Cracking the Stateful Nut

Computational Proofs of Stateful Security Protocols using the SQUIRREL Proof Assistant

David Baelde*, Stéphanie Delaune*, Adrien Koutsos† and Solène Moreau*

* Univ Rennes, CNRS, IRISA, France

† Inria Paris, France

Index Terms—Security Protocols, Formal Methods, Logic, Computational Security, Interactive Theorem Proving.

Abstract—Bana and Comon have proposed a logical approach to proving protocols in the computational model, which they call the Computationally Complete Symbolic Attacker (CCSA). The proof assistant SQUIRREL implements a verification technique that elaborates on this approach, building on a meta-logic over the CCSA base logic. In this paper, we show that this meta-logic can naturally be extended to handle protocols with mutable states (key updates, counters, *etc.*) and we extend SQUIRREL's proof system to be able to express the complex proof arguments that are sometimes required for these protocols. Our theoretical contributions have been implemented in SQUIRREL and validated on a number of case studies, including a proof of the YubiKey and YubiHSM protocols.

I. INTRODUCTION

Formally specifying and verifying security protocols has become a major field of application for formal logic and automated reasoning. Decades of intensive research in that domain have lead to mature tools and industrial successes [1]–[3]. Some tools analyze protocols in so-called symbolic models [4]–[8] where attacker capabilities are represented by a fixed set of equations or inference rules over formal terms. Others deal with the cryptographer's computational model [9]–[11] where arbitrary probabilistic Turing machines compute over bitstrings. The symbolic models have enabled the earliest successes, generally allow for highly automated analyses, and can be used to discover actual attacks. However, it is very difficult to interpret a proof in a symbolic model in terms of computational attackers, as is witnessed by the limited practical impact of research on computational soundness [12]. Originally, many tools were designed to be fully automatic, e.g. [13], [14]. Newer tools, e.g. [4], [8], still have this objective but also incorporate interactivity, and even tools designed originally to be fully automated, e.g. PROVERIF, are moving towards it [15]. It is indeed an interesting compromise to tackle the analysis of complex security protocols. The computational model is more sought after than the symbolic one, as it provides stronger and more realistic security guarantees. But these come at a cost: it involves more complex notions and specific proof techniques, and mechanized proofs in that model often involve heavy user guidance. Yet there is no clear cut between

the two approaches, and CRYPTOVERIF [9] in particular can sometimes derive proofs completely automatically.

In this paper, we consider a recent approach to proving protocols in the computational model, embodied in the proof assistant SQUIRREL [16]. It implements a meta-logic built on top of the logic introduced by Bana and Comon [17], which follows an approach called the computationally complete symbolic attacker (CCSA). The CCSA approach assumes a fixed bound on protocol traces and relies on an explicit encoding of protocol traces as terms of the CCSA logic. It yields verification goals that are tedious to prove by hand and for which no automated prover exists so far. To circumvent these problems, SQUIRREL builds on a meta-logic which internalizes the encoding of traces as terms, and allows high-level proofs that are well-suited for interactive theorem proving. Moreover, these proofs hold for any trace, without a bound on their length. SQUIRREL and its meta-logic have enabled the first completely mechanized proofs of protocols using the CCSA approach [18].

The initial successes of the SQUIRREL system lie within a limited scope, in terms of the complexity of the protocols and of the proofs of their properties, and one naturally wonders how far the approach can be taken. In the present paper, we show how the SQUIRREL approach can be extended to support protocols with mutable states. Such situations are common: new keys may be derived regularly using cryptographic hash functions, counters or timestamps may be maintained to avoid replays, *etc*. At the semantic level, the meta-logic approach is easily adapted to incorporate states. At the level of proofs, however, significant improvements appear to be needed except for the simplest case studies.

- As we shall see, significant properties of stateful protocols often rely on a rich interplay between reachability and equivalence properties, which have to be proved simultaneously by mutual induction: *e.g.*, the strong secrecy (an equivalence property) of some state values at an instant T can be shown using the fact that these values never repeat (a reachability property), which itself relies on the strong secrecy of said state values at $T' < T$, *etc*.
- Moreover, because proofs of stateful protocols are more complex, they benefit more from automated deduction support. We found (bi-)deduction to be a useful notion in that regard. Roughly, the idea is that, if (part of) some messages of a protocol can be deduced by the adversary

from past messages, they can be safely ignored: indeed, they do not bring any new knowledge to the adversary.

Our main theoretical contribution is thus the introduction of a significant generalization of the sequent calculi of [18], allowing to capture security arguments mixing reachability and equivalence properties, and the design of a rich bi-deduction proof system, which can be used to automatically simplify some proof goals. On the practical side, we have implemented these improvements in an extension of SQUIRREL, and shown that they enable new case studies. In particular, we present proofs of the YubiKey and YubiHSM protocols, adapting in the computational model the TAMARIN development of [19]. Interestingly, the SQUIRREL development involves a similar amount of human effort than the TAMARIN analysis.

A full version of this paper including the omitted rules and proofs is available [20]. Our extension of SQUIRREL, as well as our case studies are available at:

https://squirrel-prover.github.io

Outline. We briefly recall the CCSA approach in Section II, before presenting in Section III the meta-logic given in [18] enriched to support mutable states, and sequences of terms. We show in Section IV some immediate applications of this extension, and also identify limitations of the proof system of [18]. This motivates the development of several improvements of this proof system, in Section V. We put everything together in Section VI where we present our case studies on the YubiKey and YubiHSM protocols. We conclude in Section VII.

II. BACKGROUND

The CCSA approach [17], [21] allows to reason about probabilistic polynomial-time computations using the simple symbolic setting of first-order logic. In a nutshell, terms are used to model computations, and a single predicate models indistinguishability of (sequences of) terms. Standard cryptographic assumptions can then be captured by axiom schemes, from which the indistinguishabilities of interest will have to be derived. In comparison with [17], we consider only the sort message, and a single attacker symbol att. For illustration purposes, we introduce below our running example.

Example 1. *We consider a variant of the OSK protocol [22], which is an RFID protocol.*

T updates its state: $s_T := H(s_T, k)$
$T \rightarrow R : G(s_T, k')$
$R \rightarrow T : ok$ *if* $G(s_T, k') = G(H(s_R, k), k')$ *with* $s_R \in DB$
R updates the database DB: $s_R := H(s_R, k)$

We consider two keyed hash functions H *and* G *and assume that they both satisfy the* Pseudo Random Function *(PRF [23]) assumption. We will always use* H *with the key* k*, and* G *with* k'*. Each tag is associated to a mutable state* s_T *initialized with a secret value* s^0 *and can perform the following action any number of times: update* s_T *with* $H(s_T, k)$*, output* $G(s_T, k')$*. We consider that readers have access to a database, containing an entry for each authorized tag. Each entry is*

initialized with the secret value s^0 of the corresponding tag. When receiving a message x, the reader looks up in the database for an entry s_R such that $x = G(H(s_R, k), k')$. Upon success, the reader accepts the message and updates the entry s_R with $H(s_R, k)$. Intuitively, readers knowing the tag's secrets will be able to recognize its outputs, but an attacker will not be able to learn these secrets and impersonate tags.

A. Syntax of the base logic

The base logic is a first-order logic, in which terms represent probabilistic PTIME Turing machines producing bitstrings, and a single predicate $\sim$ represents computational indistinguishability. A key idea of the CCSA approach is to use a special attacker function symbol att to represent the attacker's computations, which is left unspecified to model the fact that the attacker may perform any arbitrary probabilistic PTIME computation. The logic is parameterized by a set $\mathcal{N}_B$ of name symbols (modeling random samplings), a set of variables $\mathcal{X}_B$, and a set of function symbols $\mathcal{F}_B$ (modeling *e.g.* cryptographic primitives). Terms are generated from $\mathcal{X}_B$ and $\mathcal{N}_B$ using the unary function symbol att and the function symbols of $\mathcal{F}_B$. We assume that $\mathcal{F}_B$ contains at least the following symbols, with the expected arities and usual notations: pairing $\langle \_, \_ \rangle$, equality $EQ(\_, \_)$, conditionals if _ then _ else _, and the constants empty, true and false.

Atomic formulas are of the form $u_1, \ldots, u_n \sim v_1, \ldots, v_n$ where $n \geq 1$ and $u_1, \ldots, u_n, v_1, \ldots, v_n$ are terms. Such a formula intuitively expresses that the two sequences of messages are indistinguishable. We do not use a predicate symbol for equality in the base logic: $EQ(u, v)$ is a term and we may write, for instance, the term $EQ(true, EQ(u, v))$ or the formula $EQ(u, v) \sim true$.

Example 2. *To model the protocol of Example 1, we use one name symbol* s_i^0 *for each tag i, two name symbols* k, k' *for the keys, and two function symbols* H, G*. The term* $G(H(s_i^0, k), k')$ *represents the message outputted at the first session of the tag i,* $G(H(H(s_i^0, k), k), k')$ *the message outputted at the second session, etc. Assuming two different tags with respective initial secret values* s_1^0 *and* s_2^0*, the atomic formula*

$$G(H(s_1^0, k), k'), G(H(s_2^0, k), k')$$
$$\sim G(H(s_1^0, k), k'), G(H(H(s_1^0, k), k), k'),$$

expresses a simple unlinkability property: an outside observer cannot tell the difference between the outputs of two different tags and two successive outputs of the same tag.

B. Semantics of the base logic

We are interested in the interpretation of formulas of the base logic in a specific class of first-order interpretations, called *computational models*. The domain of a computational model $\mathbb{M}$ is the set of probabilistic PTIME Turing machines that run in polynomial-time w.r.t. a security parameter η and that manipulate a pair $\rho = (\rho_s, \rho_r)$ of infinite read-only random tapes. The tape ρ_s is used to draw honestly generated random values, and is not directly accessible by the attacker,

and ρ_r is used for random values drawn by the attacker. Given a valuation σ mapping variables in t to Turing machines in $\mathbb{M}$'s domain, the interpretation $[\![t]\!]^\sigma_\mathbb{M}$ of a term as a Turing machine is defined as follows (we omit σ when t is ground, and write $[\![t]\!]$ when $\mathbb{M}$ is clear from the context).

Variables in $\mathcal{X}_\mathcal{B}$ are interpreted using σ, and each name $\mathsf{n} \in \mathcal{N}_B$ is interpreted as a machine that extracts a word of length η from ρ_s, such that different names extract disjoint parts of the tape (this ensures that syntactically distinct names are interpreted as *independent* random variables).

The symbols empty, true, false, EQ, $\langle \_, \_ \rangle$ and if $\_$ then $\_$ else $\_$ are interpreted in the expected way. For instance, for any term t_1, t_2, $[\![\mathsf{EQ}(t_1, t_2)]\!]^\sigma_\mathbb{M}$ is the Turing machine that, on input $(1^\eta, \rho)$, returns 1 if $[\![t_1]\!]^\sigma_\mathbb{M}(1^\eta, \rho)$ and $[\![t_2]\!]^\sigma_\mathbb{M}(1^\eta, \rho)$ return the same result, and 0 otherwise. The other function symbols in $\mathcal{F}_B$ are interpreted as arbitrary PTIME Turing machines that do not access the random tapes. When studying a specific protocol, we will put additional restrictions on the computational models we consider, according to the assumptions the protocol relies on: we may assume *e.g.* that a binary function symbol $\oplus$ is interpreted as exclusive or, or that a binary function symbol H is interpreted as a cryptographically secure keyed hash function. The symbol att is interpreted as a PTIME Turing machine that does not access the random tape ρ_s, but has access to ρ_r.

