

A Probabilistic Logic for Concrete Security

David Baelde ^{*}, Caroline Fontaine [‡], Adrien Koutsos [†], Guillaume Scerri [‡], Théo Vignon [‡]

^{*} Univ Rennes, CNRS, IRISA, France

[†] Inria Paris, France

[‡] Université Paris-Saclay, CNRS, ENS Paris-Saclay, LMF, Gif-sur-Yvette, France

Index Terms—Logic, Concrete Security, Crypto Protocols

Abstract—The Squirrel Prover is a proof assistant designed for the computational verification of cryptographic protocols. It implements a probabilistic logic that captures cryptographic and probabilistic arguments used in security proofs. This logic operates in the asymptotic security setting, which limits the expressiveness of formulas and proofs. As a consequence, it can only prove security for finite interactions with a protocol, falling outside of the polynomial-level of security usually expected by cryptographers. We lift all these limitations by moving to a concrete security setting. We extend the logic with concrete security predicates, and design a corresponding proof system. We show the usefulness of these extensions on a case study, and through a novel proof-transformation result which shows that a large class of asymptotic logic security proofs can be automatically rewritten into concrete logic security proofs, improving security bounds exponentially.

I. INTRODUCTION

Cryptographic protocols are crucial to get secure communications, e.g. for online payment or messaging. Strong guarantees on their security can be provided through cryptographic proofs, *via* formal mathematical analyses of the protocols and the targeted security properties. Unfortunately, properly designing a protocol remains challenging [1], [2], as cryptographic proofs can be involved and error-prone [3], [4], because of the complexity of the protocol and the intricacies of the cryptographic arguments needed. These issues have led to the development of verification tools allowing for the *mechanization* of cryptographic proofs (see [5] for a survey).

In more details, the principle of a cryptographic proof is to show that no adversary — assumed to be an arbitrary probabilistic polynomial Turing machine (PPTM) — can break the protocol security. Pen-and-paper proofs usually proceed as follows. First, the security of the protocol is expressed as a *game* between the adversary and a *challenger*, modeling the security of the protocol. Then, this game is iteratively modified by a sequence of game hops [6], where each hop is justified by a cryptographic reduction or some other probabilistic argument. Finally, the proof concludes when the prover manages to obtain a game in which security trivially holds.

Example 1. *The Private Authentication (PA) protocol [7] is a two-message protocol in which agents A and B attempt to authenticate each other and establish a shared session key. It aims to ensure privacy in the sense that an outside observer cannot tell whether B accepts to communicate with A. To do*

so, A sends to B the message $\text{enc}(\langle \text{pk}_A, n_A \rangle, \text{pk}_B, r_A)$, i.e. a randomized asymmetric encryption of A’s identity (represented by its public key) and of a fresh random nonce n_A of length η , under the public key of B — r_A is the randomness needed by the randomized asymmetric encryption enc . When he receives this message, B decrypts it and replies with $\text{enc}(\langle n_B, n_A \rangle, \text{pk}_A, r_B)$. If the received message is not valid, B replies with $\text{enc}(\langle n_B, 0^n \rangle, \text{pk}_A, r_B)$. Notice that B’s two possible answers are encryptions of same-length plaintexts, and will thus be indistinguishable for an outside observer. In this example, we want to prove that the adversary cannot know if B wanted to talk to A or not. This property holds thanks to B’s second possible answer, which is a decoy message sent when B doesn’t accept the communication with A. Note that this is independent from proving that A is able to know whether B has accepted her message or not (which is actually the case here, up to negligible probability).

Several techniques have been developed to mechanize the analysis of such cryptographic arguments. The CryptoVerif [8] tool proceeds by the automatic application of game transformations. While this tool is highly automated on simple examples, more complex protocols often require heavy user guidance. It must also be noted that CryptoVerif does not support generic mathematical reasoning: when such arguments are needed, they must be assumed in the tool through axioms, and externally proven. Probabilistic Relational Hoare Logics [9] (pRHL), upon which several tools are based (e.g. EasyCrypt [10], CryptHOL [11], and SSProve [12]), encodes games as imperative programs, and allows to express cryptographic reductions as relational properties of these programs. This is a very expressive logic, but this comes at a cost: reasoning is done at a relatively low level, which can lead to long and tedious proof developments.

In this paper, we are interested in another approach for cryptographic protocol verification called the Computationally Complete Symbolic Attacker (CCSA) model. This approach, initially introduced in [13], is based on a logic with a probabilistic semantics that can be used to encode a protocol security as an indistinguishability formula $\vec{u} \sim \vec{v}$ where, essentially, \vec{u} and \vec{v} model messages exchanged over the network (typically, \vec{u} corresponds to an execution of the protocol studied, while \vec{v} is for an idealized version of this protocol). Then, the formula $\vec{u} \sim \vec{v}$ can be shown valid using reasoning rules, therefore proving that the protocol is secure. The CCSA logic — modified to internalize the notion of

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

protocol execution and of inductive reasoning — has been implemented into a proof assistant called Squirrel [14]. Since then, Squirrel and its logic have been extended and used several times, to support and analyze stateful protocols [15], to be adapted to a post-quantum cryptography setting [16], and to support higher-order reasoning [17]. As the logic of [17] is the most expressive one, we will rely on it in this paper.

Example 2. We consider a simple scenario with two agents A and B willing to talk to each other, where each agent plays an arbitrary number of sessions. Our encoding uses some mutually recursive functions: `output` (X,N) is the output of agent $X \in \{A, B\}$ for the N-th interaction with the adversary; `choose` N lets the adversary decide whether the N-th interaction is with A or B; `input` N is the N-th input sent by the adversary; and `frame` N is the sequence of the first N outputs of the protocol, extended with the public values pk_A and pk_B .

The N-th input `input` N is any value that can be computed by the adversary using its current knowledge, i.e. the sequence `frame` (N-1) of the first N-1 outputs of the protocol. We model the adversary's computation using the adversarial function `atti`, and define `choose` (N-1) using the adversarial function `attc`:

`input` N $\stackrel{\text{def}}{=} \text{att}_i(\text{frame } (N-1))$ `choose` N $\stackrel{\text{def}}{=} \text{att}_c(\text{frame } (N-1))$

The sequence of outputs `frame` N is easily defined as:

`frame` N $\stackrel{\text{def}}{=} \begin{cases} \langle \text{frame } (N-1), \text{output } (\text{choose } N, N) \rangle & \text{if } N \geq 0 \\ pk_A, pk_B & \text{if } N = 0 \end{cases}$

To model the freshness of the random nonce n_A and encryption randomness r_A in A's message, we index these values by the interaction number N. We do the same for B.

`output` (A,N) $\stackrel{\text{def}}{=} \text{enc}(\langle pk_A, n_A N \rangle, pk_B, r_A N)$
`output` (B,N) $\stackrel{\text{def}}{=} \text{fst}(\text{dec}(\text{input } N, sk_B)) = pk_A \ \&\&$
 $\text{len}(\text{snd}(\text{dec}(\text{input } N, sk_B))) = \eta$
then $\text{enc}(\langle n_B N, \text{snd}(\text{dec}(\text{input } N, sk_B)) \rangle, pk_A, r_B N)$
else $\text{enc}(\langle n_B N, 0^\eta \rangle, pk_A, r_B N)$

Here, `fst` and `snd` denote the first and second projections w.r.t. $\langle \cdot, \cdot \rangle$, and `len` denotes the length function.

The CCSA approach operates in the *asymptotic security* model, which considers that a protocol is secure if the probability that an adversary breaks it is a negligible function of the *security parameter* η — where η can be, e.g., the length of the keys, and a function is negligible if it is asymptotically smaller than the inverse of any polynomial. In particular, predicate $\vec{u} \sim \vec{v}$ states that the advantage of any PTIME adversary in distinguishing \vec{u} from \vec{v} is negligible in η .

Example 3. Following our running example of the PA protocol, we can express the fact that it preserves user A's privacy with the unlinkability property [18] `frame` N \sim `frameid` N which states that no adversary can distinguish an execution of the protocol from an idealization of the protocol where, each time A sends a new message, we change its identity by using a different

public key $pk'_A N$ in each output. This idealized protocol is defined as the normal protocol, except for A's output which is replaced by: `outputid` (A,N) $\stackrel{\text{def}}{=} \text{enc}(\langle pk'_A N, n_A N \rangle, pk_B, r_A N)$. The identity verification in B remains the same, but it may not pass anymore even when an output of A is forwarded to B: pk_A may or may not be equal to $pk'_A N$ depending on N.

If the encryption function is IND-CCAI, it can be shown that the PA protocol is unlinkable for any constant number of sessions. More precisely, we can prove by induction over N:

$$\forall N. \text{const}(N) \Rightarrow \text{frame } N \sim \text{frame}_{id} N \quad (1)$$

a) *Limitations:* The asymptotic model allows for simple, high-level CCSA logic that completely hides probabilities and security parameter from the user. However, this limits the logic in several ways:

- 1) *Expressivity:* precise security bounds cannot be expressed by the logic, limiting the practical applicability of the logic [19]. E.g., the logic cannot be used to determine what should be the size the keys in a concrete application.
- 2) *Reasoning:* some cryptographic arguments cannot be captured in their full generality by the logic. For example, the most general version of the hybrid argument [20], [21] states that $t_0 \sim t_{P(\eta)}$ (where P is a polynomial) if it can be proved that, for every $i < P(\eta)$, we have $t_i \sim t_{i+1}$ with a uniform advantage w.r.t. the security parameter. Since the asymptotic CCSA logic cannot express precise security bounds, the uniformity conditions cannot be checked, putting such arguments out-of-reach.
- 3) *Security guarantees:* as a consequence of 2), the adversary must usually be restricted to a constant (i.e. independent from η) number of interactions with the protocol. For example, this is the case for the formula in Eq. (1) of Example 3, as it is proved by induction over N — induction is essentially an hybrid argument. Said otherwise, asymptotic CCSA logic can usually only be used to prove that a protocol provides a *parametric* level of security, instead of the stronger *polynomial* level of security expected by cryptographers.

b) *Contributions:* We lift all these limitations by moving to the concrete security setting. As a first contribution, we extend the asymptotic CCSA logic of [17] with concrete security variants of the security predicates: e.g., the new predicate $u \sim_\varepsilon v$ states that the distinguishing advantage of any adversary against $u \sim v$ is at most ε . Then, we design a new proof system for these new concrete security predicates, and show its usefulness through a case study.

Unfortunately, translating idiomatic asymptotic CCSA proofs in the concrete logic yields advantage upper-bounds which are exponential in the number of interactions with the adversary. Because of this, it could be feared that the idiomatic CCSA proof strategy — and all existing proofs relying on it — should be abandoned. We argue that this is not the case through our second contribution, a novel theoretical result which shows that a large class of (asymptotic) CCSA proofs can be automatically rewritten into concrete security proofs with an optimized

advantage which reaches the sought-after polynomial level of security. Crucially, this result is not a simple rule-by-rule translation, but involves a whole-proof transformation through non-trivial rule commutations. Finally, we show a representative practical example proof which falls in our class of proofs, and can thus be transformed to obtain polynomial security.

c) *Outline:* We describe our CCSA logic for concrete security in Section II, and its proof system in Section III. We illustrate this system on the Private Authentication protocol in Section IV. We then present in Section V our proof-transformation result, allowing to derive polynomial security from any proof in a significant fragment of our system.

II. LOGIC

We now present our extension of the probabilistic logic for cryptographic reasoning of [17] to a concrete security setting.

Terms of the logic, which are the same as in [17] (extended with `let` in constructor), are simply-typed higher-order terms with special symbols (called *names*) used to denote random samplings. A key feature of the logic is that all terms are interpreted using the same set of random bits, from which they retrieve exactly the randomness they need. This allows to track probabilistic dependencies between terms.

Formulas of [17] are first-order formulas built on top of a set of predicates capturing probabilistic and computational properties of terms. Mainly, our new logic extends the formulas of [17] with two new concrete security predicates: $[\phi]_\varepsilon$ states that the probability that ϕ does not hold is at most ε , and $t_0 \sim_\varepsilon t_1$ states that the probability that an adversary distinguishes t_0 from t_1 is at most ε .

A. Terms

We first recall the types and terms of our logic.

a) *Types:* We consider simple types, noted τ , generated from base types $\tau_b \in \mathbb{T}$ using the arrow construct $\cdot \rightarrow \cdot$. The set of *base types* \mathbb{T} must contain at least the types `bool`, `message`, `int`, `int` (modeling the set $\mathbb{N} \cup \{+\infty\}$), `bint` (modeling the set of integers $[0, N_\eta]$ where N_η is an arbitrary integer fixed by the model), and `real` (modeling the set $\mathbb{R} \cup \{-\infty, +\infty\}$).

We identify a subset $\mathbb{B} \subseteq \mathbb{T}$ of *bit-string encodable base types*, that contains at least `bool`, `message`, `bint`, `int`, and `int`. We say that a type has *order* 0 when it belongs to \mathbb{B} , and that it has order $n + 1$ when it is of the form $\tau \rightarrow \tau'$ where τ has order at most n and τ' has order at most $n + 1$. For example, `intk → message = int → ... → int → message` has order 1.

The semantics of types is described by a *type structure* \mathbb{M} which assigns to each base type $\tau_b \in \mathbb{T}$ and each value of the security parameter $\eta \in \mathbb{N}$ an interpretation domain $[\tau_b]_{\mathbb{M}}^\eta$. The interpretation of a type in \mathbb{B} must be a subset of $\{0, 1\}^*$. We force the interpretation of standard types to be the expected one, e.g. $[\text{int}]_{\mathbb{M}}^\eta$ is the set of (bit-string encodings of) integers \mathbb{N} and $[\text{message}]_{\mathbb{M}}^\eta = \{0, 1\}^*$ (for every η). Arbitrary types are then interpreted by defining $[\tau_0 \rightarrow \tau_1]_{\mathbb{M}}^\eta = [\tau_0]_{\mathbb{M}}^\eta \rightarrow [\tau_1]_{\mathbb{M}}^\eta$. We say that a type is *finite* if its interpretation is finite for every η . For example, the type `bint` is finite since for every $\eta \in \mathbb{N}$, it contains $N_\eta + 1$ elements for some arbitrary integer N_η fixed by the model.

b) *Terms:* Our *terms* are simply-typed λ -terms built upon a set of variables \mathcal{X} :

$$t ::= x \mid t t \mid \lambda(x : \tau).t \mid \forall(x : \tau).t \mid \text{let } (x : \tau) = t \text{ in } t$$

where $x \in \mathcal{X}$. Terms are considered modulo α -renaming, and we let $\text{fv}(t)$ be the free variables of t . A variable x in \mathcal{X} can be used to denote a function argument coming from a λ binder, a logical variable quantified by a \forall , but also a function symbol (e.g. integer addition `+`). We write $\lambda x_1, \dots, x_n. t$ for $\lambda x_1. \dots \lambda x_n. t$, and $t \vec{u}$ stands for $((t u_1) \dots u_n)$ when $\vec{u} = u_1, \dots, u_n$.

