Some results obtained with the Squirrel prover on various protocols.

Ongoing and Future Work. As explained in Section 3, the Squirrel tool allows cryptographic protocols to be described in a process algebra, whereas the translation from protocols to the logic relies on the definition of mutually recursive functions modeling the protocol observables depending on the execution trace. To ease formal proofs, this translation does not follow the granularity of elementary execution steps in the process algebra. Instead, it groups together elementary steps in blocks following specific patterns. Whereas the soundness of the translation is easy to obtain for a simple class of protocols, some works remains to be done for more complex protocols.

We also aim to improve the automation of the prover. One option is to discharge parts of the proof that do not rely on any cryptographic assumption to SMT solvers. Another option is to rely on type systems: techniques based on typing have mainly been used in symbolic models, but we are currently working to develop such a system to obtain computational guarantees in the CCSA framework.

We aim to extend the scope of the post-quantum work initiated in [Cremers et al. 2022] and to analyze hybrid post-quantum protocols that rely on a combination of asymmetric post-quantum algorithms with well-known and well-studied pre-quantum asymmetric cryptography, *e.g.* based on the discrete logarithm problem.

Lastly, as mentioned in Section 3, the current notion of security in Squirrel differs from the standard one in the computational model, as the number of sessions we consider is arbitrary but fixed, in the sense that it does not depend on the security parameter. We plan to overcome this limitation, notably by leveraging the composition result outlined in [Comon et al. 2020]. Furthermore, we intend to expand on our methodology to go beyond asymptotic guarantees and provide concrete security bounds.

ACKNOWLEDGMENTS

This work received funding from the France 2030 program managed by the French National Research Agency under grant agreement No. ANR-22-PECY-0006.

REFERENCES

- David Baelde, Stéphanie Delaune, Charlie Jacomme, Adrien Koutsos, and Solène Moreau. 2021. An Interactive Prover for Protocol Verification in the Computational Model. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*. IEEE, 537–554.
- David Baelde, Stéphanie Delaune, Adrien Koutsos, and Solène Moreau. 2022. Cracking the Stateful Nut: Computational Proofs of Stateful Security Protocols using the Squirrel Proof Assistant. In *35th IEEE Computer Security Foundations Symposium, CSF 2022, Haifa, Israel, August 7-10, 2022*. IEEE, 289–304.
- David Baelde, Adrien Koutsos, and Joseph Lallemand. 2023. A Higher-Order Indistinguishability Logic for Cryptographic Reasoning. In *38th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2023*. 1–13.
- David Baelde, Adrien Koutsos, and Justine Sauvage. 2024. Cryptographic Reductions By Bi-Deduction. (March 2024). <https://hal.science/hal-04511718>
- Gergei Bana and Hubert Comon-Lundh. 2012. Towards Unconditional Soundness: Computationally Complete Symbolic Attacker. In *Principles of Security and Trust - First International Conference, POST 2012, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2012, Tallinn, Estonia, March 24 - April 1, 2012, Proceedings (Lecture Notes in Computer Science)*, Pierpaolo Degano and Joshua D. Guttman (Eds.), Vol. 7215. Springer, 189–208.
- Gergei Bana and Hubert Comon-Lundh. 2014. A Computationally Complete Symbolic Attacker for Equivalence Properties. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security, Scottsdale, AZ, USA, November 3-7, 2014*. ACM, 609–620.
- Manuel Barbosa, Gilles Barthe, Karthik Bhargavan, Bruno Blanchet, Cas Cremers, Kevin Liao, and Bryan Parno. 2021. SoK: Computer-Aided Cryptography. In *42nd IEEE Symposium on Security and Privacy, SP 2021, San Francisco, CA, USA, 24-27 May 2021*. IEEE, 777–795.
- Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. 2009. Formal certification of code-based cryptographic proofs. In *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009, Savannah, GA, USA, January 21-23, 2009*, Zhong Shao and Benjamin C. Pierce (Eds.). ACM, 90–101.
- Gilles Barthe, Benjamin Grégoire, Sylvain Heraud, and Santiago Zanella Béguelin. 2011. Computer-Aided Security Proofs for the Working Cryptographer. In *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings (Lecture Notes in Computer Science)*, Phillip Rogaway (Ed.), Vol. 6841. Springer, 71–90.
- David A. Basin, Andreas Lochbihler, and S. Reza Sefidgar. 2020. CryptHOL: Game-Based Proofs in Higher-Order Logic. *J. Cryptol.* 33, 2 (2020), 494–566.
- Bruno Blanchet. 2001. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In *14th IEEE Computer Security Foundations Workshop (CSFW-14 2001), 11-13 June 2001, Cape Breton, Nova Scotia, Canada*. IEEE Computer Society, 82–96.
- Bruno Blanchet. 2008. A Computationally Sound Mechanized Prover for Security Protocols. *IEEE Trans. Dependable Secur. Comput.* 5, 4 (2008), 193–207.
- Mayla Brusò, Konstantinos Chatzikokolakis, and Jerry den Hartog. 2010. Formal Verification of Privacy for RFID Systems. In *23rd IEEE Computer Security Foundations Symposium, CSF 2010, Edinburgh, United Kingdom, July 17-19, 2010*. IEEE Computer Society, 75–88.
- Hubert Comon, Charlie Jacomme, and Guillaume Scerri. 2020. Oracle Simulation: A Technique for Protocol Composition with Long Term Shared Secrets. In *2020 ACM SIGSAC Conference on Computer and Communications Security, CCS 2020*. ACM, 1427–1444.
- Cas Cremers, Caroline Fontaine, and Charlie Jacomme. 2022. A Logic and an Interactive Prover for the Computational Post-Quantum Security of Protocols. In *43rd IEEE Symposium on Security and Privacy, SP 2022, San Francisco, CA, USA, May 22-26, 2022*. IEEE, 125–141.
- Danny Dolev and Andrew Chi-Chih Yao. 1981. On the Security of Public Key Protocols (Extended Abstract). In *22nd Annual Symposium on Foundations of Computer Science, Nashville, Tennessee, USA, 28-30 October 1981*. IEEE Computer Society, 350–357.
- Simon Meier, Benedikt Schmidt, Cas Cremers, and David A. Basin. 2013. The TAMARIN Prover for the Symbolic Analysis of Security Protocols. In *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings (Lecture Notes in Computer Science)*, Natasha Sharygina and Helmut Veith (Eds.), Vol. 8044. Springer, 696–701.
- Adam Petcher and Greg Morrisett. 2015. The Foundational Cryptography Framework. In *Principles of Security and Trust - 4th International Conference, POST 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015, Proceedings (Lecture Notes in Computer Science)*, Riccardo Focardi and Andrew C. Myers (Eds.), Vol. 9036. Springer, 53–72.

Nikhil Swamy, Catalin Hritcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Simon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss, Jean Karim Zinzindohoue, and Santiago Zanella Béguelin. 2016. Dependent types and multi-monadic effects in F. In *43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016*, Rastislav Bodík and Rupak Majumdar (Eds.). ACM, 256–270.