Finally, the predicate $\sim$ is interpreted as the computational indistinguishability $\approx$, where two terms are computationally indistinguishable if no probabilistic PTIME machine can distinguish between them with non-negligible probability. More precisely, $d_1, \ldots, d_n \approx d'_1, \ldots, d'_n$ when for any PTIME Turing machine $\mathcal{A}$, the following quantity is negligible in η:

$$|\mathbf{Pr}(\rho : \mathcal{A}(d_1(1^\eta, \rho), \ldots, d_n(1^\eta, \rho), \rho_r) = 1) -$$
$$\mathbf{Pr}(\rho : \mathcal{A}(d'_1(1^\eta, \rho), \ldots, d'_n(1^\eta, \rho), \rho_r) = 1)|.$$

Here, $\mathbf{Pr}(\rho : \mathcal{M}(\rho) = 1)$ is the probability that the machine $\mathcal{M}$ accepts w.r.t. the random tapes in ρ. As it is standard with asymptotic security, a function is said to be negligible when it is asymptotically smaller than the inverse of any polynomial.

Given a computational model $\mathbb{M}$, a valuation σ and a formula ϕ, we write $\mathbb{M}, \sigma \models \phi$ when ϕ, with its variables interpreted by σ, is satisfied in the computational model $\mathbb{M}$. We say that ϕ is valid when $\mathbb{M}, \sigma \models \phi$ for all $\mathbb{M}$ and σ.

Example 3. *Let* n *and* m *be two distinct names. The formulas* $\mathsf{n} \sim \mathsf{m}$ *and* $\mathsf{EQ}(\mathsf{n}, \mathsf{m}) \sim \mathsf{false}$ *are valid, as no attacker can distinguish between two uniform random samplings of the same length, and two such samplings that are independent have a negligible probability of being equal.*

In this approach, a proof of a security property is done using an axiomatic approach, by deriving the formula expressing the given security property from a set of sound axioms. Some axioms are sound in all computational models, *e.g.* the symmetry of $\sim$, or properties of if $\_$ then $\_$ else $\_$. Other axioms reflect cryptographic assumptions on the security

primitives, such as the PRF assumption for hash functions. Such axioms exclude any computational model that does not satisfy the assumptions under which a security protocol has been designed.

C. Limitations of the base logic approach

As mentioned in the introduction, this approach has some limitations. First, it can only deal with bounded executions, *e.g.* the equivalence formula in Example 2 expresses the unlinkability of only two sessions of the protocol. Second, there is no support for proof mechanization. The work of [18] elaborates on the CCSA approach to solve these limitations, by designing a meta-logic over the base logic, a proof system for this meta-logic and by implementing them in a new interactive prover SQUIRREL. Nevertheless, the framework of [18] does not support stateful protocols. In the next section, we extend the meta-logic syntax and semantics with mutable states, and show in Section IV some direct applications of our extension.

III. META-LOGIC

We now introduce the syntax and the semantics of our meta-logic, which extends the meta-logic of [18] with memory cells, to model stateful protocols; and sequences, which are useful to reason about sets of terms. In order to match the intuitive notion of a trace, we also add a new atomic formula happens(T), expressing that T occurs in the execution.

A. Meta-logic syntax

Our meta-logic is a first-order logic with terms of three possible sorts: *timestamp* to represent time points in an execution; *message* to represent bitstrings exchanged between protocol's participants, or stored in the participants' states; and *index* which are used to identify an element, *e.g.* a session or an item in a database. We consider infinite sets $\mathcal{T}$, $\mathcal{X}$, and $\mathcal{I}$ of variables of each sort.

We assume a set $\mathcal{F}$ of indexed function symbols to represent protocol functions and cryptographic primitives (*e.g.* encryptions or keyed hash functions). Function symbols representing cryptographic primitives have index arity 0, and a message arity depending on the primitive. In contrast, function symbols representing identities have 0 as message arity and an arbitrary index arity. We also assume a set $\mathcal{N}$ of indexed name symbols to model random samplings, and a set of indexed action symbols $\mathcal{A}$ to model particular time points.

Given a meta-logic signature $\Sigma = (\mathcal{F}, \mathcal{N}, \mathcal{A}, \mathcal{S})$ and sets of variables $\mathcal{T}$, $\mathcal{X}$, and $\mathcal{I}$, we give in Figure 1 the syntax of meta-terms of sorts timestamp (noted T) and message (noted t), and of local meta-formulas (noted ϕ). Note that the notion of *local* meta-formulas will take all its meaning in Section V where we will introduce *global* meta-formulas. Meta-terms of sort index are restricted to variables $i \in \mathcal{I}$. We denote by $\mathsf{fv}(\_)$ the free variables (of any sort) of a meta-term or meta-formula.

Messages are modeled using terms of sort message, relying on *macros* to represent specific terms involved in the protocol. Since these values change throughout protocol execution,

$$
\begin{aligned}
T \;:=\;\; & \tau && \textit{(variable, } \tau \in \mathcal{T}) \\
\mid\;\; & \mathsf{a}[i_1,\ldots,i_k] && \textit{(action, } \mathsf{a} \in \mathcal{A}) \\
\mid\;\; & \mathsf{init} \mid \mathsf{pred}(T) \\[4pt]
t \;:=\;\; & x && \textit{(variable, } x \in \mathcal{X}) \\
\mid\;\; & \mathsf{s}[i_1,\ldots,i_k]@T && \textit{(memory cell, } \mathsf{s} \in \mathcal{S}) \\
\mid\;\; & \mathsf{n}[i_1,\ldots,i_k] && \textit{(name, } \mathsf{n} \in \mathcal{N}) \\
\mid\;\; & \mathsf{f}[i_1,\ldots,i_k](t_1,\ldots,t_n) && \textit{(function, } \mathsf{f} \in \mathcal{F}) \\
\mid\;\; & \mathsf{input}@T \mid \mathsf{output}@T \mid \mathsf{frame}@T && \textit{(builtin macros)} \\
\mid\;\; & \mathsf{if}\ \phi\ \mathsf{then}\ t\ \mathsf{else}\ t' && \textit{(conditional)} \\
\mid\;\; & \mathsf{find}\ \vec{i}\ \mathsf{suchthat}\ \phi\ \mathsf{in}\ t\ \mathsf{else}\ t' && \textit{(index lookup)} \\
\mid\;\; & \mathsf{seq}(\vec{\alpha} : t) && \textit{(sequence of terms)} \\[4pt]
\alpha \;:=\;\; & i \mid \tau && \textit{(variable, } i \in \mathcal{I}, \tau \in \mathcal{T}) \\[4pt]
A \;:=\;\; & t = t' \mid i = i' \mid T = T' \mid T < T' \mid T \le T' \\
\mid\;\; & \mathsf{happens}(T) \mid \mathsf{cond}@T \mid \mathsf{exec}@T \\[4pt]
\phi \;:=\;\; & A \mid \top \mid \bot \mid \phi \wedge \phi' \mid \phi \vee \phi' \mid \phi \Rightarrow \phi' \mid \neg\phi \\
\mid\;\; & \forall i.\phi \mid \exists i.\phi \mid \forall \tau.\phi \mid \exists \tau.\phi
\end{aligned}
$$

Syntactic sugar: $(\mathsf{if}\ \phi\ \mathsf{then}\ t) \stackrel{\text{def}}{=} \mathsf{if}\ (\phi\ \mathsf{then}\ t\ \mathsf{else}\ \mathsf{empty})$ and $\mathsf{seq}(\vec{\alpha} : t \mid \phi) \stackrel{\text{def}}{=} \mathsf{seq}(\vec{\alpha} : \mathsf{if}\ \phi\ \mathsf{then}\ t)$.

Fig. 1. Syntax of meta-terms and local meta-formulas.

macros are applied to a timestamp term that indicates when the macro is evaluated. Our logic features several predefined macros: $\mathsf{input}@T$ and $\mathsf{output}@T$ refer to the input and output messages of the action executed at time point T; $\mathsf{cond}@T$ and $\mathsf{exec}@T$ encode, resp., the condition of the action T, and the conjunction of all the conditions to reach T; and $\mathsf{frame}@T$ represents all the messages known by the attacker at time point T. The special timestamp constant init stands for the initial time point, and we rely on pred to denote the predecessor of a given time point T. Finally, pred^n is used as a shortcut for n sucessive applications of pred. We assume a set $\mathcal{S}$ of indexed state macro symbols to represent memory cells. For instance, if s is a state macro symbol of arity 1, the message $\mathsf{s}[i]@T$ refers to the contents of $\mathsf{s}[i]$ at time T. Lookups generalize conditionals: $\mathsf{find}\ \vec{i}\ \mathsf{suchthat}\ \phi\ \mathsf{in}\ t\ \mathsf{else}\ t'$ evaluates to t where indices $\vec{i}$ are bound to values such that ϕ holds if such values exist, and t' otherwise.

Lastly, we often use sequences when reasoning on protocols, *e.g.* to enrich equivalences with collections of messages that are already known to the attacker or can safely be disclosed as far as the property being proved is concerned. Once meta-interpreted in a given trace model, a sequence $\mathsf{seq}(\vec{\alpha} : t)$ is simply a list built as nested pairs. However, we need to reason on these objects in the meta-logic, uniformly for all trace models, which requires the use of a specific construct.

Example 4. *In Example 2, we explained how to model the messages outputted by the tag in the base logic. In particular, we had to consider as many names s_i^0 as tag's identities. In the meta-logic, we can use indexed names to model unbounded collections of objects. We thus assume a name $\mathsf{s}^0 \in \mathcal{N}$*

of index arity 1, such that $\mathsf{s}^0[i]$ models the initial secret value associated to the tag i. We also take two state macro symbols, s_T and s_R of index arity 1, modeling the current value of the memory cell and the database entry associated to a tag. The message outputted by a tag i at time T is then modeled by the meta-term $\mathsf{G}(\mathsf{H}(\mathsf{s}_\mathsf{T}[i]@\mathsf{pred}(T), k), k')$. In this message, $\mathsf{s}_\mathsf{T}[i]@\mathsf{pred}(T)$ is the value of the memory cell just before time point T, i.e. a (possibly empty) stack of hashes $\mathsf{H}(\ldots(\mathsf{H}(\mathsf{s}^0[i]), k), \ldots), k)$. We finally model as a (local) meta-formula the condition under which a reader executing at time T' will accept the input as coming from tag i:

$$
\phi_{\mathsf{reader}}^{T'} \stackrel{\text{def}}{=} \mathsf{input}@T' = \mathsf{G}(\mathsf{H}(\mathsf{s}_\mathsf{R}[i]@\mathsf{pred}(T'), k), k').
$$

The example below illustrates the use of the $\mathsf{seq}(\vec{\alpha} : t)$ construct.