An *environment* \mathbb{E} is a finite sequence of variable *declarations* $(x : \tau)$ and *definitions* $(x : \tau = t)$. A declared or defined variable is said to be bound in \mathbb{E} , and we require that no environment declares a variable twice. We consider the same standard type system as [17], and we write $\mathbb{E} \vdash t : \tau$ (resp. $\vdash \mathbb{E}$) if t has type τ in \mathbb{E} (resp. if \mathbb{E} is well-typed). As usual, we require that terms and environments are well-typed.

We only consider environments containing at least the declarations of a number of standard functions such as Boolean connectives (e.g. $\wedge, \vee, \rightarrow$), integer operations ($+, \times, \dots$), etc.

Environments support (mutually) recursive definitions. In our examples, this is used when defining *frame* \mathbb{N} . We refer the reader to [17] for a presentation of the well-foundedness conditions guarding recursive definitions.

A *term structure* \mathbb{M} for \mathbb{E} is a type structure extended with:

- a set $\mathbb{T}_{\mathbb{M}, \eta} = \mathbb{T}_{\mathbb{M}, \eta}^a \times \mathbb{T}_{\mathbb{M}, \eta}^h$ where $\mathbb{T}_{\mathbb{M}, \eta}^a$ (resp. $\mathbb{T}_{\mathbb{M}, \eta}^h$) is the finite set of all random tapes of a given length. For example, $\mathbb{T}_{\mathbb{M}, \eta}^a$ can be the set $\{0, 1\}^{P(\eta)}$ of all random tapes of length $P(\eta)$ for some polynomial P . Tapes in $\mathbb{T}_{\mathbb{M}, \eta}^a$ are used for the *adversarial* randomness, while tapes in $\mathbb{T}_{\mathbb{M}, \eta}^h$ are used for *honest* randomness (e.g. for names).
- for any defined or declared variable x in \mathbb{E} of type τ , \mathbb{M} defines its interpretation $\mathbb{M}(x) \in \mathbb{R}\mathbb{V}_{\mathbb{M}}(\tau)$ where $\mathbb{R}\mathbb{V}_{\mathbb{M}}(\tau)$ is the set of η -indexed random variables from $\mathbb{T}_{\mathbb{M}, \eta}$ to the sampling space $[\tau]_{\mathbb{M}}^\eta$.

Variables interpretation is lifted to term interpretation $[\cdot]_{\mathbb{M}}^{\eta, \rho}$:

$$\begin{aligned} [x]_{\mathbb{M}}^{\eta, \rho} &\stackrel{\text{def}}{=} \mathbb{M}(x)(\eta)(\rho) & [t t']_{\mathbb{M}}^{\eta, \rho} &\stackrel{\text{def}}{=} [t]_{\mathbb{M}}^{\eta, \rho} ([t']_{\mathbb{M}}^{\eta, \rho}) \\ [\lambda(x : \tau_0).t]_{\mathbb{M}}^{\eta, \rho} &\stackrel{\text{def}}{=} \begin{cases} [\tau_0]_{\mathbb{M}}^\eta \rightarrow [\tau]_{\mathbb{M}}^\eta \\ a \mapsto [t]_{\mathbb{M}[x \mapsto 1_a]}^{\eta, \rho} \end{cases} \end{aligned}$$

where, in the last case, t is of type τ and $1_a^\eta \in \mathbb{R}\mathbb{V}_{\mathbb{M}}(\tau)$ is a random variable such that $1_a^\eta(\eta)(\rho) = a$ for every ρ .

A *model* \mathbb{M} for \mathbb{E} is a term structure such that, for any definition $(x : \tau = t) \in \mathbb{E}$, $[[x]]_{\mathbb{M}}^{\eta, \rho} = [[t]]_{\mathbb{M}}^{\eta, \rho}$. The existence of models is non-trivial due to recursive definitions, but [17, Theorem 1] guarantees it, thanks to well-foundedness conditions.

c) *Names:* We assume a subset $\mathcal{N} \subseteq \mathcal{X}$ of symbols called *names*, used to denote random samplings. A name $n \in \mathcal{N}$ can only be declared in \mathbb{E} , and must be of type $\tau_0 \rightarrow \tau_1$. The semantics $[[n]]_{\mathbb{M}}$ of [17] interprets a name n as a sequence (indexed by values in τ_0) of *independent identically distributed random samplings* over τ_1 . For example, $[[n 0]]_{\mathbb{M}}$ and $[[n 1]]_{\mathbb{M}}$

are independent random variables. This is also true for distinct name symbols: $\llbracket n \ i \rrbracket_{\mathbb{M}}$ and $\llbracket n' \ j \rrbracket_{\mathbb{M}}$ are always independent.

To guarantee that all instances of a name can be sampled with finite randomness, we require that τ_0 is a finite type.

When indices are not needed, we allow names of base types, corresponding to a single random sampling (independent of other names, as before). For instance, [Example 2](#) uses indexed names $r_A : \text{bint} \rightarrow \text{message}$ but also name $sk_A : \text{message}$ to model the secret key from which $pk_A = pk \ sk_A$ is derived.

B. Execution and Cost Model

Many rules of our logic will be proven by reductions and will thus have a time overhead, *i.e.* the additional time taken by the adversary involved in the reduction. In a concrete security setting, this time overhead needs to be tracked, leading to two difficulties: first, accumulated time overheads can clutter the proof, making it hard to analyze; second, bounding the time overhead in a reduction step may require proving additional cost-related sub-goals, which can be tedious. We alleviate these issues by fine-tuning our execution and cost model. The model described below is a key ingredient to obtain a simple proof system: it allows to reduce (and sometimes entirely remove) time overheads in proof rules.

We now describe PPTM, our set of machines. A machine $\mathcal{A} \in \text{PPTM}$ with l bit-string inputs and k oracles is a Turing machine over the binary alphabet with a special read-only input tape used to receive the security parameter in unary, at least l working tapes, a special *oracle input* tape, and a read-only randomness tape. The working tapes of the machine also serve as input and output tapes, as well as output tapes for the oracles. If \vec{w} is a vector of l bit-strings and \vec{f} is a vector of k bit-string oracles, *i.e.* functions from bit-strings to bit-strings:

$$\vec{w} = (w_1, \dots, w_l) \text{ where } \forall i. w_i \in \{0, 1\}^* \\ \vec{f} = (f_1, \dots, f_k) \text{ where } \forall i. f_i \in \{0, 1\}^* \rightarrow \{0, 1\}^*$$

then the result $\mathcal{A}^{\vec{f}}(1^\eta, \vec{w}, \rho_a)$ of the execution of \mathcal{A} on inputs \vec{w}, \vec{f} with security parameter η and randomness ρ_a is the bit-string obtained as follows:

- \mathcal{A} 's security parameter tape is initialized to 1^η , its randomness tape to ρ_a , its first l working tapes to \vec{w} . The rest of the tapes are all initially empty.
- Then, \mathcal{A} starts its execution, modifying the content of its tapes according to its transition table.
- At any point, \mathcal{A} can make a special transition to call oracle f_i , as follows. \mathcal{A} selects a *working tape* T (T cannot be the oracle input tape). Then, in one step $f_i(x)$ is written on tape T , where x is the content of the oracle input tape and the oracle input tape is reset to the empty tape.
- When it ends, \mathcal{A} returns the content of its output tape.

We allow \mathcal{A} to have several heads, also allowing multiple heads on the same tape, we require that all heads start at the beginning of the tape when the execution of \mathcal{A} start. The layout of the machine \mathcal{A} is static: in particular, no new tape (or head) can be spawned during the execution, and \mathcal{A} 's output tape is always

the same. We require that \mathcal{A} 's transitions are deterministic — any source of randomness must come from the random tape ρ_a .

The *time cost* $\text{time}_{\mathcal{A}}(1^\eta, \vec{w}, \vec{f}, \rho_a)$ is the number of computation steps of $\mathcal{A}^{\vec{f}}(1^\eta, \vec{w}, \rho_a)$.

One of the consequences of our model choice, which helps to simplify overhead expressions in several rules, is that a machine \mathcal{A} running in time t cannot call an oracle on an input of size larger than t . Indeed, such an input would have to be written on the oracle input tape. This relies on the fact that the oracle input tape is not a working tape, and therefore cannot be selected by \mathcal{A} to hold the result of an oracle call. However, we do not make any assumption on oracles, which can return arbitrarily large bit-strings in one step.

C. Global Formulas

We now present the formulas of our logic and their semantics. We use a standard first-order logic over our higher-order terms, with predicates that are specific to concrete security reasoning. We consider, in particular, a predicate \sim for computational indistinguishability and a predicate $[\cdot]_\varepsilon$ for almost-always truth. The syntax is as follows, where $u, \vec{u}, \vec{v}, \phi, \varepsilon, l$ and t are (sequences of) terms (see below for constraints on their types):

$$F ::= F \dot{\Rightarrow} F \mid \dot{\sim} F \mid \dot{\forall}(x : \tau). F \\ \mid \text{const}(u) \mid \text{adv}_{t, \sigma}(u) \mid [\phi]_\varepsilon \mid \vec{u} \sim_\varepsilon \vec{v} \mid \text{blen}_l(u)$$

We write $\mathbb{M} \models F$ when F holds in \mathbb{M} . This is defined as usual in first-order logic for Boolean connectives and quantifiers, and we give below the semantics of our specific predicates. As usual, we obtain $\dot{\wedge}, \dot{\vee}, \dot{\exists}$ from $\dot{\Rightarrow}, \dot{\sim}$ and $\dot{\forall}$. Remark that we use a special notation to distinguish global Boolean connectives and quantifiers $\dot{\Rightarrow}, \dot{\forall}, \dots$ from their local counter-parts $\Rightarrow, \forall, \dots$.

We now describe and give the semantics of our predicates.

a) *Almost-always truth*: The predicate $[\phi]_\varepsilon$ states that ϕ (a term of type bool) is true with a probability of error that is bounded by ε (a term of type real). In order to ease equational reasoning in our proof system, we will more specifically use the *expectation* of the random variable $\rho \mapsto \llbracket \varepsilon \rrbracket_{\mathbb{M}}^{\eta, \rho}$. Formally,

$$\mathbb{M} \models [\phi]_\varepsilon \quad \text{iff} \quad \Pr_\rho(\llbracket \phi \rrbracket_{\mathbb{M}}^{\eta, \rho}) \geq 1 - \mathbb{E}_\rho(\llbracket \varepsilon \rrbracket_{\mathbb{M}}^{\eta, \rho}) \quad \text{for all } \eta$$

For example if n_0, n_1 are names drawn uniformly at random from $\{0, 1\}^\eta$, both $[n_0 = n_1]_0$ and $[n_0 \neq n_1]_{\frac{1}{2^\eta}}$ hold (where 0 and $\frac{1}{2^\eta}$ are terms interpreted as their real counterpart).

b) *Indistinguishability*: The indistinguishability predicate only makes sense when the terms on both sides of the equivalence can be passed to a distinguisher $\mathcal{A} \in \text{PPTM}$, either as inputs or as oracles. The former case corresponds to order 0 terms (*i.e.* terms with order 0 types) and the latter to order 1 terms. We thus only consider instances $\vec{u}_l, \vec{f}_l \sim_\varepsilon \vec{u}_r, \vec{f}_r$ where:

- \vec{u}_l and \vec{u}_r are two same-length sequences of order 0 terms, where the k^{th} terms of each sequence have the same type;
- \vec{f}_l, \vec{f}_r are two sequences of order 1 terms, of length m , with similarly matching types for k^{th} elements;
- ε has type $\overline{\text{int}} \rightarrow \overline{\text{int}}^k \rightarrow \overline{\text{real}}$: intuitively, the bound ε is a function of (a) the execution time of the distinguisher

\mathcal{A} and (b) the number of calls to each of the k oracles in the execution of \mathcal{A} .

The predicate $\vec{u}_l, \vec{f}_l \sim_\varepsilon \vec{u}_r, \vec{f}_r$ states that ε upper-bounds the probability that an adversary distinguishes two scenarios: in the left scenario, the adversary receives the bit-strings \vec{u}_l and has access to the oracles \vec{f}_l ; while in the right scenario, it receives \vec{u}_r and has access to oracles \vec{f}_r . We require that the probability that any adversary \mathcal{A} distinguishes these two scenarios is bounded by the expectation of $\varepsilon(t, \vec{o})$, where t is an upper-bound on \mathcal{A} 's running time and \vec{o} is a vector of integers representing the maximal numbers of calls that \mathcal{A} can make for each oracle. Formally, for an adversary \mathcal{A} with no inputs, we let $\text{time}_{\mathcal{A}}^\eta \in \mathbb{N} \cup \{+\infty\}$ be the maximal run-time of \mathcal{A} for a fixed η :

$$\text{time}_{\mathcal{A}}^\eta \stackrel{\text{def}}{=} \sup_{\vec{u}, \vec{f}, \rho_a} \text{time}_{\mathcal{A}}(1^\eta, \vec{u}, \vec{f}, \rho_a)$$

Similarly, we let $\text{calls}_{\mathcal{A}}^\eta \in (\mathbb{N} \cup \{+\infty\})^m$ be the vector of the maximal numbers of oracle calls of \mathcal{A} for security parameter η on any oracles and random tape. All these upper bounds may be infinite for some adversaries \mathcal{A} : in that case, the bound on the advantage of \mathcal{A} will typically be uninformative, but this is an expected feature of concrete security. Then, $\mathbb{M} \models \vec{u}_l, \vec{f}_l \sim_\varepsilon \vec{u}_r, \vec{f}_r$ iff for any \mathcal{A} with $|\vec{u}_l|$ inputs and m oracles,

$$\left| \Pr(\mathcal{A}^{\vec{f}_l}_{\mathbb{M}}(1^\eta, \llbracket \vec{u}_l \rrbracket_{\mathbb{M}}^{\eta, \rho_a})) - \Pr(\mathcal{A}^{\vec{f}_r}_{\mathbb{M}}(1^\eta, \llbracket \vec{u}_r \rrbracket_{\mathbb{M}}^{\eta, \rho_a})) \right| \leq E_\rho(\llbracket \varepsilon \rrbracket_{\mathbb{M}}^{\eta, \rho}(\text{time}_{\mathcal{A}}^\eta, \text{calls}_{\mathcal{A}}^\eta))$$

where the probabilities are taken over $\rho = (\rho_a, \rho_h)$ in $\mathbb{T}_{\mathbb{M}, \eta}$.