Example 5. *Continuing our running example, we may want to consider the sequence of terms corresponding to the successive values stored in the memory cell $\mathsf{s}_\mathsf{T}[i]$ of each tag i during a valid execution. Such a sequence can be represented by the meta-term:*

$$
\mathsf{seq}(\tau, i : \mathsf{s}_\mathsf{T}[i]@\tau \mid \mathsf{happens}(\tau) \wedge \mathsf{exec}@\tau)
$$

If we only want to consider the successive values stored in the memory cell of the tag i_0, we will consider the meta-term

- $\mathsf{seq}(\tau : \mathsf{s}_\mathsf{T}[i_0]@\tau \mid \mathsf{happens}(\tau) \wedge \mathsf{exec}@\tau)$*; or*
- $\mathsf{seq}(\tau, i : \mathsf{s}_\mathsf{T}[i]@\tau \mid i = i_0 \wedge \mathsf{happens}(\tau) \wedge \mathsf{exec}@\tau)$

B. Protocols

Our local meta-logic formulas are interpreted as boolean terms of the base logic, depending on an execution of the protocol under study. We model a protocol as a finite set of actions, each action representing a basic step of the protocol where: the attacker provides an input; a condition is checked; some updates are performed; and finally, an output is emitted.

Definition 1. *An action*

$$
\mathsf{a}[\vec{i}].(\phi_{\mathsf{a}[\vec{i}]}, o_{\mathsf{a}[\vec{i}]}, \{\mathsf{s}[\vec{j}] \leftarrow u_{\mathsf{a}[\vec{i}],\mathsf{s}[\vec{j}]} \mid \mathsf{s} \in \mathcal{S}\})
$$

is formed from an action symbol a, distinct index variables $\vec{i}$, a local meta-logic formula $\phi_{\mathsf{a}[\vec{i}]}$, a meta-logic term $o_{\mathsf{a}[\vec{i}]}$ of sort message, and for each state macro symbol $\mathsf{s} \in \mathcal{S}$ and distinct index variables $\vec{j}$ (disjoint from $\vec{i}$), a meta-logic term $u_{\mathsf{a}[\vec{i}],\mathsf{s}[\vec{j}]}$ of sort message. We require that

$$
\mathsf{fv}(\phi_{\mathsf{a}[\vec{i}]}, o_{\mathsf{a}[\vec{i}]}) \subseteq \{\vec{i}\} \quad \text{and} \quad \mathsf{fv}(u_{\mathsf{a}[\vec{i}],\mathsf{s}[\vec{j}]}) \subseteq \{\vec{i}, \vec{j}\}.
$$

The formula $\phi_{\mathsf{a}[\vec{i}]}$ is called the condition of the action, $o_{\mathsf{a}[\vec{i}]}$ its output and $\{u_{\mathsf{a}[\vec{i}],\mathsf{s}[\vec{j}]} \mid \mathsf{s} \in \mathcal{S}\}$ its state update terms.

An action $\mathsf{a}[\vec{i}].(\phi_{\mathsf{a}[\vec{i}]}, o_{\mathsf{a}[\vec{i}]}, \{\mathsf{s}[\vec{j}] \leftarrow u_{\mathsf{a}[\vec{i}],\mathsf{s}[\vec{j}]} \mid \mathsf{s} \in \mathcal{S}\})$ models that, if $\phi_{\mathsf{a}[\vec{i}]}$ holds, $o_{\mathsf{a}[\vec{i}]}$ may be emitted, after having updated each memory cell $\mathsf{s}[\vec{j}]$ with the value $u_{\mathsf{a}[\vec{i}],\mathsf{s}[\vec{j}]}$. Conditional branching may be modeled using two actions, the condition of the second action being the negation of the first action's condition.

A protocol is a set of actions equipped with a dependency relation, which constrains the order of execution of actions. We require that an action only refers to timestamps of actions

occurring strictly before it, except when referring to its own input or to the value of a state macro at the current time if it is in the output term (*i.e.* after all updates have been performed).

Definition 2. *Given a finite set $\mathcal{A}$ of action symbols and a finite set $\mathcal{S}$ of state macro symbols, a protocol $\mathcal{P} = (\mathcal{P}_{\mathsf{act}}, \mathcal{U}_0, <)$ over $(\mathcal{A}, \mathcal{S})$ is a finite set $\mathcal{U}_0 \stackrel{def}{=} \{ \mathsf{s}[\vec{j}] \leftarrow u_{0,\mathsf{s}[\vec{j}]} \mid \mathsf{s} \in \mathcal{S} \}$ of initial state values, one for each state macro symbol, and a finite set $\mathcal{P}_{\mathsf{act}}$ of actions, one for each action symbol, equipped with a partial order $<$ over terms of the form $\mathsf{a}[\vec{i}]$ with $\mathsf{a} \in \mathcal{A}$. We require the following conditions.*

- *For each memory cell $\mathsf{s}[\vec{j}]$, $u_{0,\mathsf{s}[\vec{j}]}$ is a macro-free meta-logic term of sort message with free variables among $\vec{j}$.*
- *The partial order $<$ is insensitive to the choice of specific index variables, i.e. we have $\mathsf{a}_1[\vec{i_1}] < \mathsf{a}_2[\vec{i_2}]$ if, and only if, $\mathsf{a}_1[\sigma(\vec{i_1})] < \mathsf{a}_2[\sigma(\vec{i_2})]$ for any $\mathsf{a}_1, \mathsf{a}_2, \vec{i_1}$ and $\vec{i_2}$ and for any bijective variable renaming $\sigma : \mathcal{I} \to \mathcal{I}$.*
- *Actions only refer to information about previously executed actions. For every action $\mathsf{a}[\vec{i}].(\_) \in \mathcal{P}_{\mathsf{act}}$:*
  - *each subterm T of sort timestamp occurring in the condition $\phi_{\mathsf{a}[\vec{i}]}$ or in an update $u_{\mathsf{a}[\vec{i}], \mathsf{s}[\vec{j}]}$ is either (i) of the form $\mathsf{a}_0[\vec{i_0}]$ with $\mathsf{a}_0[\vec{i_0}] < \mathsf{a}[\vec{i}]$, or (ii) of the form $\mathsf{pred}^n(\mathsf{a}[\vec{i}])$ with $n \geq 1$, or (iii) of the form $\mathsf{a}[\vec{i}]$ and is applied to an input macro,*
  - *each subterm T of sort timestamp occurring in $o_{\mathsf{a}[\vec{i}]}$ either verifies (i), (ii) or (iii); or (iv) is of the form $\mathsf{a}[\vec{i}]$ and is applied to a state macro.*

Example 6. *Coming back to our running example, let $T \in \mathcal{A}$ be an action symbol of index arity 2. The action $T[i,j].(\_)$ defined below models the output performed by the tag, as well as the way the memory cells are updated. The memory cell $\mathsf{s}_\mathsf{T}[i]$ is updated with $\mathsf{H}(\mathsf{s}_\mathsf{T}[i]@\mathsf{pred}(T[i,j]), k)$, whereas $\mathsf{s}_\mathsf{T}[i']$ with $i' \neq i$, and $\mathsf{s}_\mathsf{R}[i']$, are left unchanged.*

$$T[i,j].(\mathsf{true}, \mathsf{G}(\mathsf{H}(\mathsf{s}_\mathsf{T}[i]@\mathsf{pred}(T[i,j]), k), k'),$$
$$\{ \mathsf{s}_\mathsf{T}[i'] \leftarrow \text{if } i' = i \text{ then } \mathsf{H}(\mathsf{s}_\mathsf{T}[i]@\mathsf{pred}(T[i,j]), k)$$
$$\text{else } \mathsf{s}_\mathsf{T}[i']@\mathsf{pred}(T[i,j]),$$
$$\mathsf{s}_\mathsf{R}[i'] \leftarrow \mathsf{s}_\mathsf{R}[i']@\mathsf{pred}(T[i,j]) \})$$

We similarly define the action $R[j,i].(\_)$ of a reader session j accepting tag i, and $R_1[j].(\_)$ of a reader session j rejecting its input. We take $\phi_{R[j,i]} \stackrel{def}{=} \phi_{\mathsf{reader}}^{R[j,i]}$ from Example 4, and $o_{R[j]} \stackrel{def}{=}$ ok. It remains to model the updates:

$$\mathsf{s}_\mathsf{R}[i'] \leftarrow \text{if } i' = i \text{ then } \mathsf{H}(\mathsf{s}_\mathsf{R}[i]@\mathsf{pred}(R[j,i]), k)$$
$$\text{else } \mathsf{s}_\mathsf{R}[i']@\mathsf{pred}(R[j,i])$$
$$\mathsf{s}_\mathsf{T}[i'] \leftarrow \mathsf{s}_\mathsf{T}[i']@\mathsf{pred}(R[j,i])$$

For the failure action $R_1[j]$ we take $\phi_{R_1[j]} \stackrel{def}{=} \forall i. \neg \phi_{\mathsf{reader}}^{R_1[j]}$, $o_{R_1[j]} \stackrel{def}{=}$ empty, and updates that leave cells unchanged. To complete the protocol modeling our running example we initialize cells with s^0 and specify an empty partial order:

$$\mathcal{P}_{\mathsf{ex}} \stackrel{def}{=} (\{ T[i,j].(\_), R[j,i].(\_), R_1[j].(\_) \},$$
$$\{ \mathsf{s}_\mathsf{T}[i] \leftarrow \mathsf{s}^0[i], \mathsf{s}_\mathsf{R}[i] \leftarrow \mathsf{s}^0[i] \}, \emptyset).$$

C. Semantics

To give a semantics to meta-terms and local meta-formulas, we translate them to terms of the base logic. This translation depends on the number of protocol agents and sessions, and on the interleaving of the protocol actions (*i.e.* the order in which the adversary interacts with the different protocol agents). This information is going to be encapsulated in the notion of *trace model*, which we are now going to define.

1) Protocol execution: We can instantiate the indices of an action by values to yield *concrete actions*.

Definition 3. *Given a set $\mathcal{A}$ of action symbols, a* concrete action *is the application of an action symbol $\mathsf{a} \in \mathcal{A}$ to k integers (k is the index arity of a). We lift the partial order of a protocol $\mathcal{P}$ to concrete actions. For any $\sigma : \mathcal{I} \to \mathbb{N}$, we have that $\mathsf{a}[\sigma(i_1), \ldots, \sigma(i_k)] < \mathsf{b}[\sigma(j_1), \ldots, \sigma(j_l)]$ when $\mathsf{a}[i_1, \ldots, i_k] < \mathsf{b}[j_1, \ldots, j_l]$.*

We define next the possible *interleavings* of actions for a given protocol, which over-approximates the actual possible executions by taking only dependency constraints into account.

Definition 4. *Given a protocol $\mathcal{P}$, an* interleaving *is a sequence of concrete actions $\alpha_1 \ldots \alpha_n$ in which no concrete action occurs twice, and such that, for every $1 \leq i \leq n$, for every concrete action β such that $\beta < \alpha_i$, there exists $1 \leq j < i$ such that $\beta = \alpha_j$.*

We consider meta-logic terms and formulas over $\Sigma = (\mathcal{F}, \mathcal{N}, \mathcal{A}, \mathcal{S})$, and given a finite set D of integers, the associated base logic signature $\Sigma^D = (\mathcal{F}_B, \mathcal{N}_B)$ contains exactly a name symbol $\mathsf{n}_{k_1, \ldots, k_p}$ for every $\mathsf{n} \in \mathcal{N}$ of index arity p, and every $k_1, \ldots, k_p \in D$; and a function symbol $\mathsf{f}_{k_1, \ldots, k_p}$ of arity n for every $\mathsf{f} \in \mathcal{F}$ of index arity p and message arity n, and every $k_1, \ldots, k_p \in D$.

In order to interpret meta-terms and local meta-formulas, we introduce the notion of trace model. The idea is that for each interleaving of the actions of the protocol under study, we can define a structure that will allow us to give a meaning to the macros.

Definition 5. *A trace model $\mathbb{T}$ (of a protocol $\mathcal{P}$) is a tuple $(\mathcal{D}_\mathcal{I}, \mathcal{D}_\mathcal{T}, <_\mathcal{T}, \sigma_\mathcal{I}, \sigma_\mathcal{T})$ such that:*

- *$\mathcal{D}_\mathcal{I} \subseteq \mathbb{N}$ is a finite index domain;*
- *$\mathcal{D}_\mathcal{T}$ contains two special symbols init, undef, and a subset of $\mathcal{D}_a \stackrel{def}{=} \{ \mathsf{a}[k_1, \ldots, k_n] \mid \mathsf{a} \in \mathcal{A}, k_1, \ldots, k_n \in \mathcal{D}_\mathcal{I} \}$;*
- *$<_\mathcal{T}$ is a total ordering on $\mathcal{D}_\mathcal{T} \smallsetminus \{\mathsf{undef}\}$ such that init is minimal, and such that the sequence of elements of $\mathcal{D}_\mathcal{T} \smallsetminus \{\mathsf{undef}\}$ ordered by $<_\mathcal{T}$ is an interleaving of $\mathcal{P}$;*
- *$\sigma_\mathcal{I} : \mathcal{I} \to \mathcal{D}_\mathcal{I}$ and $\sigma_\mathcal{T} : \mathcal{T} \to \mathcal{D}_\mathcal{T}$ are mappings that interpret index and timestamp variables as elements of their respective domains.*

Given a trace model, we define a predecessor function $\mathsf{pred}_\mathcal{T} : \mathcal{D}_\mathcal{T} \to \mathcal{D}_\mathcal{T}$ which maps undef to itself, init to undef, and all other elements v to the largest element v' such that $v' <_\mathcal{T} v$. Moreover, we denote by $\leq_\mathcal{T}$ the reflexive closure of $<_\mathcal{T}$ on $\mathcal{D}_\mathcal{T} \smallsetminus \{\mathsf{undef}\}$. When $\mathbb{T} = (\mathcal{D}_\mathcal{I}, \mathcal{D}_\mathcal{T}, <_\mathcal{T}, \sigma_\mathcal{I}, \sigma_\mathcal{T})$

is a trace model and $k \in \mathcal{D}_\mathcal{I}$, $\mathbb{T}\{i \to k\}$ is the trace model identical to $\mathbb{T}$ in which $\sigma_\mathcal{I}$ is updated to map i to k. We similarly define $\mathbb{T}\{\tau \to v\}$ when $v \in \mathcal{D}_\mathcal{T}$.

Example 7. *Consider the protocol $\mathcal{P}_{\mathsf{ex}}$ introduced in Example 6. A possible trace model $\mathbb{T}$ associated to this protocol is the tuple $(\mathcal{D}_\mathcal{I}, \mathcal{D}_\mathcal{T}, <_\mathcal{T}, \sigma_\mathcal{I}, \sigma_\mathcal{T})$ where:*

- $\mathcal{D}_\mathcal{I} \stackrel{def}{=} \{1, 2, 3\}$;
- $\mathcal{D}_\mathcal{T} \stackrel{def}{=} \{\mathsf{undef}, \mathsf{init}, T[1,3], T[2,1], T[2,2], R[1,2], R_1[3]\}$;
- $<_\mathcal{T}$ *is a total ordering such that:*
  $\mathsf{init} < T[2,1] < T[1,3] < R[1,2] < T[2,2] < R_1[3]$;
- $\sigma_\mathcal{I}$ *and $\sigma_\mathcal{T}$ are mappings that interpret index and timestamp variables to $\mathcal{D}_\mathcal{I}$ and $\mathcal{D}_\mathcal{T}$ respectively. For illustration purposes, we consider that $\sigma_\mathcal{I}(i) = \sigma_\mathcal{I}(j) = 2$.*

We have $\mathsf{pred}_\mathcal{T}(R_1[3]) = T[2,2]$, $\mathsf{pred}_\mathcal{T}(T[2,2]) = R[1,2]$, and so on.

2) Translation: We can now give our translation for meta-terms and local meta-formulas. This translation is similar to the one introduced in [18], but adapted to our notion of trace model in which some actions are allowed to not happen. We recall its general principle while highlighting the main differences with the translation given in [18].

The translation $(\_)_\mathcal{P}^\mathbb{T}$ is parameterized by a protocol $\mathcal{P}$ and a trace model $\mathbb{T}$ of $\mathcal{P}$. First, as expected, we have that $(\tau)_\mathcal{P}^\mathbb{T} = \sigma_\mathcal{T}(\tau)$, $(\mathsf{init})_\mathcal{P}^\mathbb{T} = \mathsf{init}$, and $(\mathsf{pred}(T))_\mathcal{P}^\mathbb{T} = \mathsf{pred}_\mathcal{T}((T)_\mathcal{P}^\mathbb{T})$. Then, each meta-logic construct is translated using its counterpart in the base logic when it is available. A name $\mathsf{n}[i_1, \ldots, i_p]$ is translated into $\mathsf{n}_{\sigma_\mathcal{I}(i_1), \ldots, \sigma_\mathcal{I}(i_p)}$. For some constructions, it is a bit more complicated. For instance, the sequence construct is not available in the base logic, but is translated using nested pairs. As an example, we show the interpretation of a sequence over a single index variable i.

$$(\mathsf{seq}(i: t))_\mathcal{P}^\mathbb{T} \stackrel{def}{=} \langle (t)_\mathcal{P}^{\mathbb{T}[i \mapsto k_1]}, \langle \ldots, (t)_\mathcal{P}^{\mathbb{T}[i \mapsto k_n]} \rangle \ldots \rangle$$

where $k_1, \ldots, k_n$ is an enumeration of $D_\mathcal{I}$. Similarly, the interpretation of (local) meta-formulas is quite straightforward using finite boolean expressions to translate quantifications over index and timestamp variables. For instance, we have that:

$$(\forall i.\phi)_\mathcal{P}^\mathbb{T} \quad \stackrel{def}{=} \quad \dot{\wedge}_{k \in D_\mathcal{I}} \, (\phi)_\mathcal{P}^{\mathbb{T}\{i \mapsto k\}}$$

Then, regarding atomic meta-formulas involving timestamps, we have that:

- $(T = T')_\mathcal{P}^\mathbb{T} = \mathsf{true}$ iff $(T)_\mathcal{P}^\mathbb{T} = (T')_\mathcal{P}^\mathbb{T}$;
- $(T < T')_\mathcal{P}^\mathbb{T} = \mathsf{true}$ iff $(T)_\mathcal{P}^\mathbb{T} <_\mathcal{T} (T')_\mathcal{P}^\mathbb{T}$;
- $(T \leq T')_\mathcal{P}^\mathbb{T} = \mathsf{true}$ iff $(T)_\mathcal{P}^\mathbb{T} \leq_\mathcal{T} (T')_\mathcal{P}^\mathbb{T}$;
- $(\mathsf{happens}(T))_\mathcal{P}^\mathbb{T} = \mathsf{true}$ iff $(T)_\mathcal{P}^\mathbb{T} \in \mathcal{D}_\mathcal{T} \setminus \{\mathsf{undef}\}$.

Lastly, we give in Figure 2 the interpretation of macros. Roughly, an output macro is replaced by the meta-term as specified by the protocol, and is then interpreted in the trace model $\mathbb{T}$. The cond macro has a similar treatment and produces a base logic formula corresponding to the condition of the action. The macro exec is translated as the conjunction of all

$$(\mathsf{s}[\vec{l}]@T)_\mathcal{P}^\mathbb{T} \stackrel{def}{=}$$

$$\begin{cases} \mathsf{empty} & \text{if } (T)_\mathcal{P}^\mathbb{T} = \mathsf{undef} \\ (u_{0,\mathsf{s}[\vec{j}]})_\mathcal{P}^{\mathbb{T}\{\vec{j} \mapsto \vec{k}_0\}} & \text{if } (T)_\mathcal{P}^\mathbb{T} = \mathsf{init}, (\vec{l})_\mathcal{P}^\mathbb{T} = \vec{k}_0 \\ & \text{and } (\mathsf{s}[\vec{j}] \leftarrow u_{0,\mathsf{s}[\vec{j}]}) \in \mathcal{U}_0 \\ (u_{\mathsf{a}[\vec{i}],\mathsf{s}[\vec{j}]})_\mathcal{P}^{\mathbb{T}\{\vec{i} \mapsto \vec{k}, \vec{j} \mapsto \vec{k}_0\}} & \text{if } (T)_\mathcal{P}^\mathbb{T} = \mathsf{a}[\vec{k}] \in \mathcal{D}_\mathcal{T}, (\vec{l})_\mathcal{P}^\mathbb{T} = \vec{k}_0 \end{cases}$$

$$(m@T)_\mathcal{P}^\mathbb{T} \stackrel{def}{=}$$

$$\begin{cases} \mathsf{empty} & \text{if } (T)_\mathcal{P}^\mathbb{T} = \mathsf{undef} \\ & \text{and } m \in \{\mathsf{input}, \mathsf{output}, \mathsf{frame}\} \\ \mathsf{false} & \text{if } (T)_\mathcal{P}^\mathbb{T} = \mathsf{undef} \text{ and } m \in \{\mathsf{cond}, \mathsf{exec}\} \\ (m_\mathsf{init})_\mathcal{P}^\mathbb{T} & \text{if } (T)_\mathcal{P}^\mathbb{T} = \mathsf{init} \\ (m_{\mathsf{a}[\vec{i}]})_\mathcal{P}^{\mathbb{T}\{\vec{i} \mapsto \vec{k}\}} & \text{if } (T)_\mathcal{P}^\mathbb{T} = \mathsf{a}[\vec{k}] \in \mathcal{D}_\mathcal{T} \end{cases}$$

where m_init and $m_{\mathsf{a}[\vec{i}]}$ are defined as follows:

$$\mathsf{cond}_\mathsf{init} = \mathsf{exec}_\mathsf{init} = \mathsf{true}$$

$$\mathsf{input}_\mathsf{init} = \mathsf{frame}_\mathsf{init} = \mathsf{output}_\mathsf{init} = \mathsf{empty}$$

$$\mathsf{output}_{\mathsf{a}[\vec{i}]} = o_{\mathsf{a}[\vec{i}]} \quad \mathsf{exec}_{\mathsf{a}[\vec{i}]} = \mathsf{cond}@\mathsf{a}[\vec{i}] \wedge \mathsf{exec}@\mathsf{pred}(\mathsf{a}[\vec{i}])$$

$$\mathsf{cond}_{\mathsf{a}[\vec{i}]} = \phi_{\mathsf{a}[\vec{i}]} \quad \mathsf{input}_{\mathsf{a}[\vec{i}]} = \mathsf{att}\big(\mathsf{frame}@\mathsf{pred}(\mathsf{a}[\vec{i}])\big)$$

$$\mathsf{frame}_{\mathsf{a}[\vec{i}]} = \langle \mathsf{exec}@\mathsf{a}[\vec{i}], \langle \mathsf{if } \mathsf{exec}@\mathsf{a}[\vec{i}] \mathsf{ then } \mathsf{output}@\mathsf{a}[\vec{i}]$$
$$\mathsf{else } \mathsf{empty}, \mathsf{frame}@\mathsf{pred}(\mathsf{a}[\vec{i}]) \rangle \rangle$$

Fig. 2. Interpretation of macros.

past conditions. The translation of the frame macro gathers (using nested pairs) all the information available to the attacker at some execution point: for each past action, the attacker observes if the execution continues and, if that is the case, obtains the output. Finally, to model the fact that the attacker controls the network, the input macro is interpreted using the attacker symbol att, to which we pass the current frame.