Example 4. Let us assume that enc is interpreted as an IND-CPA encryption with associated public key derivation pk , over a type of plaintexts that is finite. Assuming that $0_{\text{len}} : \text{message} \rightarrow \text{message}$ is interpreted as the function $x \mapsto 0^{|x|}$, we have

$$\lambda i. \text{enc}(i, \text{pk } k, r i) \sim_\varepsilon \lambda i. \text{enc}(0_{\text{len}} i, \text{pk } k, r i)$$

provided that $\varepsilon : \text{int} \rightarrow \text{int} \rightarrow \text{real}$ is such that $\varepsilon t n$ upper-bounds the concrete advantage of an IND-CPA adversary against enc , running in time at most t and making at most n calls to the CPA oracle.

c) *Constancy*: The predicate $\text{const}(u)$ states that the semantics of an arbitrary term u does not depend on the security parameter η or the random tape ρ . Formally, $\mathbb{M} \models \text{const}(u)$ iff. there exists c such that $\llbracket u \rrbracket_{\mathbb{M}}^{\eta, \rho} = c$ for all η and $\rho \in \mathbb{T}_{\mathbb{M}, \eta}$.

d) *Deterministic values*: For any term u of any type and order, the predicate $\text{det}(u)$ states that u is a deterministic value that, as opposed to $\text{const}(u)$, can depend on the security parameter η . More precisely, for any model \mathbb{M} , we have $\mathbb{M} \models \text{det}(u)$ iff. there exists a sequence of values $(c_\eta)_{\eta \in \mathbb{N}}$ such that $\llbracket u \rrbracket_{\mathbb{M}}^{\eta, \rho} = c_\eta$ for any η and ρ .

e) *Bounded length*: The predicate $\text{blen}_l(u)$ states that the interpretation of u is of length at most l . Its precise meaning depends on the order of u , which must be at most 1.

When u is of order 0, l must have type $\overline{\text{int}}$, and we have $\mathbb{M} \models \text{blen}_l(u)$ if for all η and ρ , $\llbracket \text{len}(u) \rrbracket_{\mathbb{M}}^{\eta, \rho} \leq \inf_\rho \llbracket l \rrbracket_{\mathbb{M}}^{\eta, \rho}$. Note that we handle the possible dependency of $\llbracket l \rrbracket$ on ρ by

taking the infimum of $\llbracket l \rrbracket_{\mathbb{M}}^{\eta, \rho}$ over all tapes (for a given η). In practice, proofs only need constant length values.

When u has order 1, l must also be a function: intuitively, it provides a bound on the length of $u \vec{v}$, for all \vec{v} , as a function of the lengths of \vec{v} . For example, we have $\mathbb{M} \models \text{blen}_{\lambda n, m. n+m}(\lambda x, y. \langle x, y \rangle)$ in models where tupling is interpreted as concatenation.

f) *Adversarial computability*: Predicate $\text{adv}_{t, \vec{o}}(u)$ roughly expresses that u is computable by an adversary in time at most t , with at most \vec{o} oracle calls. Again, its meaning depends on the order of u , which must be at most 2; \vec{o} must be empty when u is of order at most 1.

When u has order 0, term t must have type $\overline{\text{int}}$, and $\mathbb{M} \models \text{adv}_t(u)$ means that there exists a machine $\mathcal{A} \in \text{PPTM}$ such that $\llbracket u \rrbracket_{\mathbb{M}}^{\eta, \rho} = \mathcal{A}(1^\eta, \rho_a)$ for all η and ρ , such that $\text{time}_{\mathcal{A}}(1^\eta) \leq \inf_\rho(\llbracket t \rrbracket_{\mathbb{M}}^{\eta, \rho})$. Importantly, the adversary cannot access the honest tape ρ_h here.

This is then generalized to order 1 by asking that a machine can compute the output of the function for each input, in time that depends on the size of order 0 inputs (Meaning that the bound on time t become of type $\overline{\text{int}} \rightarrow \overline{\text{int}}$). We further extend adversarial computability to order-2 terms by making use of oracles: for example, $\lambda f, x. f x$ where x has order 0 and f has order 1 is computable by an adversary taking a function f as oracle and an input x , and returns the output of f on x . The term t in that case imposes a bound on the execution time, as a function of the length of order-0 inputs and of length-bounding functions (of type $\overline{\text{int}} \rightarrow \overline{\text{int}}$) for each of the order-1 inputs. Terms \vec{o} bound the number of calls to each order-1 input. See [Appendix A](#) for more details.

III. PROOF SYSTEM

We now adapt the proof system given for asymptotic computational security in [17] to a concrete security setting. Our logic supports reasoning at two different levels, local and global, according to the kind of formulas being manipulated. This is reflected by the judgement of our logic. A *global* judgement $\mathbb{E}; \Theta \vdash F$ is composed of an environment \mathbb{E} , a set of global formulas Θ (the hypotheses), and a global formula F (the conclusion). A *local* judgement $\mathbb{E}; \Theta; \Gamma \vdash_\varepsilon \phi$ is composed of an environment \mathbb{E} , a set of global formulas Θ (the global hypotheses), a set of local formulas Γ (the local hypotheses), a local formula (i.e. a term of type bool) ϕ (the conclusion), and a term ε of type $\overline{\text{real}}$ (the advantage upper-bound). All formulas (local and global) occurring in a judgement (local and global) must be well-typed in the judgement environment. The validity of global and local judgements is defined by:

$$\begin{aligned} \models \mathbb{E}; \Theta \vdash F & \quad \text{iff.} \quad \models (\tilde{\wedge} \Theta) \Rightarrow F \\ \models \mathbb{E}; \Theta; \Gamma \vdash_\varepsilon \phi & \quad \text{iff.} \quad \models (\tilde{\wedge} \Theta) \Rightarrow [(\wedge \Gamma) \Rightarrow \phi]_\varepsilon \end{aligned}$$

Our global judgements are the same as in [17] since the global logic remains unchanged (we only added new predicates). However, local judgements have been extended with an explicit advantage upper-bound, reflecting the move from $[\phi]$ to $[\phi]_\varepsilon$.

Example 5. Recall that $\text{choose } N$ represents the agent the adversary decides to interact with at step N of the protocol

execution. We can model the fact that there are only two agents in our protocol by adding the following global formula as an axiom:

$$[\forall N.(\text{choose } N = A) \vee (\text{choose } N = B)]_0.$$

The validity of this formula amounts to that of the local judgement:

$$\vdash_0 (\text{choose } N = A) \vee (\text{choose } N = B).$$

This formula means that `choose N` is always either equal to A or B. Remark that this is not equivalent to the following global formula:

$$\tilde{\forall} N. [\text{choose } N = A]_0 \tilde{\vee} [\text{choose } N = B]_0,$$

which states that `choose N` is either always equal to A or always equal to B, and is an unrealistic assumption, as it prevents the adversary from choosing A or B in a randomized fashion.

A. Local Proof System

Our *local proof system* provides a set of generic reasoning rules allowing to deal with Boolean connectives and higher-order quantification, and is essentially an extension of the rules of [17] with an explicit advantage upper-bound. Most local rules of our concrete security proof system are straightforward adaptations of the corresponding asymptotic rules, simply propagating probabilities in the expected way. As an example, we show two typical rules below:

$$\frac{\text{L}_\varepsilon.\text{R-}\wedge \quad \mathbb{E}; \Theta; \Gamma \vdash_{\varepsilon_0} \phi \quad \mathbb{E}; \Theta; \Gamma \vdash_{\varepsilon_1} \psi}{\mathbb{E}; \Theta; \Gamma \vdash_{\varepsilon_0 + \varepsilon_1} \phi \wedge \psi} \quad \frac{\text{L}_\varepsilon.\text{R1-}\vee \quad \mathbb{E}; \Theta; \Gamma \vdash_{\varepsilon} \phi}{\mathbb{E}; \Theta; \Gamma \vdash_{\varepsilon} \phi \vee \psi}$$

The $\text{L}_\varepsilon.\text{R-}\wedge$ rule states that the probability that two formulas do not jointly hold is bounded by the sum of the probabilities that each formula does not hold; while $\text{L}_\varepsilon.\text{R1-}\vee$ states that the probability that a disjunction $\phi \vee \psi$ does not hold is bounded by the probability that the left disjunct ϕ does not hold (the corresponding right rule is not shown here).

The probability that a formula does not hold lies in the real interval $[0; 1]$. Nonetheless, we choose to bound such probabilities using the type `real` corresponding to $\mathbb{R} \cup \{-\infty, +\infty\}$, as having a set that is stable by the standard arithmetic operations ($+$, $-$, and countable sums) usually allows for simpler and more elegant rules. E.g., if we use $[0; 1]$, then the probability upper-bound $\varepsilon + \varepsilon'$ in the $\text{L}_\varepsilon.\text{R-}\wedge$ rule should be replaced by $\min(\varepsilon + \varepsilon', 1)$. Still, there are a few cases where this requires additional checks on the probability bounds, as in the right weakening rule shown below:

$$\frac{\text{L}_\varepsilon.\text{WEAK}_0 \quad \mathbb{E}; \Theta; \Gamma \vdash_{\varepsilon} \psi}{\mathbb{E}; \Theta; \emptyset \vdash_0 \varepsilon \leq \varepsilon'} \quad \frac{\text{L}_\varepsilon.\text{WEAK}_\varepsilon^s \quad \mathbb{E}; \Theta; \Gamma \vdash_{\varepsilon} \psi}{\mathbb{E}; \Theta; \emptyset \vdash_{\varepsilon_0} \varepsilon \leq \varepsilon' \quad \mathbb{E}; \Theta; \emptyset \vdash_0 \varepsilon \leq 1}{\mathbb{E}; \Theta; \Gamma \vdash_{\varepsilon'} \psi \quad \mathbb{E}; \Theta; \Gamma \vdash_{\varepsilon' + \varepsilon_0} \psi}$$

The $\text{L}_\varepsilon.\text{WEAK}_0$ rule allows to replace a probability upper-bound ε by ε' as long as ε' is *always* greater or equal to ε . The $\text{L}_\varepsilon.\text{WEAK}_\varepsilon^s$ rule can be used if the inequality $\varepsilon \leq \varepsilon'$ does not hold unconditionally, simply by adding to ε the bound ε_0 on

the probability that $\varepsilon \leq \varepsilon'$ does not hold. Moreover, because of the usage of probabilistic expectation in the logic semantics (recall that $[\phi]_\varepsilon$ holds if $\Pr(\neg[\![\phi]\!])$ is bounded by $\mathbb{E}([\![\varepsilon]\!])$), we also need to check that the bound in the premises is smaller than one. We show that such a check is necessary below.

Example 6. Let n be a name (without argument) of type message, and assume that names over message are uniform random samplings among bit-strings of length η . Considering the term $\varepsilon \stackrel{\text{def}}{=} 2^\eta \cdot \mathbb{1}_{n=0}$, where $\mathbb{1}_\phi$ is syntactic sugar for (if ϕ then 1 else 0), we have

$$\mathbb{E}([\![\varepsilon]\!]_{\mathbb{M}}^{\eta, \rho}) = 2^\eta \cdot \Pr([\![n]\!]_{\mathbb{M}}^{\eta, \rho} = 0) = 2^\eta \cdot \frac{1}{2^\eta} = 1$$

and $\mathbb{E}; \emptyset; \emptyset \vdash_\varepsilon \perp$ holds. Moreover, $\mathbb{E}; \emptyset; \emptyset \vdash_{\frac{1}{2^\eta}} \varepsilon \leq \frac{1}{2^\eta}$ holds. If the $\text{L}_\varepsilon.\text{WEAK}_\varepsilon^s$ rule did not require that the probability bound in the premises is smaller than 1, we would obtain that $\mathbb{E}; \emptyset; \emptyset \vdash_{2 \cdot \frac{1}{2^\eta}} \perp$ is valid, stating that \perp is true with non-zero probability. This is absurd.

For the full set of local rules, see the appendices of the long version [22].

B. Global Proof System

Our *global proof system* comprises the usual generic reasoning rules (for Boolean connectives and quantifiers), as well as rules dedicated to the predicates of our logic (such as computational indistinguishability). Our global formulas and the semantics of their connectives are the same as in [17] — we only added new predicates. Thus, the generic logical global rules of [17] remain valid in our extension. We recall them in [Appendix B](#) for the sake of completeness, but do not describe them any further. The situation is different for rules related to specific predicates, which must be adapted.

a) *Notation:* Recall that in an indistinguishability formula $\vec{u} \sim_\varepsilon \vec{w}$, the advantage bound ε is a function of the adversary time t and of the number of oracle calls \vec{o} , where \vec{o} contains as many entries as there are terms of order 1 in \vec{u} . For the sake of clarity, we lift the usual mathematical operations to this kind of function terms: e.g., if ε and ε' are two such bounding terms, then $\varepsilon + \varepsilon'$ denotes the term $\lambda t, \vec{o}. \varepsilon t \vec{o} + \varepsilon' t \vec{o}$.

b) *Rewriting:* Any equality, possibly with some error probability, can be used to rewrite terms occurring in a predicate of the logic, by adding the error probability of the equality to the bound in the predicate. For instance, we can rewrite the terms involved in an indistinguishability formula:

$$\frac{\mathbb{E}; \Theta \vdash \vec{u}\{v_0\} \sim_\varepsilon \vec{w} \quad \mathbb{E}; \Theta \vdash [v_0 = v_1]_{\varepsilon_0}}{\mathbb{E}; \Theta \vdash \vec{u}\{v_1\} \sim_{\varepsilon + \varepsilon_0} \vec{w}}$$

This rule only allows to rewrite terms on one side, as rewriting on both sides requires to pay the error probability twice.

Remark 1. Let us show that, when rewriting terms in an equivalence formula, we cannot simultaneously rewrite on the left and right side of the equivalence without paying the equality error twice. Said otherwise, the following rule is unsound:

$$\frac{\mathbb{E}; \Theta \vdash \vec{u}\{v_0\} \sim_\varepsilon \vec{w}\{v_0\} \quad \mathbb{E}; \Theta \vdash [v_0 = v_1]_{\varepsilon_0}}{\mathbb{E}; \Theta \vdash \vec{u}\{v_1\} \sim_{\varepsilon + \varepsilon_0} \vec{w}\{v_1\}} \quad (2)$$

To fix this rule, it is necessary to pay ε_0 twice, i.e. to replace the bound in the conclusion by $\varepsilon + 2\varepsilon_0$. Now, considering three names n_0, n_1, n_2 of type message and 0 : message the zero bitstring, Assuming that names are uniform random samplings among bitstrings of length η , we get that:

$$((n_0 \neq 0) \vee ((n_1 = 0) \wedge (n_2 = 0))) \sim_\alpha \neg(n_1 = 0 \wedge n_2 = 0) \quad (3)$$

with $\alpha \stackrel{\text{def}}{=} \frac{1}{2^\eta} - \frac{1}{2^{3*\eta}} - \frac{1}{2^{2*\eta}}$ as the exact optimal upper-bound. Indeed since any adversary cannot be better than the statistical difference between the two underlining distribution. (And the adversary that accept when the boolean it gets is true and reject otherwise as this exact advantage). Therefore, for any $\eta \in \mathbb{N}^*$,

$$\begin{aligned} \alpha &= \left| \Pr_{\rho \in \mathbb{T}_{M,\eta}} (\llbracket n_0 \neq 0 \rrbracket_M^{\eta,\rho} \vee (\llbracket n_1 = 0 \rrbracket_M^{\eta,\rho} \wedge \llbracket n_2 = 0 \rrbracket_M^{\eta,\rho})) \right. \\ &\quad \left. - \Pr_{\rho \in \mathbb{T}_{M,\eta}} (\llbracket n_1 \neq 0 \rrbracket_M^{\eta,\rho} \vee \llbracket n_2 \neq 0 \rrbracket_M^{\eta,\rho}) \right| \\ &= \left| \Pr_{\rho \in \mathbb{T}_{M,\eta}} (\llbracket n_0 \neq 0 \rrbracket_M^{\eta,\rho}) \Pr_{\rho \in \mathbb{T}_{M,\eta}} (n_1 \neq 0 \vee n_2 \neq 0) \right. \\ &\quad \left. - \Pr_{\rho \in \mathbb{T}_{M,\eta}} (\llbracket n_1 = 0 \vee \llbracket n_2 = 0 \rrbracket_M^{\eta,\rho} \rrbracket_M^{\eta,\rho}) \right| \\ &= \frac{1}{2^\eta} - \frac{1}{2^{3*\eta}} - \frac{1}{2^{2*\eta}} \end{aligned}$$

and

$$((n_0 \neq 0)) \sim_{\frac{1}{2^\eta}} \top \quad (4)$$

and this bound is tight. We know that the probability that both n_1 and n_2 are equal to 0 at the same time is exactly $\frac{1}{2^{2\eta}}$, so $\llbracket (n_1 = 0) \wedge (n_2 = 0) \rrbracket_{\frac{1}{2^{2\eta}}}$ holds. Rewriting this equality with help of the rule given in Eq. (2) on both the left and right side of Eq. (3), we get $((n_0 \neq 0) \vee \perp) \sim_{\alpha + \frac{1}{2^\eta}} \neg \perp$, which is equivalent to $((n_0 \neq 0)) \sim_{\alpha + \frac{1}{2^\eta}} \top$ which is exactly the formula of Eq. (4), but with a tighter bound since $\alpha + \frac{1}{2^{2\eta}} = \frac{1}{2^\eta} - \frac{1}{2^{(k+1)\eta}} < \frac{2}{2^\eta}$. This contradicts the tightness of the bound of Eq. (4), showing that the rule in Eq. (2) is unsound.