Example 8. *Considering the trace model of Example 7 on our running example, we have the following translation:*

$$(\mathsf{output}@T[i,j])_{\mathcal{P}_\mathsf{ex}}^\mathbb{T} = (o_{T[i,j]})_{\mathcal{P}_\mathsf{ex}}^\mathbb{T}$$
$$= (\mathsf{G}(\mathsf{H}(\mathsf{s}_\mathsf{T}[i]@\mathsf{pred}(T[i,j]), k), k'))_{\mathcal{P}_\mathsf{ex}}^\mathbb{T}$$
$$= \mathsf{G}(\mathsf{H}((\mathsf{s}_\mathsf{T}[i]@\mathsf{pred}(T[i,j]))_{\mathcal{P}_\mathsf{ex}}^\mathbb{T}, k), k').$$

We slightly abuse notation by using the same function symbols G, H, k and k' in the base logic and the meta-logic.

Then, the translation $(\mathsf{s}_\mathsf{T}[i]@\mathsf{pred}(T[i,j]))_{\mathcal{P}_\mathsf{ex}}^\mathbb{T}$ is unfolded to a stack of hashes until reaching the initial time point init. Here, after some basic reasoning in the base logic regarding the $\mathsf{if}\_\mathsf{then}\_\mathsf{else}$ construct, we obtain:

$$(\mathsf{s}_\mathsf{T}[i]@\mathsf{pred}(T[i,j]))_{\mathcal{P}_\mathsf{ex}}^\mathbb{T} = \mathsf{H}(s_2^0, k).$$

A local meta-logic formula ϕ is valid w.r.t. a protocol $\mathcal{P}$ when, for any trace model $\mathbb{T}$, the base logic formula $(\phi)_\mathcal{P}^\mathbb{T} \sim \mathsf{true}$ is valid (*i.e.* is satisfied in all computational models).

Example 9. *The following meta-formula models an authentication property which is valid in any computational model interpreting* G *and* H *as* EUF-CMA *secure hash functions.*

$$\forall j, i. \; \mathsf{happens}(R[j,i]) \Rightarrow \mathsf{cond}@R[j,i] \Rightarrow$$
$$(\exists j'. T[i,j'] < R[j,i] \wedge \mathsf{output}@T[i,j'] = \mathsf{input}@R[j,i])$$

This formula states that, whenever the condition of the reader's action $R[j,i]$ holds, there must be some session j' of the tag i (which uses the seed $s_0[i]$ to initialize its memory cell) that has been executed before $R[j,i]$. Moreover, the output of the tag's action coincides with the input of the reader's action. Here, we use quantifications to express our security goal. In the base logic, we would have to first fix a trace model, and then to explicitly enumerate all possible values for j, i, and j'.

Comparison with [18]: In the notion of trace model of [18], the set $\mathcal{D}_\mathcal{T}$ contains init and *all* concrete actions of the set $\mathcal{D}_a$ (see Definition 5). In our new definition, $\mathcal{D}_\mathcal{T}$ contains only a subset of all possible concrete actions, which correspond to the actions that happen. For example, this change allows to model conflicts between actions: for instance, the axiom $\forall \vec{i}, \vec{j}. \; \neg(\mathsf{happens}(a_1[\vec{i}]) \wedge \mathsf{happens}(a_2[\vec{j}]))$ rules out traces where the two actions are scheduled.

IV. APPLICATIONS AND LIMITATIONS

We present in Table I some case studies corresponding to protocols with global mutable states. These case studies use the extension of the meta-logic presented above to support mutable states, but rely on the same proof system as [18]. For each protocol, we give the security properties that we proved and the intermediate lemmas (also proved with SQUIRREL) used to carry out the proofs. We will comment below on these lemmas. All files are between 100 and 500 lines long (for both modeling and proof script). The length of each development mostly depends on the number of intermediate lemmas used.

Protocol	Intermediate Lemmas	Security Properties
Toy Counter [15]	Counter increase	Secrecy
CANAuth [15]	Counter increase	Authentication, Injectivity
Running example	Last update, Disjoint chains	Authentication
SLK06 [24]	-	Authentication
YPLRK05 [24]	No update	Authentication

TABLE I
CASE STUDIES OF SOME STATEFUL PROTOCOLS

1) Protocols with counters: We first consider two stateful protocols from [15], which rely on counter values that are incremented throughout the protocols execution. Toy Counter is a toy protocol where two agents A and B access a shared counter. We prove that this protocol provides some form of secrecy (expressed as a reachability property). CANAuth is a message authentication protocol where every agent, who can be both initiator and responder, stores a counter in memory which is incremented at each action of the agent. Proving the security of these protocols required reasoning on counter values. To that end, we axiomatized the counters ordering

relation in SQUIRREL. This can be done using, *e.g.*, the axioms (i.e. global meta-formulas) in Listing 1.

```
abstract Succ : message → message
abstract (~<) : message→ message→ boolean

axiom orderTrans (n1,n2,n3:message):
  n1 ~< n2 ⇒  n2 ~< n3 ⇒  n1 ~< n3.

axiom orderStrict (n1,n2:message):
  n1 = n2 ⇒  n1 ~< n2 ⇒  false.

axiom orderSucc (n1,n2:message):
  n1 = n2 ⇒  n1 ~< Succ(n2).
```

Listing 1. Axioms in SQUIRREL.

This highlights an advantage of our method, which lets the user axiomatize any theory needed to conduct the security analysis. In contrast, symbolic approaches are less flexible, and only let the user provide axioms in some restricted ways.

To analyze the security of these two protocols, we proved some intermediate lemmas about state values, *e.g.* that counters are monotonically increasing. Our security analysis is similar to the one in GSVerif [15], and we obtain the same results, except that we provide computational guarantees (assuming that the MAC function used in CANAuth is EUF-CMA secure).

2) Protocols with stacks of hashes: We consider next a few RFID protocols which use a hash function to repeatedly refresh a secret value: the protocol of Example 1 and simplified versions of the SLK06 and YPLRK05 protocols of [24]. We prove authentication properties in all cases. These protocols are intended to be used with unkeyed hash functions, in the random oracle model (ROM), but we model them here with keyed hash functions (satisfying the PRF assumption) — a stronger assumption. For the SLK06 and YPLRK05 protocols we even assume that each tag uses its own hash key (shared with the reader) which rules out confusions between the hashes of different tags and thus simplifies proofs. For our running example, the proof of authentication is more complex, involving in particular a lemma showing that the chains of hash stacks of distinct agents remain disjoint, despite their use of the same hash key. For these protocols, it has sometimes proved useful to do some basic reasoning about updates of the memory cells, such as proving that the value of a state at a given timestamp is either equal to the initial value, or equal to the update term of the last action that updated this particular state; or proving that the value of a state remains constant between two consecutive updates.

3) Limitations: In order to faithfully model the ROM in our setting, it would be natural to provide the attacker with a hashing oracle, *i.e.* include in the protocol, for each hash function $H(\_, k)$ in use, an action which upon input x outputs $H(x, k)$. With this modeling, we can still use the EUF-CMA and PRF tactics associated to hash functions, but the exploration of possible occurrences of hashes now includes this additional oracle process. In order to prove an authentication property, we use the unforgeability of some hash function

(rule EUF-CMA): if the reader receives a valid hash $H(m, k)$, it must have been produced by some action of the protocol. In presence of an oracle, we would have to consider the possibility that the hash has been obtained by feeding m to the oracle. This possibility can however be dismissed because m is not known to the attacker. Proving such *weak secrecy* properties directly is difficult in our setting, but we can prove the strong secrecy of m (i.e. that it is indistinguishable from a fresh name) and conclude from there that it is also weakly secret. Unfortunately, the proof system of [18] does not give the possibility to derive reachability properties (here, weak secrecy) from equivalence properties (strong secrecy). Our extension of the proof system presented in the next section allows to lift this limitation.

V. PROOF SYSTEM

We now generalize the proof system of [18] in order to enable complex reasoning where reachability and equivalence properties are simultaneously established, and to improve automated reasoning.

A. Global meta-logic formulas

The *local* meta-logic formulas of Section III express properties that hold for all traces of a protocol. We now introduce *global* meta-logic formulas, which will make it possible to express properties involving multiple protocols, and will be the basis for our extended sequent calculus. In order to be able to simultaneously talk of the traces of multiple protocols, we need these protocols to be compatible.

Definition 6. *Two protocols are* compatible *if they use the same set $\mathcal{A}$ of action names and have the same partial order.*

The syntax of global meta-formulas is shown next, where we write α for a variable of any sort, including message, and $\vec{t}, \vec{t}'$ are sequences of meta-terms and meta-formulas.

$$F \quad := \quad [\phi]_{\mathcal{P}} \mid [\vec{t} \sim \vec{t}']_{\mathcal{P},\mathcal{P}'} \mid \tilde{\perp} \mid \tilde{\top} \mid$$
$$\tilde{\forall}\alpha.F \mid \tilde{\exists}\alpha.F \mid F \tilde{\Rightarrow} F' \mid F \tilde{\wedge} F' \mid F \tilde{\vee} F'$$

To avoid the confusion with local meta-formulas we use other symbols for quantifiers and connectives. An atom of a global meta-formula may be a local formula relative to some protocol, or an equivalence relative to two protocols. We impose that all protocols used in a global meta-formula are compatible.

Similarly to local meta-formulas, a global meta-formula is said to be valid when, for any trace model, its meta-interpretation as a base logic *formula* is valid. This meta-interpretation is defined as follows — notice the specific treatment of quantification over messages.

$$(\tilde{\perp})^{\mathbb{T}} \stackrel{\text{def}}{=} \perp \qquad\qquad (\tilde{\forall}i.F)^{\mathbb{T}} \stackrel{\text{def}}{=} \wedge_{v \in \mathcal{D}_{\mathcal{I}}} (F)^{\mathbb{T}\{i \mapsto v\}}$$

$$([\phi]_{\mathcal{P}})^{\mathbb{T}} \stackrel{\text{def}}{=} (\phi)_{\mathcal{P}}^{\mathbb{T}} \sim \text{true} \qquad (\tilde{\forall}\tau.F)^{\mathbb{T}} \stackrel{\text{def}}{=} \wedge_{v \in \mathcal{D}_{\mathcal{T}}} (F)^{\mathbb{T}\{\tau \mapsto v\}}$$

$$(F \tilde{\Rightarrow} F')^{\mathbb{T}} \stackrel{\text{def}}{=} (F)^{\mathbb{T}} \Rightarrow (F')^{\mathbb{T}} \quad (\tilde{\forall}x.F)^{\mathbb{T}} \stackrel{\text{def}}{=} \forall x.(F)^{\mathbb{T}}$$

$$([\vec{t} \sim \vec{t}']_{\mathcal{P},\mathcal{P}'})^{\mathbb{T}} \stackrel{\text{def}}{=} (\vec{t})_{\mathcal{P}}^{\mathbb{T}} \sim (\vec{t}')_{\mathcal{P}'}^{\mathbb{T}}$$

The other cases are similarly handled.

Definition 7. *Two protocols $\mathcal{P}_1$ and $\mathcal{P}_2$ are* observationally equivalent *when they are compatible and the following global meta-logic formula is valid:*

$$\tilde{\forall}\tau. \, [\text{happens}(\tau)]_{\mathcal{P}_1} \tilde{\Rightarrow} [\text{frame@}\tau \sim \text{frame@}\tau]_{\mathcal{P}_1,\mathcal{P}_2}.$$

Note that $[\text{happens}(\tau)]_{\mathcal{P}_1}$ and $[\text{happens}(\tau)]_{\mathcal{P}_2}$ are equivalent as $\text{happens}(\tau)$ only talks about the trace model.

B. Sequents

We consider two kinds of sequents, general and reachability sequents, which are respectively of the form

$$\Sigma; \Theta \vdash F \qquad \text{and} \qquad \Sigma; \Theta : \Gamma \vdash_{\mathcal{P}} \phi$$

where Σ is a sequence of variables (of any sort), Θ is a set of global meta-formulas, Γ is a set of local meta-formulas, ϕ is a local meta-formula and F is a global meta-formula. We require that sequents are closed, *i.e.* Σ binds all variables occurring in the rest of the sequent.

These sequents translate to global meta-formulas, which indirectly defines their semantics. When $\Theta = \{F_1, \dots, F_n\}$ we write $\Theta \tilde{\Rightarrow} F$ for $F_1 \tilde{\Rightarrow} \dots \tilde{\Rightarrow} F_n \tilde{\Rightarrow} F$. We define similarly the notation $\Gamma \tilde{\Rightarrow} \phi$. Finally, when $\Sigma = \{\alpha_1, \dots, \alpha_n\}$, we write $\tilde{\forall}\Sigma.F$ for $\tilde{\forall}\alpha_1 \dots \tilde{\forall}\alpha_n.F$. The translation of our two kinds of sequents is then defined as follows.

$$\Sigma; \Theta \vdash F \quad \rightsquigarrow \quad \tilde{\forall}\Sigma.(\Theta \tilde{\Rightarrow} F)$$
$$\Sigma; \Theta : \Gamma \vdash_{\mathcal{P}} \phi \quad \rightsquigarrow \quad \tilde{\forall}\Sigma.(\Theta \tilde{\Rightarrow} [\Gamma \Rightarrow \phi]_{\mathcal{P}})$$

A sequent is valid when its translation as a global meta-formula is valid.

Example 10. *To understand the distinction between Θ and Γ hypotheses in reachability sequents, note that the validity of $\Sigma; : \phi \vdash_{\mathcal{P}} \psi$ implies that of $\Sigma; [\phi]_{\mathcal{P}} : \vdash_{\mathcal{P}} \psi$, but the converse does not generally hold. In the former case we state that the implication $\phi \Rightarrow \psi$ is true with overwhelming probability; in the latter, we state that if ϕ is true with overwhelming probability then so is ψ.*

When $\vec{u}$ and $\vec{v}$ are sequences of meta-terms or meta-formulas of the same length, we write $\Sigma; \Theta \vdash_{\mathcal{P},\mathcal{P}'} \vec{u} \sim \vec{v}$ for $\Sigma; \Theta \vdash [\vec{u} \sim \vec{v}]_{\mathcal{P},\mathcal{P}'}$.

C. Structural rules

Reachability sequents (resp. general sequents) generalize the reachability (resp. equivalence) sequents of [18], and the proof rules for that earlier work can naturally be adapted for our generalized sequents. This initial set of proof rules includes, for both kinds of sequents, the usual first-order reasoning rules and proof by induction for both kinds of sequents (see Appendix A) but also rules that embody cryptographic assumptions or reasoning about names. We shall not comment further these (essentially) unchanged rules, but refer the reader to [18] for more details.

The semantics of local meta-formulas is not boolean, but probabilistic: a formula is valid if its meta-interpretation is

$$\frac{\Sigma; \Theta : \Gamma, \phi \vdash_{\mathcal{P}} \psi \qquad \Sigma; \Theta : \Gamma, \phi' \vdash_{\mathcal{P}} \psi}{\Sigma; \Theta : \Gamma, \phi \vee \phi' \vdash_{\mathcal{P}} \psi} \qquad \frac{\Sigma; \Theta, [\phi]_{\mathcal{P}} : \Gamma \vdash_{\mathcal{P}} \psi \qquad \Sigma; \Theta, [\phi']_{\mathcal{P}} : \Gamma \vdash_{\mathcal{P}} \psi}{\Sigma; \Theta, [\phi \vee \phi']_{\mathcal{P}} : \Gamma \vdash_{\mathcal{P}} \psi}$$

In the right rule, we require that one of the two disjuncts is a pure trace formula.

Fig. 3. Left disjunction rules at the local and global levels.

GLOBAL-LOCAL
$$\frac{\Sigma; \Theta \vdash [\phi]_{\mathcal{P}}}{\Sigma; \Theta : \vdash_{\mathcal{P}} \phi}$$

LOCAL-GLOBAL
$$\frac{\Sigma; \Theta : \vdash_{\mathcal{P}} \phi}{\Sigma; \Theta \vdash [\phi]_{\mathcal{P}}}$$

EQUIV-TERM
$$\frac{\Sigma; \Theta : \vdash_{\mathcal{P}} t = t' \qquad \Sigma; \Theta \vdash_{\mathcal{P}, \mathcal{P}'} \vec{C}[t'] \sim \vec{v}}{\Sigma; \Theta \vdash_{\mathcal{P}, \mathcal{P}'} \vec{C}[t] \sim \vec{v}}$$

EQUIV-FORM
$$\frac{\Sigma; \Theta : \vdash_{\mathcal{P}} \phi \Leftrightarrow \psi \qquad \Sigma; \Theta \vdash_{\mathcal{P}, \mathcal{P}'} \vec{C}[\psi] \sim \vec{v}}{\Sigma; \Theta \vdash_{\mathcal{P}, \mathcal{P}'} \vec{C}[\phi] \sim \vec{v}}$$

REWRITE-EQUIV
$$\frac{\Sigma; \Theta \vdash_{\mathcal{P}, \mathcal{P}'} (\Gamma \Rightarrow \phi) \sim (\Delta \Rightarrow \psi) \qquad \Sigma; \Theta : \Delta \vdash_{\mathcal{P}'} \psi}{\Sigma; \Theta : \Gamma \vdash_{\mathcal{P}} \phi}$$

Fig. 4. Inference rules involving mixed kinds of sequents.

true with overwhelming probability. It is thus remarkable that the usual rules of first-order logic are sound for these formulas. But this only holds when we remain within the "local" part of local sequents, as illustrated in Figure 3. The second rule in that figure would not be valid without the proviso that one of the disjuncts is a pure trace formula. A pure trace formula is a local meta-formula which does not feature any macros, and is always true or false when evaluated in a trace model.

Definition 8. *A local meta-formula ϕ is a pure trace formula w.r.t. protocol $\mathcal{P}$ if ϕ does not contain any macro and if for any trace model $\mathbb{T}$, $(\phi)^{\mathbb{T}}_{\mathcal{P}} \sim$ true or $(\phi)^{\mathbb{T}}_{\mathcal{P}} \sim$ false is valid.*

A simple syntactic criterion guaranteeing this property is that the formula does not feature any macro or message, *i.e.* its atoms are only comparisons of indices and timestamps.

Essentially, assuming that $\phi \vee \phi'$ is true with overwhelming probability, we cannot generally conclude that one of ϕ and ϕ' is also true with overwhelming probability; the pure trace formula proviso allows one to fix this gap because a pure trace formula is either always true or always false (depending on the trace model). Proofs in our case studies must sometimes be carefully structured to be able to accommodate this subtlety.

Our proof system also features rules which articulate the relationship between our two kinds of sequents, shown in Figure 4. Rules GLOBAL-LOCAL and LOCAL-GLOBAL are obvious conversions. Rules EQUIV-TERM and EQUIV-FORM are immediate generalizations of the identically named rules in [18], with the new environment Θ allowing to carry hypotheses

from the equivalence sequent in conclusion to the reachability sequent in first premise. Conversely, the new rule REWRITE-EQUIV shows how an equivalence judgement can be used to help derive reachability judgements. As expected all these rules have been shown to be sound.

Proposition 1. *The rules in Figures 3 and 4 are sound.*

Example 11. *To illustrate the utility of this articulation between equivalence and reachability, we show how to prove a (weak) secrecy property (expressed as a reachability sequent) using a strong secrecy property (an equivalence) as hypothesis. Consider a protocol $\mathcal{P}$ involving a memory cell s of index arity 1. Given a fresh name m, we consider the reachability sequent $\Sigma; \Theta_{\mathsf{hap}}, \Theta : \vdash_{\mathcal{P}} \phi$ where $\Sigma \stackrel{def}{=} \tau, \tau', i$ and:*

$$\Theta_{\mathsf{hap}} \stackrel{def}{=} [\mathsf{happens}(\mathsf{pred}(\tau))]_{\mathcal{P}}$$
$$\Theta \stackrel{def}{=} [\mathsf{frame}@\tau, s[i]@\tau' \sim \mathsf{frame}@\tau, m]_{\mathcal{P}, \mathcal{P}}$$
$$\phi \stackrel{def}{=} \mathsf{input}@\tau \neq s[i]@\tau'$$

This sequent states that, if the values stored in the memory cell s are strongly secret (this is the global meta-formula Θ), then we know that $s[i]@\tau'$ is (weakly) secret, i.e. the attacker cannot deduce its value (this is the local meta-formula ϕ). In order to derive this sequent, we use the rule REWRITE-EQUIV with $\psi := \mathsf{input}@\tau \neq m$ and empty environments Γ and Δ. We thus have to derive the following two sequents.

$$\Sigma; \Theta_{\mathsf{hap}}, \Theta \vdash_{\mathcal{P}, \mathcal{P}} \mathsf{input}@\tau \neq s[i]@\tau' \sim \mathsf{input}@\tau \neq m$$

$$\Sigma; \Theta_{\mathsf{hap}}, \Theta : \vdash_{\mathcal{P}} \mathsf{input}@\tau \neq m$$

The second one can be easily derived using the freshness of the name m. Concerning the first one, we shall see that we can use the equivalence in Θ to deduce the right side of the sequent: this is the purpose of the next subsection. The file corresponding to this example in [16] is `running-ex-secrecy.sp`.

Echoing the limitations discussed at the end of Section IV, the articulation illustrated above is also critical to prove the strong secrecy of the states of tags in a variant of the OSK protocol with oracles but no readers. Interestingly, strong and weak secrecy must be derived simultaneously in that proof, in a mutual induction. See the file `running-ex-oracle.sp` at [16] for more details.