Moreover, this rule does not allow to rewrite in the advantage upper-bound ε . Allowing such rewritings can be done, but it requires an additional check to ensure that the initial bound is no greater than 1 (for the same reason than the one given in Example 6):

$$\frac{\mathbb{E}; \Theta \vdash \vec{u} \sim_{\varepsilon_0} \vec{v} \quad \mathbb{E}; \Theta \vdash [\varepsilon_0 = \varepsilon_1]_{\varepsilon_e} \quad [\varepsilon_0 \leq 1]_0}{\mathbb{E}; \Theta \vdash \vec{u} \sim_{\varepsilon_1 + \varepsilon_e} \vec{v}}$$

For the complete version of those rules, see the appendix in the long version [22].

c) *Indistinguishability Rules*: We designed concrete security versions of the rules provided for asymptotic computational indistinguishability in [17]. We show a selected set of our rules in Fig. 1 (the full set is in see the appendices of the long version [22]).

The $G_\varepsilon.E:TRANS$ states that advantage bounds must be added in a transitive step, and $G_\varepsilon.E:CS$ allows to perform a case study over the branching test of a conditional term, summing the advantage bounds of both branches. Remaining rules of Fig. 1 capture reduction-based arguments, and are

$$\begin{array}{c} G_\varepsilon.E:TRANS \\ \mathbb{E}; \Theta \vdash \vec{u} \sim_{\varepsilon_0} \vec{w} \\ \mathbb{E}; \Theta \vdash \vec{w} \sim_{\varepsilon_1} \vec{v} \\ \hline \mathbb{E}; \Theta \vdash \vec{u} \sim_{\varepsilon_0 + \varepsilon_1} \vec{v} \end{array} \quad \begin{array}{c} G_\varepsilon.E:FA-BASE \\ \mathbb{E}; \Theta \vdash \vec{u}_l \sim_\varepsilon \vec{u}_r \quad \mathbb{E}; \Theta \vdash \text{adv}_{t_v}(v) \\ \hline \mathbb{E}; \Theta \vdash \vec{u}_l, v \sim_{\varepsilon'} \vec{u}_r, v \\ \text{where } \varepsilon' \stackrel{\text{def}}{=} \lambda t, \vec{o}. \varepsilon (t + t_v) \vec{o} \end{array}$$

$$\frac{\mathbb{E}; \Theta \vdash \vec{u}_l, f_l \sim_\varepsilon \vec{u}_r, f_r}{\mathbb{E}; \Theta \vdash \vec{u}_l, f_l, f_l \sim_{\varepsilon'} \vec{u}_r, f_r, f_r} G_\varepsilon.E:DUP-FUN$$

where $\varepsilon' \stackrel{\text{def}}{=} \lambda t, \vec{o}. o_1, o_2. \varepsilon t \vec{o} (o_1 + o_2)$

$$\frac{\mathbb{E}; \Theta \vdash \vec{u}_l \sim_\varepsilon \vec{u}_r}{\mathbb{E}; \Theta \vdash \vec{u}_l, n_f \sim_{\varepsilon'} \vec{u}_r, n_f} G_\varepsilon.E:FA-NAMES$$

where $\varepsilon' \stackrel{\text{def}}{=} \lambda t, \vec{o}. o_n. \varepsilon (t + o_n \cdot t_n) \vec{o}$

$$\begin{array}{c} G_\varepsilon.E:CS \\ \mathbb{E}; \Theta \vdash \vec{u}_l, b_l, v_l \sim_{\varepsilon_1} \vec{u}_r, b_r, v_r \\ \mathbb{E}; \Theta \vdash \vec{u}_l, b_l, w_l \sim_{\varepsilon_2} \vec{u}_r, b_r, w_r \\ \hline \mathbb{E}; \Theta \vdash \vec{u}_l, \text{if } b_l \text{ then } v_l \text{ else } w_l \sim_{\varepsilon_1 + \varepsilon_2} \vec{u}_r, \text{if } b_r \text{ then } v_r \text{ else } w_r \end{array}$$

In the rules above, v is an order-0 and f_l, f_r are order-1.

In $G_\varepsilon.E:FA-NAMES$, n_f only occurs in its declaration in \mathbb{E} , and t_n is an upper-bound on the time needed by n_f for a single sampling.

Figure 1. Selected concrete security rules for indistinguishability.

more interesting as they require to carefully track bounds on running time, number of oracle calls, and error probabilities.

$G_\varepsilon.E:FA-BASE$ allows to remove a term v (of order zero) appearing on both sides of an equivalence as long as v can be computed by the adversary. To understand this rule's advantage bound, let us consider an adversary \mathcal{A} against the conclusion, running in time at most t and calling its oracles at most \vec{o} times. By hypothesis, there exists a machine \mathcal{M} computing v in time at most t_v . We can thus build an adversary \mathcal{B} against the premise $\vec{u}_l \sim_\varepsilon \vec{u}_r$ by composing \mathcal{M} with \mathcal{A} . This adversary runs in time $t + t_v$, and calls its oracle \vec{o} times (as did \mathcal{A}). We conclude that \mathcal{B} 's advantage must be at most $\varepsilon (t + t_v) \vec{o}$.

The other reduction-based rules have similar features. The rule $G_\varepsilon.E:DUP-FUN$ allows to get rid of a duplicated oracle, by having the total number of calls to this oracle be the sum of the number of calls to each of its copies. The rule $G_\varepsilon.E:FA-NAMES$ removes a name n_f from an equivalence by letting the adversary sample it itself, at the cost of a time overhead $o_n \cdot t_n$ in the advantage bound, where o_n is the maximal number of adversary calls to the name oracle n_f .

Example 7. *The bound on the case-study rule is tight, and cannot be improved. To see why this is the case, let us consider two Booleans b and b' representing two independent coin-flips (i.e. both have 50/50 chance to be true or false). Let us show that:*

$$\begin{aligned} &\text{if } b \text{ then (if } b \wedge b' \text{ then } \perp \text{ else } \top) \text{ else } \top \\ &\sim \text{if } b \text{ then } \top \text{ else (if } \neg b \wedge b' \text{ then } \perp \text{ else } \top) \end{aligned}$$

is valid using the case-study rule. After the application of the

rule, we are left to prove:

$$b, \text{ if } b \wedge b' \text{ then } \perp \text{ else } \top \sim b, \top$$

$$\text{and } b, \text{ if } \neg b \wedge b' \text{ then } \perp \text{ else } \top \sim b, \top$$

It is easy to check that both premises hold with an upper-bound of $\frac{1}{4}$, since the adversary with the best advantage is the one returning its second input unchanged. Summing both advantages yields a probability upper-bound of $\frac{1}{2}$, as does our rule.

d) *Bi-deduction*: To ease the proof transformations that we will presented in [Section V](#), it is desirable to reduce the number of different rules while keeping the same expressive power. To that end, we designed a single rule, called $G_\varepsilon.\text{E:BI-DEDUCE}$, that captures many rules relying on reduction-based arguments of our proof-system. As the full bi-deduction rule is very technical, we only present a simplified version of the rule which is sufficient to show most of its key aspects. The full rule is described in [Appendix B-A](#).

Let $\vec{u} \stackrel{\text{def}}{=} u_1, \dots, u_m$, $\vec{v} \stackrel{\text{def}}{=} v_1, \dots, v_m$ and $\vec{l} \stackrel{\text{def}}{=} l_1, \dots, l_m$ be three sequences of terms of the same length. Let \mathbb{E} be an environment such that any declared symbol x in \mathbb{E} is only used in eta-long form and is such that $\mathbb{E}; \Theta \vdash \text{adv}_{+\infty}(x)$ is valid. Then the following rule is valid:

$$\frac{\begin{array}{l} G_\varepsilon.\text{E:BI-DEDUCE}^s \\ \mathbb{E}; \Theta \vdash \vec{u} \sim_\varepsilon \vec{v} \quad \mathbb{E}; \Theta \vdash \text{adv}_{t_c, \vec{o}_c}(C) \\ \mathbb{E}; \Theta \vdash \bigwedge_{i \leq n} \text{blen}_{l_i}(u_i) \tilde{\wedge} \text{blen}_{l_i}(v_i) \end{array}}{\mathbb{E}; \Theta \vdash C \vec{u} \sim_{\varepsilon'} C \vec{v}}$$

where ε' must precisely account for the simulation times:

$$\varepsilon' \stackrel{\text{def}}{=} \begin{cases} \lambda t. \varepsilon (t + t_c \vec{l}) (\vec{o}_c) & ((C \vec{u}) \text{ of order } 0) \\ \lambda t, o. \varepsilon (t + o \cdot t_c \vec{l}) (o \cdot \vec{o}_c) & ((C \vec{u}) \text{ of order } 1) \end{cases}$$

Recall that $C \vec{u}$ stand for $((t u_1) \dots u_n)$ when $\vec{u} = u_1, \dots, u_n$.

This rule subsumes the $G_\varepsilon.\text{E:FA-BASE}$ rule, and has a similar structure. It allows to remove a computable (by the adversary) context C from both side of an equivalence. Here, since C can be of order at most two (\vec{u} and \vec{v} contains terms of order 0 or 1), the predicate $\text{adv}(C)$ take two subscripts arguments t and \vec{o}_C . The first argument is the computation time of C and in particular take the length of the terms in \vec{u} and \vec{v} as argument. Those length are bounded by the vector \vec{l} and the premise on $\text{blen}(\cdot)$. The second one takes into account the number of calls to the first-order terms in \vec{u} and \vec{v} .

Example 8. Consider name symbols k, k', r, n, n' of type message, representing independent samplings of length η . We show the following toy formula:

$$\text{enc}(0_{\text{len}} n, r, \text{pk } k) \sim_{\varepsilon'} \text{enc}(0_{\text{len}} n', r, \text{pk } k').$$

First, the formula above can be rewritten without error into:

$$\begin{aligned} & (\lambda x, y, z. \text{enc}(0_{\text{len}} x, y, \text{pk } z)) n r k \\ & \sim_{\varepsilon'} (\lambda x, y, z. \text{enc}(0_{\text{len}} x, y, \text{pk } z)) n' r k' \end{aligned}$$

Let C be the context defined by:

$$C \stackrel{\text{def}}{=} \lambda x, y, z. \text{enc}(0_{\text{len}} x, y, \text{pk } z)$$

Applying $G_\varepsilon.\text{E:BI-DEDUCE}^s$ leaves us with the main premise

$$n, r, k \sim_\varepsilon n', r, k'.$$

and additional premises $\text{adv}_{t_C}(C)$ (for some t_C to be determined), and $\text{blen}_\eta(m)$ for m any of the previous names. The $\text{blen}_\eta(m)$ premise can be easily proven under our assumptions on the size of names. For the $\text{adv}_{t_C}(C)$ premise, assume that t_{enc}, t_0 and t_{pk} are upper-bounds on the computation times of the underlying functions, and also that l_0 and l_{pk} represent the lengths of the outputs of the functions. We can prove $\text{adv}_{t_C}(C)$ with: $t_C \stackrel{\text{def}}{=} \lambda l_x, l_y, l_z. t_{\text{enc}}(l_0 l_x) l_y l_{\text{pk}} l_z + t_0 l_x + t_{\text{pk}} l_z$

(Here, since all of input terms of C are of order 0, \vec{o}_C does not appear.) Finally, we obtain: $\varepsilon' \stackrel{\text{def}}{=} \lambda t. \varepsilon(t + t_C \eta \eta \eta)$.

The full bi-deduction rule (in [Appendix B-A](#)) generalizes the rule above in several ways: i) we provide additional fresh names to the contexts; ii) we allow for an arbitrary number of contexts instead of a single one. These two extensions make our bi-deduction rule general enough to subsume many reduction-based CCSA rules: e.g. adding names captures the $G_\varepsilon.\text{E:FA-NAMES}$ rule capabilities, while having many contexts captures the duplication rules (such as $G_\varepsilon.\text{E:DUP-FUN}$).

C. Global Induction

The asymptotic logic of [17] only supports induction for a constant number of steps, which can be expressed using the following *asymptotic logic* induction rule:

$$\frac{\begin{array}{l} \mathbb{E}; \Theta \vdash \text{well-founded}_\tau(<) \wedge \text{det}(<) \\ \mathbb{E}; \Theta \vdash \tilde{V}(x : \tau). \text{const}(x) \Rightarrow \\ (\tilde{V}(x_1 : \tau). \text{const}(x_1) \Rightarrow [x_1 < x]_{\text{negl}} \Rightarrow F\{x \mapsto x_1\}) \Rightarrow F \end{array}}{\mathbb{E}; \Theta \vdash \tilde{V}(x : \tau). \text{const}(x) \Rightarrow F}$$

where $[\phi]_{\text{negl}}$ holds iff the probability that ϕ does not hold is negligible, i.e. $\Pr_\rho([\neg\phi]_{\mathbb{M}}^{\eta, \rho}) \in \text{negl}(\eta)$, and $\text{well-founded}_\tau(<)$ is a global formula stating that $([\tau]_{\mathbb{M}}^\eta, [\cdot]_{\mathbb{M}}^\eta)$ is well-founded for every η (see the long version [22] for details).

Remark 2. Actually, the asymptotic logic induction rule of [17] is not restricted to values x such that $\text{const}(x)$ holds. This is a mistake, as the rule is unsound without this, as shown in the example below.