D. Automating proofs by bi-deduction

In the sequent $\Sigma; \Theta_{\mathsf{hap}}, \Theta \vdash_{\mathcal{P}, \mathcal{P}} \phi \sim \psi$ of Example 11, the left side of the equivalence in Θ contains all the necessary information to compute ϕ: there exists a machine $\mathcal{B}$ computing the latter from the former. Moreover, the *same* algorithm $\mathcal{B}$ computes ψ from the right side of the equivalence in Θ. Intuitively, the proof is done here: if you can break the equivalence $\phi \sim \psi$, then you can break the equivalence Θ by composing the distinguisher for $\phi \sim \psi$ with $\mathcal{B}$. The existence of such a $\mathcal{B}$ proves that an equivalence is entailed by another equivalence. Informally:

$$\frac{\exists \mathcal{B} \text{ s.t. } \mathcal{B} \text{ computes } \vec{v}_i \text{ from } \vec{u}_i \text{ w.r.t. } \mathcal{P}_i \text{ for any } i \in \{1, 2\}}{\Sigma; \Theta, [\vec{u}_1 \sim \vec{u}_2]_{\mathcal{P}_1, \mathcal{P}_2} \vdash [\vec{v}_1 \sim \vec{v}_2]_{\mathcal{P}_1, \mathcal{P}_2}} \quad (1)$$

9

We shall now define formally this notion of simultaneous computation, which we call *bi-deduction* and note $\Sigma; \Theta; \#(\vec{u}_1, \vec{u}_2) \triangleright_{\mathcal{P}_1, \mathcal{P}_2} \#(\vec{v}_1, \vec{v}_2)$.

1) Bi-terms: A bi-term $t_{\#}$ is an element of the form $\#(t_1, t_2)$ where t_1 and t_2 are meta-terms. Similarly, a bi-formula $\#(\phi_1, \phi_2)$ is a pair of local formulas. We often use compact notations to refer to bi-terms (and bi-formulas) sharing some top-level constructs.

For example, $f(\#(t_1, t_2), g(\#(s_1, s_2)))$ is the bi-term $\#(f(t_1, g(s_1)), f(t_2, g(s_2)))$, and $\mathsf{seq}(\vec{\alpha}: \#(t_1, t_2))$ is the bi-term $\#(\mathsf{seq}(\vec{\alpha}: t_1), \mathsf{seq}(\vec{\alpha}: t_2))$. We may also note t for the bi-term $\#(t, t)$. We similarly lift boolean connectives to bi-terms: *e.g.* $\#(\phi_1, \phi_2) \wedge \#(\psi_1, \psi_2)$ is $\#(\phi_1 \wedge \psi_1, \phi_2 \wedge \psi_2)$.

2) Bi-deduction: A bi-deduction judgement is an element of the form $\Sigma; \Theta; \vec{u}_{\#} \triangleright_{\mathcal{P}_1, \mathcal{P}_2} \{v_{\#} \mid \phi_{\#}\}$ where Σ is a sequence of variables of any sort, Θ is a set of global meta-formulas, $\mathcal{P}_1, \mathcal{P}_2$ are protocols, $\vec{u}_{\#}$ is a sequence of bi-terms, v is a bi-term, and $\phi_{\#}$ is a bi-formula. We require that Σ binds all variables in the judgement. We omit the protocols $\mathcal{P}_1, \mathcal{P}_2$ from the judgement when they are clear from the context, and write $\Sigma; \Theta; \vec{u}_{\#} \triangleright v_{\#}$ when $\phi_{\#} = \top$. We also write $\Sigma; \Theta; \vec{u}_{\#} \triangleright \vec{v}_{\#}$ when $\Sigma; \Theta; \vec{u}_{\#} \triangleright v_{\#}^i$ for each element $v_{\#}^i$ in $\vec{v}_{\#}$.

A judgement $\Sigma; \Theta; \#(\vec{u_1}, \vec{u_2}) \triangleright_{\mathcal{P}_1, \mathcal{P}_2} \{\#(v_1, v_2) \mid \#(\phi_1, \phi_2)\}$ is valid when, for any trace model $\mathbb{T}$, there exist probabilistic PTIME machines $\mathcal{B}_o$ and $\mathcal{B}_c$ such that, for any computational model $\mathbb{M}$ and valuation σ such that $\mathbb{M}, \sigma \models (\Theta)^{\mathbb{T}}$, and for any $i \in \{1, 2\}$, the following equalities hold with overwhelming probability:

$$\begin{aligned}
& [\![\mathsf{EQ}]\!]_{\mathbb{M}}\big(\mathcal{B}_c([\![(\vec{u}_i)_{\mathcal{P}_i}^{\mathbb{T}}]\!]_{\mathbb{M}}^{\sigma}), \ [\![(\phi_i)_{\mathcal{P}_i}^{\mathbb{T}}]\!]_{\mathbb{M}}^{\sigma}) = 1 \\
& [\![\mathsf{EQ}]\!]_{\mathbb{M}}\big(\mathcal{B}_o([\![(\vec{u}_i)_{\mathcal{P}_i}^{\mathbb{T}}]\!]_{\mathbb{M}}^{\sigma}), \ [\![(\text{if } \phi_i \text{ then } v_i)_{\mathcal{P}_i}^{\mathbb{T}}]\!]_{\mathbb{M}}^{\sigma}) = 1
\end{aligned} \quad (2)$$

Said otherwise, given $\vec{u}_i$ as input, the machines $\mathcal{B}_c$ and $\mathcal{B}_o$ respectively compute ϕ_i and (if ϕ_i then v_i) (up to a negligible probability of failure) whenever Θ holds, for $i = 1$ and protocol $\mathcal{P}_1$, and for $i = 2$ and protocol $\mathcal{P}_2$. Note that the *same* machines are used on *both* sides.

Example 12. *Let i', i be index variables. The bi-deduction judgement $i', i; \ ; \ \triangleright i' = i$ is valid because, for any trace model $\mathbb{T}$, we can choose $\mathcal{B}$ to be the machine that always returns true if $(i' = i)^{\mathbb{T}}$ is true, and false otherwise. The dependence of $\mathcal{B}$ on $\mathbb{T}$ is crucial here.*

Example 13. *Let $a[\_]$ be some action name of index arity one, τ a timestamp variable and i, j index variables. The bi-deduction judgement:*

$$\tau, j; \ ; \ \mathsf{seq}(i: n[i] \mid a[i] \leq \tau) \triangleright \{n[j] \mid a[j] < \tau\}$$

is valid because the sequence $\mathsf{seq}(i: n[i] \mid a[i] \leq \tau)$ contains the names $n[i]$ for any $a[i] \leq \tau$, and consequently contains $n[j]$ whenever $a[j] < \tau$. Hence we can easily construct a machine $\mathcal{B}$ extracting $n[j]$ from the sequence.

3) Rules: We can now give a formal version of Equ. (1), *i.e.* state the rule that derives equivalences from bi-deductions:

$$\frac{\Sigma; \Theta; \#(\vec{u}_1, \vec{u}_2) \triangleright_{\mathcal{P}_1, \mathcal{P}_2} \#(\vec{v}_1, \vec{v}_2)}{\Sigma; \Theta, [\vec{u}_1 \sim \vec{u}_2]_{\mathcal{P}_1, \mathcal{P}_2} \vdash [\vec{v}_1 \sim \vec{v}_2]_{\mathcal{P}_1, \mathcal{P}_2}} \ \text{\small BI-DEDUCE}$$

If, given some known bi-terms $\#(\vec{u}_1, \vec{u}_2)$, we can bi-deduce some other bi-terms $\#(\vec{v}_1, \vec{v}_2)$ (for a pair of protocols $\mathcal{P}_1, \mathcal{P}_2$), then we know that the equivalence $[\vec{u}_1 \sim \vec{u}_2]_{\mathcal{P}_1, \mathcal{P}_2}$ entails $[\vec{v}_1 \sim \vec{v}_2]_{\mathcal{P}_1, \mathcal{P}_2}$.

Example 14. *Let us come back to Example 11 in which we had to derive the following sequent:*

$$\Sigma; \Theta_{\mathsf{hap}}, [\mathsf{frame}@\tau, \mathsf{s}[i]@\tau' \sim \mathsf{frame}@\tau, m]_{\mathcal{P}, \mathcal{P}}$$
$$\vdash_{\mathcal{P}, \mathcal{P}} \mathsf{input}@\tau \neq \mathsf{s}[i]@\tau' \sim \mathsf{input}@\tau \neq m$$

Using the rule BI-DEDUCE*, we are left with the following bi-deduction judgement:*

$$\Sigma; \Theta_{\mathsf{hap}}; \mathsf{frame}@\tau, \#(\mathsf{s}[i]@\tau', m) \triangleright_{\mathcal{P}, \mathcal{P}} \mathsf{input}@\tau \neq \#(\mathsf{s}[i]@\tau', m)$$

That judgement is valid according to our semantics: indeed there is a machine that can (1) obtain $x := \mathsf{frame}@\tau$ and $y := \mathsf{s}[i]@\tau'$ when $i = 1$ (resp. $y := m$ when $i = 2$); (2) build $z := \mathsf{input}@\tau$ from $\mathsf{frame}@\tau$; (3) compute $z \neq y$.

We present in Figure 5 some rules for bi-deducibility; a complete presentation is given in Figure 13, Appendix B. Rule SEQ-DED allows to bi-deduce a bi-term $\#(v_1, v_2)$, whenever $\#(\phi_1, \phi_2)$ holds, from a sequence $\mathsf{seq}(\vec{\alpha}: \#(u_1, u_2) \mid \#(\psi_1, \psi_2))$ if it appears in the sequence *simultaneously* on both sides: there must exists $\vec{\alpha}$ such that, for any $i \in \{1, 2\}$, $u_i = v_i$ and the sequence condition ψ_i holds whenever ϕ_i holds. Rule PURE-DED states that pure trace formulas are always bi-deducible: indeed, once the trace model is fixed, they have a constant value (either always true or always false). Rule PAIR-PROJ allows to exploit the fact that a machine can recover a pair's components. The FRAME rule states that $\mathsf{frame}@T$ contains all $\mathsf{frame}@T'$ for any $T' \leq T$. Combined with other rules, this makes it possible to bi-deduce any past input from a frame. The rule FA says that the image of a function $f(\vec{v}_{\#})$ is bi-deducible whenever its arguments $\vec{v}_{\#}$ are bi-deducible. In case f is of arity zero, we must also check that the condition $\phi_{\#}$ is bi-deducible — see rule FA$_0$. Finally, FA-ITE allows one to derive a bi-deduction with a conditional on its right, taking advantage of its condition to refine the premises: we know that we must deduce the then branch only if the condition ψ holds, and similarly for the else branch.