Example 9. Consider a model \mathbb{M} such that type nat is interpreted as the set of natural numbers, and $<$ as the standard order over \mathbb{N} . For all $i \in \mathbb{N}$, we let $(X^i(\eta))_{\eta \in \mathbb{N}}$ be the η -indexed sequence of random variables over \mathbb{N} defined by:

$$X^i(\eta)(\rho) \stackrel{\text{def}}{=} \max(\eta - i, 0). \quad (\text{for any } \rho \in \mathbb{T}_{\mathbb{M}, \eta})$$

For every $i < j$, there exists a rank η_0 such that for every $\eta \geq \eta_0$, we have $X^i(\eta) > X^j(\eta)$ (e.g. $\eta_0 = j$). Thus:

$$\Pr_\rho(X^i(\eta)(\rho) > X^j(\eta)(\rho)) = 1. \quad (\text{for } \eta \text{ large enough})$$

Hence, the set of η -indexed sequence of random variables $\mathbb{R}\mathbb{V}_{\mathbb{M}}(\text{nat})$ ordered by: “ X is smaller than Y ” iff:

$$\Pr_{\rho} \left(\llbracket \neg(x < y) \rrbracket_{\mathbb{M}[x \mapsto X, y \mapsto Y]}^{\eta, \rho} \right) \in \text{negl}(\eta)$$

is not well-founded, even though $<$ is well-founded over \mathbb{N} .

a) *Hybrid arguments*: Actually, the issue shown above is a well-known pitfall of hybrid arguments, which are the way inductions are usually used in cryptographic proofs. Essentially, an hybrid argument is the specialization of the induction principle to the case where the property to be shown is of the form $\vec{u}(n) \sim \vec{u}(n+1)$. Informally, it says that:

$$\vec{u}(1) \sim \vec{u}(2) \sim \dots \sim \vec{u}(n) \quad \text{implies} \quad \vec{u}(1) \sim \vec{u}(n).$$

In its asymptotic formulation, the above argument is sound if n is constant, or more generally if the bound $\varepsilon(\eta, i)$ on the advantage of any adversary against $\vec{u}(i) \sim \vec{u}(i+1)$ is *uniformly* bounded by a negligible function $\varepsilon'(\eta)$ — here, uniformly means independently from i . The counter-example shown in [Example 9](#) shows a typical case of the issues arising when uniformity does not hold. (We refer the reader to [20] for a detailed discussion of uniformity and hybrid arguments.)

Since *asymptotic* CCSA logics do not provide explicit advantage bounds, they cannot capture the more expressive variant of the hybrid argument, and must restrict themselves to the constant case where the number of induction steps does not depend on the security parameter η . This limits the logics applicability: e.g., they can only be used to prove the security of a protocol for a constant number of sessions, instead of the stronger polynomial-level of security (this is exactly what happens in [Example 3](#)).

b) *Concrete security induction*: Our concrete logic does not suffer from such a restriction, as the advantages can be explicitly established during an induction proof. More precisely, our logic allows for the following global induction principle:

$$\frac{\text{G}_{\varepsilon}.\text{INDUCTION} \quad \begin{array}{l} \mathbb{E}; \Theta \vdash \text{well-founded}_{\tau}(<) \\ \mathbb{E}; \Theta \vdash \tilde{\forall}(x : \tau). (\tilde{\forall}(x_1 : \tau). [x_1 < x]_0 \Rightarrow F\{x \mapsto x_1\}) \Rightarrow F \end{array}}{\mathbb{E}; \Theta \vdash \tilde{\forall}(x : \tau). F}$$

Note that we require $[x_1 < x]_0$, avoiding the pitfall described in [Example 9](#).

D. Freshness and Cryptographic Rules

A fourth class of rules deals with cryptographic assumptions and reasoning about probabilistic independence. We present here two examples of such rules.

A crucial tool for such rules are freshness conditions, as defined in [17]. Intuitively, the formula $\phi_{\text{fresh}}^{n,v}(\vec{u})$ is an approximation of the conditions under which the name $n v$ can appear in the *generalized* subterms of \vec{u} , i.e. subterms considering the expansion of (recursive) definitions. If $\llbracket \phi_{\text{fresh}}^{n,v}(\vec{u}) \rrbracket_{\mathbb{M}}^{\eta, \rho}$ holds, then the term u can be computed without sampling the name n at index $\llbracket v \rrbracket_{\mathbb{M}}^{\eta, \rho}$. A full definition can be found in the appendices of the long version [22].

The general idea of the fresh rule is that a name that is not used anywhere can be replaced by another one. In addition to

adding explicit advantages, we generalize the rule from [17] to work under a context (to ease proof transformations):

$$\frac{\text{G}_{\varepsilon}.\text{E}:\text{FRESH} \quad \begin{array}{l} \mathbb{E}; \Theta \vdash \left[\phi_{\text{fresh}}^{n,i}(\vec{u}, C(n_{\text{fresh}}())) \tilde{\forall} \phi_{\text{fresh}}^{n_{\text{fresh}},i}(\vec{u}, C(n_{\text{fresh}}())) \right]_{\varepsilon'} \\ \mathbb{E}; \Theta \vdash \vec{u}, C(n_{\text{fresh}}()) \sim_{\varepsilon} \vec{v} \end{array}}{\mathbb{E}; \Theta \vdash \vec{u}, C(n i) \sim_{\varepsilon + \varepsilon'} \vec{v}}$$

The freshness premise ensures that $n i$ and $n_{\text{fresh}}()$ are not sampled at the same time in any evaluation of C (with an error at most ε'). E.g., we can rewrite (if b then $n i$ else $n t$) into (if b then $n i$ else $n_{\text{fresh}}()$) as, in this context, $n_{\text{fresh}}()$ is only sampled under condition $\neg b$, and $n i$ is only sampled under condition b .

We present in [Fig. 2](#) a simplified version of the rule for indistinguishability against chosen-ciphertext attack (CCA1), that states that an attacker cannot learn anything about the content of a ciphertext except its length. We model it by replacing the plaintext inside the ciphertext with the same length of zeros. The full version can be found in the long version of the paper [22] with other cryptographic rules.

A key ingredient for this rule is the ability to simulate all relevant terms during the reduction for the soundness proof. This is done through the predicate \vdash_t^c , which has a lot of similarity with the adv_t predicate since they both represent the ability of the adversary to compute something. However, in the \vdash_t^c predicate, the adversary can simulate the randomness of the protocol, while it cannot in the adv_t function. This predicate is an adaptation of the \vdash_{pptm} of [17] to the explicit advantage settings. Details can be found in [Appendix B-B](#).

We also need some condition on key and randomness usage in the context. We reuse the conditions from [17]: intuitively, ϕ_{key} means that all needed terms can be computed without using the secret key except for decryption, while ϕ_{rand} ensures that the randomness used in the ciphertext is fresh, and ϕ_{dec} ensures that the context does not decrypt with the secret key. Formal definitions can be found in the long version [22].

This rule is a generalization of the rule in [17] since it allows to apply the CCA1 hypothesis under a context. Again, this will be important in our proof transformation result.

IV. APPLICATION TO PRIVATE AUTHENTICATION

In this section, we prove the unlinkability of the PA protocol with our proof system. We first describe an idiomatic CCSA proof, without specifying upper-bounds on advantages. Then we analyze these bounds and explain why the usage of case studies in proofs by induction yields advantage upper-bounds which are exponential in the number of inductive steps, failing to provide a polynomial level of security. Finally, we describe how our proof can be modified to fix this issue, hinting at the general result of the next section.

A. Idiomatic Proof by Case Study

Following up on [Example 2](#), we are going to prove, in our concrete security logic, that PA is unlinkable for any

$G_\varepsilon.CCA1^s$

$$\frac{\begin{array}{l} \mathbb{E}; \Theta; \emptyset \vdash_{t_0}^c \vec{u}, b, m, i_r, i_k \quad \mathbb{E}; \Theta; \emptyset \vdash_{t_C}^c C \quad \mathbb{E}; \Theta \vdash \det(i_r) \bar{\wedge} \det(i_k) \\ \mathbb{E}; \Theta \vdash \text{blen}_{l_e}(\text{enc } m (r i_r) (\text{pk}(k i_k))) \bar{\wedge} \text{blen}_{l_e}(\text{enc } (0_{\text{len}}(m)) (r i_r) (\text{pk}(k i_k))) \\ \mathbb{E}; \Theta; \emptyset \vdash_{\varepsilon_\phi} \phi_{\text{key}}^{k, i_k}(\vec{w}, C) \wedge \phi_{\text{rand}}^{r, i_r}(\vec{w}, C) \wedge \phi_{\text{dec}}^{k, i_k}(C) \quad \mathbb{E}; \Theta \vdash \vec{u}, \text{if } b \text{ then } C (\text{enc } (0_{\text{len}}(m)) (r i_r) (\text{pk}(k i_k))) \text{ else } u_e \sim_\varepsilon \vec{v} \end{array}}{\mathbb{E}; \Theta \vdash \vec{u}, \text{if } b \text{ then } C (\text{enc } m (r i_r) (\text{pk}(k i_k))) \text{ else } u_e \sim_{\varepsilon_f} \vec{v}}$$

where \vec{v} is of order 0 and $\vec{w} \stackrel{\text{def}}{=} \vec{u}, b, m, i_r, i_k$ (note that u_e is *not* in \vec{w}) and $\varepsilon_f \stackrel{\text{def}}{=} \lambda t, \vec{d}. \varepsilon t \vec{d} + \varepsilon_\phi + \varepsilon_{CCA}(t + t_C l_e + \vec{1} \cdot \vec{t}_0), \vec{1} \cdot \vec{w} = \sum_i w_i$ is the scalar product of \vec{w} with the sequence $1, \dots, 1$ of the same length.

Figure 2. Simplified rule for the CCA1 cryptographic game.

deterministic number of sessions N that can be computed by the adversary in time $(t N)$. More precisely, we must derive

$$\mathbb{E}; \Theta \vdash \tilde{\forall} t, N. \det(N) \Rightarrow \text{adv}_{t N}(N) \Rightarrow F_N$$

where $F_N \stackrel{\text{def}}{=} \text{frame } N \sim_{\varepsilon N} \text{frame}_{\text{id}} N$ and where $\det(t)$ is a new predicate of the logic that states that t is a deterministic value that can depend on the security parameter η (see [Appendix A](#)). The bound εN will be described in the next subsection. In fact, all bounds are omitted for now.

Using an induction principle specialized on integers that only applies to deterministic variables N , it suffices to prove, after some minor book-keeping, that

$$\mathbb{E}; \Theta, \det(0), \text{adv}_{t_0}(0) \vdash F_0 \quad \text{and} \quad \mathbb{E}_0; \Theta_0 \vdash F_{N+1}$$

where \mathbb{E}_0 declares $N : \text{bint}$ and $\Theta_0 \stackrel{\text{def}}{=} \Theta, \det(N), \text{adv}_{t N}(N), F_N$.

We focus on the proof of F_{N+1} , which is the interesting case. From definition of **frame** and **frame_{id}**, and using bi-deduction (*i.e.* the $G_\varepsilon.E:BI-DEDUCE^s$ rule) to remove the tupling function $\langle \cdot, \cdot \rangle$, we must show that:

$$\mathbb{E}_0; \Theta_0 \vdash \begin{array}{l} \text{frame } N, \text{output } (\text{choose } (N+1), N+1) \\ \sim \text{frame}_{\text{id}} N, \text{output}_{\text{id}} (\text{choose}_{\text{id}} (N+1), N+1) \end{array} \quad (5)$$

From now on, we omit \mathbb{E}_0 and Θ_0 , as they remain unchanged throughout the rest of the proof. We assume that the adversary can only choose to interact with A or B, *i.e.* we use the modeling axiom:

$$\tilde{\forall} N_0. [\text{choose } N_0 = A \vee \text{choose } N_0 = B]_0.$$

Using this axiom for $N_0 = N+1$, and by rewriting (with no error) in the terms on the left of the equivalence in [Eq. \(5\)](#), we obtain the (left) terms:

$$\text{frame } N, \text{if } \text{choose } (N+1) = A \text{ then output } (A, N+1) \\ \text{else output } (B, N+1)$$

We do the same on the terms on the right of \sim , which yields the same conditional term, except that it uses the idealized versions of **frame**, **output**, and **choose**.

Then, doing a case study with $G_\varepsilon.E:CS$ on $b \stackrel{\text{def}}{=} (\text{choose } (N+1) = A)$ on the left and $b_{\text{id}} \stackrel{\text{def}}{=} (\text{choose}_{\text{id}} (N+1) = A)$ on the right, we get the two equivalences:

$$\begin{array}{l} \text{frame } N, b, \text{output } (A, N+1) \\ \sim \text{frame}_{\text{id}} N, b_{\text{id}}, \text{output}_{\text{id}} (A, N+1) \end{array} \quad (\dagger\text{-A})$$

$$\begin{array}{l} \text{frame } N, b, \text{output } (B, N+1) \\ \sim \text{frame}_{\text{id}} N, b_{\text{id}}, \text{output}_{\text{id}} (B, N+1) \end{array} \quad (\dagger\text{-B})$$

We focus on case $(\dagger\text{-A})$ (case $(\dagger\text{-B})$ is similar). By definition, **choose** $(N+1)$ and **choose_{id}** $(N+1)$ are equal to, resp.,

$$\text{att}_c(\text{frame } N) \quad \text{and} \quad \text{att}_c(\text{frame}_{\text{id}} N).$$

Using bi-deduction to remove the computation $\lambda x, y. (x = \text{att}_c(y))$, and replacing the outputs by their definitions, we get:

$$\begin{array}{l} \text{frame } N, \text{enc}(\langle \text{pk}_A, n_A(N+1) \rangle, \text{pk}_B, r_A(N+1)) \\ \sim \text{frame}_{\text{id}} N, \text{enc}(\langle \text{pk}'_A(N+1), n_A(N+1) \rangle, \text{pk}_B, r_A(N+1)) \end{array}$$

We apply the CCA1 rule twice, once per side, to zero-out the content of both encryptions using $0_{\text{len}}(\cdot)$ — the verification of the premises that are not an indistinguishability predicate of this rule is not central, and we omit it. For readability's sake, we presented the mutually recursive functions **frame**, **choose**, **input** and **output** in [Example 2](#) without using the **let** in construct. However while not done explicitly here, for an exact proof, the frame should be bound by a **let** in in order to not compute it multiple times and make all these functions polynomial-time computable (with the same semantic).

Thanks to some basic modeling assumptions on lengths, we then get that both plaintexts are of the same length u_{len} , which we assume to be a publicly known quantity (*i.e.* such that $\text{adv}_{t_u}(u_{\text{len}})$ holds for some efficient time t_u). Then, using bi-deduction and rewriting without error to remove the encryption computation, the zeroing function $0_{\text{len}}(\cdot)$ and the publicly known length u_{len} , we get:

$$\text{frame } N, \text{pk}_B, r_A(N+1) \sim \text{frame}_{\text{id}} N, \text{pk}_B, r_A(N+1).$$

The name $r_A(N+1)$ is fresh in this context, as it is never sampled in **frame** N and **frame_{id}** N , and thus can be sampled by the adversary itself. This reasoning can be captured using the freshness rules $G_\varepsilon.E:FRESH$ and $G_\varepsilon.E:FA-NAMES$, which yields, after discharging some minor proof obligations:

$$\text{frame } N, \text{pk}_B \sim \text{frame}_{\text{id}} N, \text{pk}_B$$

Observe that pk_B can be obtained by computation from the frame on both sides, as it is included in **frame** 0 and **frame_{id}** 0. Thus, by bi-deduction, it remains to prove that **frame** $N \sim \text{frame}_{\text{id}}$ N , which follows from the induction hypothesis F_N .

The proof of the branch $(\dagger\text{-B})$ follows very similar steps, and also uses the CCA1 cryptographic rule exactly once.

B. Analysis of the Advantage Upper-Bound

We now analyze the advantage bound obtained from the previous proof. We are not interested in its precise expression, but seek to analyze how it grows with the number of sessions.

As this is a proof by induction, we need to find the recurrence relation between the advantages ε_N and ε_{N+1} (the initial value ε_0 is of little interest here). Our proof of the inductive step is of the form described in the figure on the right, where F_{N+1}^A and F_{N+1}^B are, resp., the formulas in Eq. (†-A) and Eq. (†-B), and the branching comes from the case study on whether `choose = A` or `choose = B`.

$$\frac{\frac{F_N}{\Pi_A}}{F_{N+1}^A} \quad \frac{\frac{F_N}{\Pi_B}}{F_{N+1}^B}$$

$$\frac{\Pi}{F_{N+1}}$$

Figure 3. Shape of the proof of the inductive step for the PA protocol.

The advantage bounds are modified by the proof as follows, when *descending* from the induction hypothesis F_N to the conclusion F_{N+1} :

- Both leaves start at advantage ε_N .
- *Simple reduction-based* steps (e.g. bi-deductions) add a time overhead to the current advantage: roughly, the advantage $\lambda t. \varepsilon_0 t$ in premise is changed into $\lambda t. \varepsilon_0 (t+t_0)$, where t_0 is the time needed to evaluate the added computations.
- *Cryptographic* steps increase the advantage by the advantage in breaking the cryptographic assumption under consideration: e.g. here, the advantage is increased by $\varepsilon_{CCA}(t+t_{CCA})$, where t_{CCA} is a bound on the time needed to simulate the particular context used in our proof.
- *Logical* steps, such as rewriting without error or the logical reasoning rules, leave the advantage (mostly) unchanged.
- Last but not least, the case study rule $G_\varepsilon.E:CS$ adds the advantages of both premises.

Putting everything together, we get the relation:

$$\varepsilon(N+1) = \lambda t. \left(\varepsilon_N(t+t_A) + 2 \varepsilon_{CCA}(t+t_{CCA}^A) \right) + \left(\varepsilon_N(t+t_B) + 2 \varepsilon_{CCA}(t+t_{CCA}^B) \right) \quad (6)$$

where, very roughly, t_A is the sum of the time overheads in Π_A and Π , t_{CCA}^A is the time needed to simulate the context when the CCA1 rule is applied (plus any time overhead added later); t_B and t_{CCA}^B are similar, but for the right branch corresponding to the case `choose = B`. The top quantity of Eq. (6) comes from the left branch of the proof in Fig. 3, while the bottom quantity comes from the right branch. Any advantage ε_N satisfying the recurrence relation in Eq. (6) is exponential in N (as it must contain 2^N CCA1 advantages upper-bounds ε_{CCA}).

To summarize, our idiomatic proof yields the parametric level of security that is today's standard with Squirrel: the advantage is negligible in η for any N that does not depend on η . However, it is not satisfying when N is polynomial in η , *i.e.* when the attacker can dynamically choose for how long it interacts with the protocol: in that case, the CCA1 advantage is multiplied by 2^η , canceling out assumption that this cryptographic advantage is negligible. To obtain a polynomial level of security, we need

a proof with (in particular) an at-most polynomial factor for the CCA1 advantage.

C. Fixing the PA Proof

The crucial problem in our previous proof is the fact that the induction hypothesis is used twice. This stems from the fact that we use a case study to split the (†-A) and (†-B) cases. In order to fix this problem, we do away with the case study and apply the $G_\varepsilon.CCA1^s$ rule under the context.

Starting from the equivalence goal we have in the naive proof, between the terms

```
frame N, if choose (N + 1) = A then output (A, N + 1)
      else output (B, N + 1)
```

and their idealized versions, we need to define contexts under which we can apply the $G_\varepsilon.CCA1^s$ rule.

Replacing the `output` by its definition, the $G_\varepsilon.CCA1^s$ rule is applied for the A side with the trivial context $C = (\lambda x.x)$, condition ($b = \text{choose}(N+1) = A$), ignored term ($u_e = \text{output}(B, N+1)$), and vector of outside terms `frame N`. We obtain the left side of the equivalence goal where the encryption in the A branch is zeroed out. We do the same transformation in the B branch, and on the right side of the equivalence. Finally, using some lengths reasoning and bi-deduction, we get the goal:

```
frame N, choose (N + 1), pk_B, pk_A, r_A (N+1), r_B (N+1)
~ frame_id N, choose (N + 1), pk_B, pk_A, r_A (N+1), r_B (N+1).
```

We conclude by remarking that `choose(N+1), pk_B, pk_A` can be computed from `frame N` and the fact that $r_A(N+1), r_B(N+1)$ are not used in `frame N` and therefore are fresh sampling. We can now directly use the induction hypothesis exactly once.

Summing up the advantages of the proof, and taking t_{CCA} as an upper-bound of all time overheads in the different uses of the CCA rule, we get the relation:

$$\varepsilon(N+1) = \lambda t. (\varepsilon_N(t+t'_{AB}) + 4 \cdot \varepsilon_{CCA}(t+t_{CCA})).$$

This relation is satisfied by an ε that uses a linear number (in N) of ε_{CCA} and with some polynomial time overheads in both N and η , and is thus negligible for N polynomial in η .

More generally, the technique we used in this proof, applying cryptographic rules under context rather than using case studies (*i.e.* commuting case studies and cryptographic rules) can be generalized to a general proof rewriting strategy applicable to most existing idiomatic proofs, as shown in the next section.

V. PROOF REWRITING STRATEGY

We now present our proof transformation result, which can be used to improve the advantage bound of a large class of proofs, allowing to reach a polynomial level of security. Our result proceeds by rewriting the proof to be improved: i) we design a set of local proof transformations; ii) we design a proof transformation strategy that terminates on *admissible* proofs — a syntactic proof fragment defined below — and guarantees that normalized proofs only use their induction hypothesis once; and iii), we prove that if Π is an admissible proof that

only features terms that can be computed in polynomial-time, and if Π only relies on cryptographic assumptions that hold with a negligible advantage, then a normalization of Π by our proof transformation strategy provides polynomial security.

Our proof transformation strategy operates on the proof of the inductive step of a proof of security of a protocol, and is based on two key ideas, exposed below.

i) *Proof commutations*: we move the applications of the case study rule $G_\varepsilon.E:CS$ upward across the proof-tree using proof commutations. There are some rules that do not commute with $G_\varepsilon.E:CS$, most notably the bi-deduction rule $G_\varepsilon.E:BI-DEDUCE$: such rules also need to be moved upward, pushed by rising case study rules. This naturally separates our rules in two disjoint sets, ascending rules and descending rules, For each pair of ascending rule A and descending rule D , there must exist a proof commutation \blacktriangleright_{AD} . An ideal commutation would look like the figure on the right. Of course, our commutations are almost never that simple, as the rules to commute can have several premises, and the proof transformation may need to introduce new rules when modifying the proof.

ii) *Collapsing case-studies*: once all ascending rules, and thus application of the case-study rule, are at the top of the proof-tree, we start a second round of proof transformations $\blacktriangleright_{col}$, this time with the goal of *collapsing* applications of the case study rule. For example, a case study CS followed by an application of the bi-deduction B and induction I rules in both branches is an unnecessary detour, that can be replaced by a single application of a bi-deduction and induction hypothesis. We schematize this in the figure on the right, where $F_1 = F'_1$ up-to some errorless rewriting.

$$\frac{\frac{\frac{F_3}{F_2} D}{F_1} A}{F_1} \blacktriangleright_{AD} \frac{\frac{F_3}{F'_2} A}{F_1} D$$

$$\frac{\frac{\frac{F}{F_2} I}{F_1} B}{F_1} \frac{\frac{F}{F_3} I}{F_1} B}{F_1} CS \blacktriangleright_{col} \frac{\frac{F}{F'_1} I}{F_1} B$$

Figure 4. A case-study collapse.

A. The Proof Fragment of Admissible Proofs

We now define the proof fragment our result applies to. To that end, we first partition rules according to their behavior during the proof transformations. Then, we describe the restrictions defining the proof fragment of admissible proofs.

a) *Rule partitioning*: We partition our set of rules in three: *leaf rules* are the leaves of our proofs, *ascending rules* are the rules that will be pushed upwards by our proof transformations, while *descending rules* will be moved downwards.

We have four *leaf rules*: the reflexivity of equivalence, the application of an axiom, and the trivial rules with a false hypothesis. We have three *ascending rules*: the case study rule, the bi-deduction rule and the hypothesis weakening rule. Finally, *descending rules* are any other rule of the logic that can have an equivalence in conclusion, at the exception of: induction, transitivity, rewriting with errors, the upper-bound

weakening rules, the left disjunction rules for $\tilde{\vee}$ and \vee , and the excluded-middle rule.

Moreover, all rules subsumed by the bi-deduction rule are excluded from consideration here, without loss of generality. All those rules, as well as a quick discussion on why we do not support some rules in our proof-transformation result, and the impact that removing these rules has, can be found in the long version [22].

All ascending or descending rule (except the case-study rule, and two others) have at most one indistinguishability premise, which we call the *principal premise* of the rule. The rest of the premises are the *auxiliary premises* of the rule. For the two other rules, the premise with the larger context is the principal one, the other premise is auxiliary. Finally, $G_\varepsilon.E:CS$ is the only rule with two principal premises.

b) *Admissible proofs*: We present and justify the class of proofs to which our result applies using our running example. We can identify in our example proof in Section IV-A a proof structure which is standard in CCSA proofs (e.g. [15], [23], [24]). Looking at the proof of the induction step from the top to the bottom, we see that we start with the induction hypothesis $\text{frame } N \sim \text{frame}_{id} N$ which is then iteratively augmented with additional elements (say, \vec{u}_0 on the left and \vec{v}_0 on the right) and surrounded by a computing context (say C), which yields intermediate equivalences of the form $C(\text{frame } N) \vec{u}_0 \sim C(\text{frame}_{id} N) \vec{v}_0$. Then, the proof modifies the context C or the elements in \vec{u}_0 and \vec{v}_0 , e.g. by rewriting, bi-deduction, application of cryptographic rules, and by merging distinct branches using the case study rule. The proof ends when it reaches the target equivalence $\text{frame}(N+1) \sim \text{frame}_{id}(N+1)$.

Interestingly, We add the induction hypothesis $\text{frame } N \sim \text{frame}_{id} N$ to ours hypothesis at the beginning of our proof and only use it at the end to close our induction. Further, we can observe that the case study over $\text{choose} = A$ in Section IV-A deals with a condition that can be computed by the adversary from the induction hypothesis terms $\text{frame } N$ and $\text{frame}_{id} N$. Essentially, our proof fragment is the subset of proofs of this form. Roughly, we say that a proof is $(\vec{u}; \vec{v})$ -admissible if:

- it is only a single leaf rule with $\vec{u} \sim_\varepsilon \vec{v}$ (for some arbitrary term ε) or its root is an ascending or descending rule, and the principal premise of this rule is proven by an $(\vec{u}; \vec{v})$ -admissible proof.
- The application of the case study rule must only be done on a branching condition that can be computed (i.e. bi-deduced) from \vec{u} in the left side of \sim , (respectively \vec{v} in right side)
- All intermediate equivalences along the *main trunk* (i.e. the part of the proof-tree which only considers principal premises of rules, starting from the root) of the proof are of the form $\vec{u}, \vec{u}_0 \sim_\varepsilon \vec{v}, \vec{v}_0$ where \vec{u}_0, \vec{v}_0 and ε are arbitrary terms. Said otherwise, \vec{u} and \vec{v} are not modified in the main trunk.

The full definition of $(\vec{u}; \vec{v})$ -admissibility can be found in the long version [22]. A proof is said to be admissible if it is $(\vec{u}; \vec{v})$ -admissible for some terms $(\vec{u}; \vec{v})$. Finally, an auxiliary

| | | Descending Rules R | |
|-----------------|----|--|---|
| | | R is one of the RD rules | Other rules R |
| Ascending rules | CS | $\frac{\frac{\Pi_1 \text{ Aux}_R}{\cdot} R \quad \Pi_2}{\cdot} \text{CS} \triangleright_{AD} \frac{\frac{\Pi_1}{\cdot} \quad \frac{\Pi_2}{\cdot} W^*}{\cdot} \text{CS} \frac{\text{Aux}_R}{\cdot} R$ | $\frac{\frac{\Pi_1 \text{ Aux}_R}{\cdot} R \quad \Pi_2}{\cdot} \text{CS} \triangleright_{AD} \frac{\frac{\frac{\Pi_1}{\cdot} W^* \quad \frac{\Pi_2}{\cdot} W^*}{\cdot} \text{CS} \quad \frac{\text{Aux}_R}{\cdot} A^*}{\cdot} \text{RD}^* \frac{\text{Aux}_R}{\cdot} A^* R \text{RD}^*$ |
| | B | $\frac{\frac{\Pi \text{ Aux}_R}{\cdot} R \quad \text{Aux}_B}{\cdot} B \triangleright_{AD} \frac{\frac{\Pi}{\cdot} \quad \frac{\text{Aux}_B}{\cdot} WR}{\cdot} B \frac{\text{Aux}_R}{\cdot} R$ | $\frac{\frac{\Pi \text{ Aux}_R}{\cdot} R \quad \text{Aux}_B}{\cdot} B \triangleright_{AD} \frac{\frac{\frac{\Pi}{\cdot} W^* \quad \frac{\text{Aux}_B}{\cdot} A^*}{\cdot} B \quad \frac{\text{Aux}_R}{\cdot} A^*}{\cdot} \text{RD}^* \frac{\text{Aux}_R}{\cdot} A^* R \text{RD}^*$ |
| | W | $\frac{\frac{\Pi \text{ Aux}_R}{\cdot} R}{\cdot} W \triangleright_{AD} \frac{\frac{\Pi}{\cdot} W \quad \frac{\text{Aux}_R}{\cdot} W}{\cdot} R$ | $\frac{\frac{\Pi \text{ Aux}_R}{\cdot} R}{\cdot} W \triangleright_{AD} \frac{\frac{\Pi}{\cdot} W \quad \frac{\text{Aux}_R}{\cdot} W}{\cdot} R$ |

Legend:

CS : the case-study rule B : the bi-deduce rule W : The weakening of hypothesis rule
RD : the rewriting without error or the duplication of hypothesis rules A : any rule

Convention: Ascending rules in purple. Auxiliary sub-proofs are denoted by Aux.

Figure 5. Shapes of the proof transformations commuting descending and ascending rules on the trunk of an admissible proof.

sub-proof of an admissible proof Π is any proof anchored in auxiliary premises of the trunk of Π .

B. Commuting Ascending and Descending Rules

We design a set of proof transformations \triangleright_{AD} which can be used to commute any pair of descending and ascending rules occurring in the main trunk of an admissible proof: if Π_1 is an admissible proof then $\Pi_1 \triangleright_{AD} \Pi_2$ if Π_2 can be obtained from Π_1 using one of our proof commutations applied on a rule on the trunk of Π_1 . We can notice that in this case Π_2 is also an admissible proof since \triangleright_{AD} preserve admissible proofs. The shapes of our transformation steps \triangleright_{AD} , omitting the formulas to improve readability, are shown in Fig. 5 (detailed transformations are given in the long version [22]).

Lemma 1. *The relation \triangleright_{AD} terminates on any admissible proof Π . Moreover, if Π is an admissible proof irreducible w.r.t. \triangleright_{AD} then no ascending rule appears below a descending rule on the trunk of Π .*

We prove this by showing that admissibility is preserved by \triangleright_{AD} , and that some well-chosen numerical quantity $\text{Value}_{\triangleright_{AD}}(\Pi)$ strictly decreases after each proof transformation of the shape described in Fig. 5. See the long version [22] for a proof sketch and the definition of $\text{Value}_{\triangleright_{AD}}(\Pi)$.

C. Collapsing Proofs

We consider a proof of the main inductive step of a larger proof with induction hypothesis $\vec{u} \sim \vec{v}$. W.l.o.g., we assume that $\vec{u} \neq \vec{v}$ (otherwise, the hypothesis is a triviality). Our goal is to reduce the number of applications of the induction hypothesis

until there remains only one, through a sequence of proof transformation collapsing applications of the case-study rule.

A proof is (\vec{u}, \vec{v}) -collapsible if it is $(\vec{u}; \vec{v})$ -admissible and if the rewriting without error rule is the only descending rule appearing in its trunk. This represents the top of a normalized proof w.r.t. \triangleright_{AD} , where all descending rules have been moved down the proof-tree. We allow the (descending) rewriting without error rule to appear in collapsible proofs as new occurrences of this rule may be added by collapse proof transformations.

We design a set of collapsing proof transformations \triangleright_{col} that applies on the trunk of a collapsible proof, transforming it into another collapsible proof. For example, we show the shape of a key transformation for the case study rule in Fig. 4, which allows to merge two occurrences of the axiom rule above a case study into a single axiom rule. We refer the reader to the long version [22] for more details on the \triangleright_{col} transformations and a proof sketch of the next lemma.

Lemma 2. *Let $\vec{u} \sim_\varepsilon \vec{v}$ be an equivalence predicate such that $\vec{u} \neq \vec{v}$. Let Π be an (\vec{u}, \vec{v}) -collapsible proof where all occurrences of the axiom rule in the trunk are on $\vec{u} \sim_\varepsilon \vec{v}$. Then, the rewrite relation \triangleright_{col} terminates on Π and yields proofs with at most one application of this axiom rule in the trunk.*

D. Polynomial security

We now seek to find conditions on proof by induction over natural numbers that make it possible to derive polynomial security from it. In essence, we will adapt the proof in order to obtain better advantage upper-bounds. This is done in the key lemma presented below, which relies on a specialized version of induction that we introduce first.

Consider the following induction rule specialized to integers:

$$\frac{\begin{array}{l} \mathbb{E}'; \Theta, u n \sim_{x_\varepsilon} v n \vdash \\ u n, u' (n+1) \sim_{\varepsilon_{ih}} v n, v' (n+1) \\ \mathbb{E}; \Theta \vdash u_0 \sim_{\varepsilon_0} v_0 \end{array}}{\mathbb{E}; \Theta \vdash \tilde{v}(n : \text{bint}). \text{det}(n) \stackrel{\text{def}}{\Rightarrow} u n \sim_{\varepsilon n} v n} \text{N-IND}$$

where \mathbb{E}' declares x_ε and $n : \text{bint}$. As for u, v and ε , they are defined by recurrence:

$$\begin{aligned} u 0 &\stackrel{\text{def}}{=} u_0 & v 0 &\stackrel{\text{def}}{=} v_0 & \varepsilon 0 &\stackrel{\text{def}}{=} \varepsilon_0 \\ u (n+1) &\stackrel{\text{def}}{=} \langle u n, u' (n+1) \rangle & v (n+1) &\stackrel{\text{def}}{=} \langle v n, v' (n+1) \rangle \\ \varepsilon (n+1) &\stackrel{\text{def}}{=} \lambda t. (\varepsilon_{ih} \{x_\varepsilon \mapsto \varepsilon n\}) (t + t_{\langle \rangle}) \end{aligned} \quad (7)$$

and where the main inductive premise has already been simplified by removing the top-level pair in $u (n+1)$ and $v (n+1)$ (remark that we account for the time of simulating the pair computation in the advantage εn , where $t_{\langle \rangle}$ must bound the time needed to compute each application of $\langle \rangle$).

Theorem 1. *Let Π be an $(u n; v n)$ -admissible proof proving the inductive premise of the rule above where: the only occurrence of the axiom rule in the trunk is on the induction hypothesis $u n \sim_{\varepsilon n} v n$; and where the induction hypothesis and x_ε are only used in the trunk.*

Then there exists Π_{poly} proving the inductive case of N-IND with an improved upper-bound $\varepsilon_{ih}^{\text{negl}}$, i.e. Π_{poly} proves:

$$\mathbb{E}'; \Theta, u n \sim_{x_\varepsilon} v n \vdash u n, u' (n+1) \sim_{\varepsilon_{ih}^{\text{negl}}} v n, v' (n+1)$$

where $\varepsilon_{ih}^{\text{negl}}$ is such that the final advantage $\varepsilon^{\text{negl}}$, defined from $\varepsilon_{ih}^{\text{negl}}$ as described in Eq. (7), is such that for any \mathbb{M} where:

- Π is in $\varepsilon / \text{poly}_{\mathbb{M}}^{x_\varepsilon, n}$, i.e. (roughly) the advantage terms (resp. time and length terms) appearing in all auxiliary sub-proofs of Π are negligible (resp. polynomial) in \mathbb{M} , for any number of session n and adversarial time t which are polynomial in η .
- The initial advantage bound ε_0 and the cryptographic advantages $\varepsilon_{\text{CCA1}}$ and ε_{PRF} must be negligible for adversaries running in time polynomial in η .

Then the term $\varepsilon^{\text{negl}}$ is a negligible term w.r.t. n and \mathbb{M} , i.e. for any polynomials $P_n, P_t \in \mathbb{N}[\eta]$ bounding, resp., the number of sessions and the execution time of the adversary:

$$\mathbb{E}_\rho(\llbracket \varepsilon^{\text{negl}} n \rrbracket_{\mathbb{M}[n \mapsto P_n(\eta)]}^{\eta, P_t}) (P_t(\eta)) \in \text{negl}(\eta).$$

See the long version [22] for both the formal definition of $\varepsilon / \text{poly}_{\mathbb{M}}^{x_\varepsilon, n}$ and the proof sketch for this lemma.

We can check that this result applies to the proof of Section IV-A establishing the main inductive step of the security of PA. This shows that PA provides a polynomial-level of security, though the initial proof did not.

VI. CONCLUSION

The CCSA approach has already been studied and extended in many directions, but always for asymptotic analyses and estimations. In this paper, we propose the first concrete logic for this approach, providing precise advantages bounds inside

the proofs reasoning steps. Also, whereas it is of major interest in many papers to improve the tightness of the estimation of the adversary advantages, it is usually done in ad-hoc ways, and mainly for pen-and-paper proofs. To our knowledge, it is the first time that such precise advantages are managed in an automatic and generic way for the derivation of security proofs. As a future work, we plan to tackle the sizeable task of implementing our concrete security logic in Squirrel. Another line of future work could be to extend the class of admissible proofs, one possible extension is the proof with mutual inductions, that are out of scope for now.

REFERENCES

- [1] K. Bhargavan, A. Delignat-Lavaud, C. Fournet, A. Pironti, and P. Strub, “Triple handshakes and cookie cutters: Breaking and fixing authentication over TLS,” in *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2014, pp. 98–113.
- [2] D. Adrian, K. Bhargavan, Z. Durumeric, P. Gaudry, M. Green, J. A. Halderman, N. Heninger, D. Springall, E. Thomé, L. Valenta, B. VanderSloot, E. Wustrow, S. Z. Béguelin, and P. Zimmermann, “Imperfect forward secrecy: how diffie-hellman fails in practice,” *Commun. ACM*, vol. 62, no. 1, pp. 106–114, 2019.
- [3] V. Shoup, “OAEP reconsidered,” *J. Cryptol.*, vol. 15, no. 4, pp. 223–249, 2002.
- [4] M. Barbosa, G. Barthe, C. Doczkal, J. Don, S. Fehr, B. Grégoire, Y. Huang, A. Hülsing, Y. Lee, and X. Wu, “Fixing and mechanizing the security proof of fiat-shamir with aborts and dilithium,” in *CRYPTO (5)*, ser. Lecture Notes in Computer Science, vol. 14085. Springer, 2023, pp. 358–389.
- [5] M. Barbosa, G. Barthe, K. Bhargavan, B. Blanchet, C. Cremers, K. Liao, and B. Parno, “Sok: Computer-aided cryptography,” in *2021 IEEE Symposium on Security and Privacy (SP)*, 2021, pp. 777–795.
- [6] V. Shoup, “Sequences of games: a tool for taming complexity in security proofs,” Cryptology ePrint Archive, Paper 2004/332, 2004, <https://eprint.iacr.org/2004/332>. [Online]. Available: <https://eprint.iacr.org/2004/332>
- [7] M. Abadi and C. Fournet, “Private authentication,” *Theor. Comput. Sci.*, vol. 322, no. 3, pp. 427–476, 2004.
- [8] B. Blanchet, “A computationally sound mechanized prover for security protocols,” in *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2006, pp. 140–154.
- [9] G. Barthe, B. Grégoire, and S. Z. Béguelin, “Formal certification of code-based cryptographic proofs,” in *POPL*. ACM, 2009, pp. 90–101.
- [10] G. Barthe, B. Grégoire, S. Héraud, and S. Z. Béguelin, “Computer-aided security proofs for the working cryptographer,” in *CRYPTO*, ser. Lecture Notes in Computer Science, vol. 6841. Springer, 2011, pp. 71–90.
- [11] D. A. Basin, A. Lochbihler, and S. R. Sefidgar, “CryptHOL: Game-based proofs in higher-order logic,” *J. Cryptol.*, vol. 33, no. 2, pp. 494–566, 2020.
- [12] C. Abate, P. G. Haselwarter, E. Rivas, A. V. Muyllder, T. Winterhalter, C. Hritcu, K. Maillard, and B. Spitters, “Ssprover: A foundational framework for modular cryptographic proofs in coq,” in *CSF*. IEEE, 2021, pp. 1–15.
- [13] G. Bana and H. Comon-Lundh, “A computationally complete symbolic attacker for equivalence properties,” in *CCS*. ACM, 2014, pp. 609–620.
- [14] D. Baelde, S. Delaune, C. Jacomme, A. Koutsos, and S. Moreau, “An interactive prover for protocol verification in the computational model,” in *SP*. IEEE, 2021, pp. 537–554.
- [15] D. Baelde, S. Delaune, A. Koutsos, and S. Moreau, “Cracking the stateful nut: Computational proofs of stateful security protocols using the squirrel proof assistant,” in *CSF*. IEEE, 2022, pp. 289–304.
- [16] C. Cremers, C. Fontaine, and C. Jacomme, “A logic and an interactive prover for the computational post-quantum security of protocols,” in *SP*. IEEE, 2022, pp. 125–141.
- [17] D. Baelde, A. Koutsos, and J. Lallemand, “A higher-order indistinguishability logic for cryptographic reasoning,” in *LICS*, 2023, pp. 1–13.
- [18] M. Arapinis, L. I. Mancini, E. Ritter, and M. D. Ryan, “Analysis of privacy in mobile telephony systems,” *Int. J. Inf. Sec.*, vol. 16, no. 5, pp. 491–523, 2017.

- [19] M. Bellare, “Practice-oriented provable security,” in *Lectures on Data Security*, ser. Lecture Notes in Computer Science, vol. 1561. Springer, 1998, pp. 1–15.
- [20] M. Fischlin and A. Mittelbach, “An overview of the hybrid argument,” *IACR Cryptol. ePrint Arch.*, p. 88, 2021.
- [21] M. Bellare, A. Boldyreva, and S. Micali, “Public-key encryption in a multi-user setting: Security proofs and improvements,” in *EUROCRYPT*, ser. Lecture Notes in Computer Science, vol. 1807. Springer, 2000, pp. 259–274.
- [22] D. Baelde, C. Fontaine, A. Koutsos, G. Scerri, and T. Vignon, “A Probabilistic Logic for Concrete Security (long version of the present paper),” 2024. [Online]. Available: <https://hal.science/hal-04577828>
- [23] H. Comon and A. Koutsos, “Formal computational unlinkability proofs of RFID protocols,” in *CSF*. IEEE Computer Society, 2017, pp. 100–114.
- [24] A. Koutsos, “The 5G-AKA authentication protocol privacy,” in *EuroS&P*. IEEE, 2019, pp. 464–479.

APPENDIX A SEMANTICS OF PREDICATES

We give here the full definitions of the predicates described in Section II and in rules of Section III.

a) *Bounded length*: For u an order-1 term of type $\tau_0 \rightarrow \tau_1$, l a term of type $\overline{\text{int}} \rightarrow \overline{\text{int}}$, the predicate $\text{blen}_l(u)$ holds if, for any η and $a \in \llbracket \tau_0 \rrbracket_{\mathbb{M}}^{\eta}$ of length $n_a \in \mathbb{N}$, and for every $\rho \in \mathbb{T}_{\mathbb{M}, \eta}$, we have $|\llbracket u \rrbracket_{\mathbb{M}}^{\eta, \rho}(a)| \leq \inf_{\rho} \llbracket l \rrbracket_{\mathbb{M}}^{\eta, \rho}(n_a)$. The definition extends naturally to order-1 terms with more than one argument.

b) *Adversarial computability*: Let u be a term of type $\vec{\tau}_1 \rightarrow \vec{\tau}_0 \rightarrow \tau$ with $\vec{\tau}_1$ are types of order 1, $\vec{\tau}_0$ and τ have order 1. Let t be a term of type $(\overline{\text{int}} \rightarrow \overline{\text{int}})^n \rightarrow \overline{\text{int}}^m \rightarrow \overline{\text{int}}$, and \vec{o} of type $\overline{\text{int}}^n$, with $n = |\vec{\tau}_1|$ and $m = |\vec{\tau}_0|$. The predicate $\text{adv}_{t, \vec{o}}(u)$ holds if there exists a Turing machine \mathcal{A} with n oracles and m inputs such that, for all

- $\eta \in \mathbb{N}$ and $\rho \in \mathbb{T}_{\mathbb{M}, \eta}$,
- $\vec{f} \in \llbracket \vec{\tau}_1 \rrbracket_{\mathbb{M}}^{\eta, \rho}$ and \vec{l} such that $|f_i(x)| \leq l_i(|x|)$ for all i and x ,
- $\vec{w} \in \llbracket \vec{\tau}_0 \rrbracket_{\mathbb{M}}^{\eta, \rho}$ and \vec{s} such that $w_i = |s_i|$ for all i ,

we have:

$$\begin{aligned} \llbracket u \rrbracket_{\mathbb{M}}^{\eta, \rho}(\vec{f}, \vec{w}) &= \mathcal{A}^{\vec{f}}(1^{\eta}, \vec{w}, \rho_a) \\ \text{time}_{\mathcal{A}}(1^{\eta}, \vec{w}, \vec{f}) &\leq \inf_{\rho} \llbracket t \vec{l} \vec{s} \rrbracket_{\mathbb{M}}^{\eta, \rho} \\ \text{calls}_{\mathcal{A}}^{\eta} &\leq \inf_{\rho} \llbracket \vec{o} \rrbracket_{\mathbb{M}}^{\eta, \rho} \end{aligned}$$

c) *Well-foundedness*: Let τ be a type and $<_{\tau}$ be a symbol of type $\tau \rightarrow \tau \rightarrow \text{bool}$. We let $\text{well-founded}_{\tau}(<_{\tau})$ be the following global formula, which checks if type τ , ordered by $<_{\tau}$, is well-founded :

$$\text{well-founded}_{\tau}(<_{\tau}) \stackrel{\text{def}}{=} \llbracket \forall(l : \text{nat} \rightarrow \tau). \neg(\forall i, j. i < j \rightarrow l j <_{\tau} l i) \rrbracket_0$$

where nat and $< : \text{nat} \rightarrow \text{nat} \rightarrow \text{bool}$ are always interpreted as, resp., the set of natural numbers and the standard order over natural numbers.

APPENDIX B PROOF SYSTEM

We present some additional rules of the proof system in Fig. 6.

Mixed judgements rules

$$\begin{array}{c} \text{L}_{\varepsilon}.\text{BYGLOB} \\ \frac{\mathbb{E}; \Theta \vdash [\psi]_{\varepsilon}}{\mathbb{E}; \Theta; \emptyset \vdash_{\varepsilon} \psi} \end{array} \quad \begin{array}{c} \text{G}_{\varepsilon}.\text{BYLOC} \\ \frac{\mathbb{E}; \Theta; \emptyset \vdash_{\varepsilon} \psi}{\mathbb{E}; \Theta \vdash [\psi]_{\varepsilon}} \end{array} \quad \begin{array}{c} \text{L}_{\varepsilon}.\text{LOCALISE} \\ \frac{\mathbb{E}; \Theta; \Gamma, \phi \vdash_{\varepsilon_0} \psi}{\mathbb{E}; \Theta, [\phi]_{\varepsilon_1}; \Gamma \vdash_{\varepsilon_0 + \varepsilon_1} \psi} \end{array}$$

$$\begin{array}{c} \text{L}_{\varepsilon}.\text{REWRITE-EQUIV} \\ \frac{\mathbb{E}; \Theta; \Gamma_1 \vdash_{\varepsilon_1} \psi_1 \quad \mathbb{E}; \Theta \vdash (\Gamma_0 \Rightarrow \psi_0) \sim_{\varepsilon_0} (\Gamma_1 \Rightarrow \psi_1)}{\mathbb{E}; \Theta; \Gamma_0 \vdash_{\varepsilon_1 + \varepsilon_0(1)} \psi_0} \end{array}$$

Global judgement: local and global relations.

$$\begin{array}{c} \text{G}_{\varepsilon}.\text{L-LOC} : \Rightarrow \\ \frac{\mathbb{E}; \Theta, [\phi]_0 \dot{\Rightarrow} [\psi]_{\varepsilon} \vdash F \quad \mathbb{E}; \Theta \vdash \text{const}(\phi)}{\mathbb{E}; \Theta, [\phi \Rightarrow \psi]_{\varepsilon} \vdash F} \end{array} \quad \begin{array}{c} \text{G}_{\varepsilon}.\text{L-}\check{\forall}\text{-}\check{\forall} \\ \frac{\mathbb{E}; \Theta, \check{\forall}(x : \tau). [\psi]_{\varepsilon} \vdash F}{\mathbb{E}; \Theta, [\check{\forall}(x : \tau). \psi]_{\varepsilon} \vdash F} \end{array}$$

$$\begin{array}{c} \text{G}_{\varepsilon}.\text{L-}\check{\forall}\text{-}\check{\forall} \\ \frac{\mathbb{E}; \Theta, [\check{\forall}(x : \tau). \psi]_{\varepsilon} \vdash F}{\mathbb{E}; \Theta, \check{\forall}(x : \tau). [\psi]_{\varepsilon} \vdash F} \end{array}$$

Figure 6. Selected additional rules.

A. Bi-Deduction

The bi-deduction rule shown in Section III-B is a generalization of several reductionistic rules of our proof system. More precisely, this rule captures the Function Applications rules $\text{G}_{\varepsilon}.\text{E:FA-BASE}$ and the other variants (see the long version [22] for those), as well as the weakening rule.

We now present our full bi-deduction rule, which generalizes the rule presented in the body of this paper in several ways:

- we provide additional fresh names to the contexts;
- we allow for an arbitrary number of contexts instead of a single context.

These two extensions make our bi-deduction rule general enough to subsume more standard CCSA rules: adding names captures the $\text{G}_{\varepsilon}.\text{E:FA-NAMES}$ rule capabilities, while having many contexts captures the duplication rules $\text{G}_{\varepsilon}.\text{E:DUP-FUN}$ and the version with terms of order 0.

a) *Generalized bi-deduction*: We now describe our first generalization of the bi-deduction rule. Let \mathbb{E} be an environment such that any declared symbols x in \mathbb{E} is only used in eta-long form and is such that $\mathbb{E}; \Theta \vdash \text{adv}_{+\infty}(x)$ is valid, and let Θ be a set of global hypotheses.

Let $\vec{u} \stackrel{\text{def}}{=} u_1, \dots, u_m$ and $\vec{v} \stackrel{\text{def}}{=} v_1, \dots, v_m$ be two sequences of terms of the same length such that, for every i , terms u_i and v_i have the same types. Let $\vec{l} \stackrel{\text{def}}{=} l_1, \dots, l_m$ be a sequence of length terms for \vec{u} and \vec{v} :

- we assume that the types of \vec{l} are compatible with \vec{u} and \vec{v} , i.e. l_i has type $\overline{\text{int}}$ if u_i and v_i are of order 0, and has type $\overline{\text{int}} \rightarrow \overline{\text{int}}$ if u_i and v_i are of order 1;
- we will require $\text{blen}_{l_i}(u_i) \check{\wedge} \text{blen}_{l_i}(v_i)$ for every i .

Similarly, let \vec{w}_0 and \vec{w}_1 be two sequences of terms of the same length with compatible types, and let \vec{l}^w be a sequence of length terms for \vec{w}_0 and \vec{w}_1 .

Let C_1, \dots, C_n be sequences of terms representing the contexts that are to be simulated by the adversary. To account for the cost aspects of the rule, we consider, for every $i \leq n$:

- a term t_i representing the execution time of C_i ;
- a sequence of terms $\bar{\sigma}_i$ representing the number of calls that C_i makes to its arguments of order 1.

We will require that $\text{adv}_{t_i, \bar{\sigma}_i}(C_i)$ holds for every $i \leq n$, and that the terms C_1, \dots, C_n are *without names*.

Let $\vec{n} \stackrel{\text{def}}{=} n_1, \dots, n_p$ a sequence of names that are to be provided to C_1, \dots, C_n . We require that:

- names in \vec{n} are fresh, i.e. we require that no name in \vec{n} appears in $\vec{w}_0, \vec{w}_1, \vec{u}, \vec{v}, C_1, \dots, C_n$, and Θ ;
- names in \vec{n} do not appear in \mathbb{E} , except in their declarations;
- we are provided, for every $i \leq p$, with a term t_j^n upper-bounding the time needed by n_j to do a single sampling in any model;
- we are given a sequence of length terms $\vec{l}^n = l_1^n, \dots, l_p^n$ for \vec{n} , where we will require that $\text{blen}_{l_i^n}(n_i)$ holds for every $i \leq p$.

We let j_i, \dots, j_q be such that

$$(C_{j_1} \vec{w}_0 \vec{u} \vec{n}), \dots, (C_{j_q} \vec{w}_0 \vec{u} \vec{n})$$

is the sub-sequence, in the same order, of terms of order 1 in $(C_1 \vec{w}_0 \vec{u} \vec{n}), \dots, (C_n \vec{w}_0 \vec{u} \vec{n})$. Furthermore, for every context C_i , we decompose $\bar{\sigma}_i$ into $\bar{\sigma}_i^a, \bar{\sigma}_i^n$ where: $\bar{\sigma}_i^a$ bounds the number of calls from C_i to each order-1 argument in \vec{w}_0, \vec{u} (or \vec{w}_1, \vec{v}); $\bar{\sigma}_i^n$ is of length p and bounds the number of calls from C_i to each name symbol in \vec{n} .

Then, we have the rule:

G_ε.E:BI-DEDUCE

$$\frac{\begin{array}{l} \mathbb{E}; \Theta \vdash \vec{w}_0, \vec{u} \sim_\varepsilon \vec{w}_1, \vec{v} \quad \mathbb{E}; \Theta \vdash \bigwedge_{i \leq n} \text{adv}_{t_i, \bar{\sigma}_i}(C_i) \\ \mathbb{E}; \Theta \vdash \text{blen}_{\vec{l}^n}(\vec{w}_0) \tilde{\wedge} \text{blen}_{\vec{l}^n}(\vec{w}_1) \\ \mathbb{E}; \Theta \vdash \text{blen}_{\vec{l}}(\vec{u}) \tilde{\wedge} \text{blen}_{\vec{l}}(\vec{v}) \quad \mathbb{E}; \Theta \vdash \text{blen}_{\vec{l}^n}(\vec{n}) \end{array}}{\mathbb{E}; \Theta \vdash \vec{w}_0, (C_1, \dots, C_n) \vec{w}_0 \vec{u} \vec{n} \sim_{\varepsilon'} \vec{w}_1, (C_1, \dots, C_n) \vec{w}_1 \vec{v} \vec{n}}$$

where $(C_1, \dots, C_n) \vec{w}_0 \vec{u} \vec{n}$ denotes the sequence terms

$$C_1 \vec{w}_0 \vec{u} \vec{n}, \dots, C_n \vec{w}_0 \vec{u} \vec{n},$$

and ε' is the advantage bound term:

$$\lambda t, a_1, \dots, a_q. \varepsilon \left(t + t_{\text{oracle}} + \sum_{i \leq n} \text{cpt}_i \cdot t_i \vec{l}^n \vec{l}^n \right) \left(\sum_{i \leq n} \text{cpt}_i \cdot \bar{\sigma}_i^a \right) \quad (8)$$

where t_{oracle} is the time needed to compute all random samplings for names \vec{n} defined by:

$$t_{\text{oracle}} \stackrel{\text{def}}{=} \sum_{i \leq n, j \leq p} t_j^n \cdot \bar{\sigma}_i^n[j]$$

and where, for every $i \leq n$, the term cpt_i represents the number of times we need to simulate C_i , and is defined as:

$$\text{cpt}_i \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } (C_i \vec{u} \vec{n}) \text{ is of order 0} \\ a_r & \text{if } (C_i \vec{u} \vec{n}) \text{ is of order 1, where } r \text{ is s.t. } j_r = i. \end{cases}$$

In **C.NAME**, and t_n and l_n are upper-bounds on, resp., the time needed by n for a single sampling, and the length of a single sampling.

C.APP₀

$$\frac{\mathbb{E}; \Theta; \vec{\alpha} \vdash_{t_u, \bar{\sigma}}^c u :- l_u \quad \mathbb{E}; \vec{\alpha} \vdash v :- \tau \quad \tau \text{ of order 0} \quad u \notin \vec{\alpha}}{\mathbb{E}; \Theta; \vec{\alpha} \vdash_{t_u l_u + t_v, \bar{\sigma}}^c u v :- l_u l_v}$$

C.NAME

$$\frac{n \in \mathcal{N} \quad \mathbb{E}; \Theta; \vec{\alpha} \vdash_{t_i}^c i :- l}{\mathbb{E}; \Theta; \vec{\alpha} \vdash_{t_i + t_n}^c n i :- l_n}$$

C.ADV

$$\frac{\mathbb{E}; \Theta \vdash \text{adv}_{t, \bar{\sigma}}(u) \quad \mathbb{E}; \Theta \vdash \text{blen}_l(u)}{\mathbb{E}; \Theta; \emptyset \vdash_{t, \bar{\sigma}}^c u :- l}$$

C.LAMBDA₀

$$\frac{\mathbb{E}, l_x : \text{int}; \Theta; \vec{\alpha}, (x : \tau :- l_x) \vdash_{t, \bar{\sigma}}^c u :- l \quad \tau \text{ of order 0}}{\mathbb{E}; \Theta; \vec{\alpha} \vdash_{\lambda l_x. t, \bar{\sigma}}^c \lambda x. u :- \lambda l_x. l}$$

Figure 7. Rules for the computability predicate $\mathbb{E}; \Theta; \vec{\alpha} \vdash_{t, \bar{\sigma}}^c u$.

B. The Computability Predicate $\mathbb{E}; \Theta; \vec{\alpha} \vdash_{t, \bar{\sigma}}^c u$

For u a term of order at most 2, the computability predicate $\mathbb{E}; \Theta; \vec{\alpha} \vdash_{t, \bar{\sigma}}^c u$ represents the fact that $\llbracket u \rrbracket_{\mathbb{M}}^{\eta, \rho}$ can be simulated from $\vec{\alpha}$, in time t and calling its first order arguments at most $\bar{\sigma}$ times (where t and $\bar{\sigma}$ are well-type in \mathbb{E}). It is an extension of the \vdash_{pptm} predicate from [17]. Practically this predicate is similar to $\text{adv}_{t, \bar{\sigma}}(u)$ with two major differences:

- we allow for simulation of names, that is the distribution of the simulation should be equal to the distribution of $\llbracket u \rrbracket_{\mathbb{M}}^{\eta, \rho}$, as opposed to $\text{adv}()$ where we require equality for every ρ ;
- we fix the evaluation strategy for the simulation, i.e. the simulation evaluates u in a standard way, as opposed to $\text{adv}()$ where we only require that a simulator exists.

In order to be able to provide a proof system that allows for simulating quantifiers, we need a new predicates. For any type τ of order 0, we consider an additional global predicate $\text{enum}_\tau(t_e)$, which is satisfied in a model \mathbb{M} iff. there exists a machine $\mathcal{M} \in \text{PPTM}$ such that, for any η , \mathcal{M} enumerates all elements of type τ in time at most $\inf_\rho(\llbracket t_e \rrbracket_{\mathbb{M}}^{\eta, \rho})$. By enumerating, we mean that there exists a particular working tape T of \mathcal{M} such that, during an execution of \mathcal{M} , the tape T will successively contains all values of type τ , where \mathcal{M} enters a special state each time it emits a value of type τ on tape T (and \mathcal{M} enter this state only when it emits such a value).

We provide some selected rules for $\mathbb{E}; \Theta; \vec{\alpha} \vdash_{t, \bar{\sigma}}^c u$ in 7.