Proposition 2. *The* BI-DEDUCE *rule and the rules in Figure 5 (as well as the complete version of the rules given in Figure 13 in Appendix) are sound.*

Example 15. *We now formally derive the bi-deduction judgement of Example 14. We first apply* FA *on the $\neq$ function symbol. This gives us two bi-deduction judgements:*

1) $\Sigma; \Theta_{\mathsf{hap}}; \mathsf{frame}@\tau, \#(\mathsf{s}[i]@\tau', m) \triangleright_{\mathcal{P}, \mathcal{P}} \mathsf{input}@\tau$; and
2) $\Sigma; \Theta_{\mathsf{hap}}; \mathsf{frame}@\tau, \#(\mathsf{s}[i]@\tau', m) \triangleright_{\mathcal{P}, \mathcal{P}} \#(\mathsf{s}[i]@\tau', m)$.

$$\text{SEQ-DED}$$
$$\frac{\Sigma;\ \Theta \vdash \tilde{\exists}\vec{\alpha}.\ \begin{pmatrix} [\phi_1 \Rightarrow (\psi_1 \wedge u_1 = v_1)]_{\mathcal{P}_1} \tilde{\wedge} \\ [\phi_2 \Rightarrow (\psi_2 \wedge u_2 = v_2)]_{\mathcal{P}_2} \end{pmatrix} \qquad \Sigma;\ \Theta;\ \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: \#(u_1, u_2) \mid \#(\psi_1, \psi_2)) \rhd \#(\phi_1, \phi_2)}{\Sigma;\ \Theta;\ \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: \#(u_1, u_2) \mid \#(\psi_1, \psi_2)) \rhd \{\#(v_1, v_2) \mid \#(\phi_1, \phi_2)\}}$$

$$\text{PURE-DED}$$
$$\frac{v, \phi \text{ pure trace formula}}{\Sigma;\ \Theta;\ \vec{u}_\# \rhd \{\#(v, v) \mid \#(\phi, \phi)\}}$$

$$\text{PAIR-PROJ}$$
$$\frac{\Sigma;\ \Theta;\ \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: u_\#^1 \mid \psi_\#), \mathsf{seq}(\vec{\alpha}: u_\#^2 \mid \psi_\#) \rhd \{v_\# \mid \phi_\#\}}{\Sigma;\ \Theta;\ \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: \langle u_\#^1, u_\#^2\rangle \mid \psi_\#) \rhd \{v_\# \mid \phi_\#\}}$$

$$\text{FRAME}$$
$$\frac{\Sigma;\ \Theta;\ \vec{u}_\#, \mathsf{seq}(\vec{\alpha}, T': \mathsf{frame}@T' \mid \psi_\# \wedge T' \leq T) \rhd \{v_\# \mid \phi_\#\}}{\Sigma;\ \Theta;\ \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: \mathsf{frame}@T \mid \psi_\#) \rhd \{v_\# \mid \phi_\#\}}$$

$$\text{FA-ITE}$$
$$\text{FA} \qquad\qquad \text{FA}_0 \qquad\qquad \frac{\begin{array}{c}\Sigma;\ \Theta;\ \vec{u}_\# \rhd \{\psi_\# \mid \phi_\#\} \\ \Sigma;\ \Theta;\ \vec{u}_\# \rhd \{t_\# \mid \phi_\# \wedge \psi_\#\} \\ \Sigma;\ \Theta;\ \vec{u}_\# \rhd \{t'_\# \mid \phi_\# \wedge \neg\psi_\#\}\end{array}}{\Sigma;\ \Theta;\ \vec{u}_\# \rhd \{\mathsf{if}\ \psi_\#\ \mathsf{then}\ t_\#\ \mathsf{else}\ t'_\# \mid \phi_\#\}}$$

$$\frac{\Sigma;\ \Theta;\ \vec{u}_\# \rhd \{\vec{v}_\# \mid \phi_\#\}}{\Sigma;\ \Theta;\ \vec{u}_\# \rhd \{f(\vec{v}_\#) \mid \phi_\#\}} \qquad\qquad \frac{\Sigma;\ \Theta;\ \vec{u}_\# \rhd \phi_\#}{\Sigma;\ \Theta;\ \vec{u}_\# \rhd \{f() \mid \phi_\#\}}$$

We identify a single term u with the degenerate sequence $\mathsf{seq}(\epsilon: u \mid \top)$ so that the left rules apply not only to sequences. The FA rule only applies when f is of non-zero arity.

Fig. 5. Bi-deduction rules for protocols $\mathcal{P}_1, \mathcal{P}_2$.

The second judgement is derived using SEQ-DED*, since the bi-term we must deduce is already present in the left side of the judgement. For the first one, we replace* input@τ *by* att(frame@pred(τ))*: this is done by unfolding the* input *macro (rule not shown). Then, we get rid of the* att *function symbol using* FA*. Applying* FRAME *then gives:*

$$\Sigma;\ \Theta_{\mathsf{hap}};\ \mathsf{seq}(T: \mathsf{frame}@T \mid T \leq \tau) \rhd_{\mathcal{P},\mathcal{P}} \mathsf{frame}@\mathsf{pred}(\tau).$$

We conclude with the rule SEQ-DED*, noticing that we can instantiate* T *by* pred(τ)*.*

Using our bi-deduction rules, we implemented a procedure $\mathsf{DED}(\cdot)$ which automatically checks whether a bi-deduction judgement $\Sigma;\ \Theta;\ \vec{u}_\# \rhd \{v_\# \mid \phi_\#\}$ is valid. Our procedure is sound but incomplete: it relies on heuristics to decide on which order it applies the rules. Studying the decidability and complexity of the bi-deduction logic is left as future work.

This procedure has been incorporated in SQUIRREL `apply` tactic: when applying a lemma H whose conclusion is an equivalence, we use our automatic procedure to check that the current goal is bi-deducible from H's conclusion. This improvement plays a crucial role to ease the proof effort for the case studies presented in this paper, but has also allowed to simplify older case studies from [18].

a) Discussion: Some bi-deduction proofs can be done directly using equivalence judgements and existing rules. Roughly, the idea is to construct an explicit context computing the wanted terms $\vec{v}$ from the known terms $\vec{u}$, using DUP, FA.

But this is not always possible, in particular when sequences are involved. E.g., using BI-DEDUCE, it is possible to prove (the two sides are syntactically equal but interpreted w.r.t. distinct protocols):

$$\tau;\ \begin{bmatrix} \mathsf{frame}@\tau \sim \\ \mathsf{frame}@\tau \end{bmatrix}_{\mathcal{P}_1,\mathcal{P}_2} \vdash \begin{bmatrix} \mathsf{seq}(\tau_1: \mathsf{input}@\tau_1 \mid \tau_1 < \tau) \sim \\ \mathsf{seq}(\tau_1: \mathsf{input}@\tau_1 \mid \tau_1 < \tau) \end{bmatrix}_{\mathcal{P}_1,\mathcal{P}_2}$$

using the fact that frame@τ contains frame@τ_0 for all $\tau_0 \leq \tau$, and the definition of input — see Appendix B-B for a detailed proof. With our current set of rules, we believe this judgement cannot be proved without BI-DEDUCE, as we cannot construct an explicit context extracting frame@τ_0 from frame@τ for $\tau_0 \leq \tau$. Indeed, such a context would have to perform projections on frames a number of times which depends on the trace models, which cannot be done.

E. Induction and bi-deduction

The automatic bi-deduction procedure derived from the bi-deduction proof system presented so far is very efficient and helpful. In some cases, though, it is not enough to conclude, at least not with a reasonable proof effort.

Example 16. *Consider the protocol of our running example, with tags only for simplicity. Call it $\mathcal{P}_1$, and let $\mathcal{P}_2$ be a modified version where the states $\mathsf{s}_\mathsf{T}[i]$ are initialized with $n^0[i]$ instead of $s^0[i]$. Consider the equivalence*

$$\begin{array}{c} \mathsf{frame}@t, k, \mathsf{seq}(i: s^0[i]) \sim_{\mathcal{P}_1,\mathcal{P}_2} \\ \mathsf{frame}@t, k, \mathsf{seq}(i: n^0[i]) \end{array} \qquad (3)$$

where we reveal, in addition to the frames, the key k and the initial values for all i. This equivalence obviously holds because it compares terms that only differ by a bijective renaming. Consider now an even stronger equivalence where we reveal all past states:

$$\begin{array}{c} \mathsf{frame}@\tau, k, \mathsf{seq}(i, \tau': \mathsf{if}\ \tau' \leq \tau\ \mathsf{then}\ \mathsf{s}_\mathsf{T}[i]@\tau') \sim_{\mathcal{P}_1,\mathcal{P}_2} \\ \mathsf{frame}@\tau, k, \mathsf{seq}(i, \tau': \mathsf{if}\ \tau' \leq \tau\ \mathsf{then}\ \mathsf{s}_\mathsf{T}[i]@\tau') \end{array} \qquad (4)$$

Such enriched equivalences are typically needed for use with REWRITE-EQUIV*, as is the case in our YubiHSM case study. In our simple example, the enriched equivalence (4) obviously holds because it is a renaming, but it is difficult to derive it without an ad-hoc renaming rule, which would be inapplicable in realistic cases such as our YubiHSM case study. Intuitively,*

equivalence (4) follows from (3) because the attacker can bi-deduce the successive states from their initial values and k. However, the number of steps of the construction of $\mathsf{s_T}[i]@\tau'$ depends on the position of τ' in a trace model, so we need some sort of inductive reasoning to establish this bi-deduction.

At a high-level, this inductive reasoning would proceed as follows. If τ is init, the sequence of states is composed of initial values only, which are given in the first equivalence. If τ is undef, the comparison $\tau' \leq \tau$ cannot hold hence the sequence is essentially empty. Otherwise, τ is of the form $T[i', j]$. By induction hypothesis we have already bi-deduced the part of the sequence corresponding to values of τ' such that $\tau' < \tau$. It remains to bi-deduce the states for $\tau' = \tau = T[i', j]$. But, for any i, we have (for both $\mathcal{P}_1$ and $\mathcal{P}_2$)

$$\mathsf{s_T}[i]@T[i', j] \stackrel{def}{=} \text{if } i = i' \text{ then } \mathsf{H}(\underline{\mathsf{s_T}[i']@\mathsf{pred}(T[i', j])}, \underline{k})$$
$$\text{else } \underline{\mathsf{s_T}[i]@\mathsf{pred}(T[i', j])}$$

which can be bi-deduced, notably because the parts underlined above have already been shown to be bi-deducible.

The reasoning performed in the previous example cannot be formally carried out with the bi-deduction inference rules presented so far, because they do not include any form of inductive reasoning. Of course, an induction and case analysis over τ could be done at the general sequent level, but this would require the user to supply the inductive invariant and carry out a very tedious manual proof, as shown in `running-ex-deduction.sp` on the previous example.

Instead, we extend our automated procedure for bi-deduction with inductive reasoning. The soundness of our procedure relies on the following induction rule:

IND-▷
$$\frac{\Sigma, \tau; \Theta; \vec{u}_\#, \mathsf{seq}(\tau' : \vec{v}_\#[\tau'] \mid \phi_\# \land \tau' < \tau) \triangleright \{\vec{v}_\#[\tau] \mid \phi_\#\}}{\Sigma; \Theta; \vec{u}_\# \triangleright \{\mathsf{seq}(\tau : \vec{v}_\#[\tau]) \mid \phi_\#\}}$$

Assume we want to prove the bi-deduction judgement $\Sigma; \Theta; u_\# \triangleright v_\#$. First, our procedure computes an invariant of the form $U_\# \stackrel{def}{=} \lambda\tau, u_\#[\tau]$ such that $U_\#$ is an invariant of the protocol:

$$\Sigma, \tau; \Theta; \vec{u}_\#, \mathsf{seq}(\tau' : U_\#[\tau'] \mid \tau' < \tau) \triangleright U_\#[\tau]$$

This invariant $U_\#$ is computed automatically, using abstract interpretation techniques [25]. Finally, we check whether

$$\Sigma; \Theta; u_\#, \mathsf{seq}(\tau : U_\#[\tau]) \triangleright v_\#$$

using our automated procedure $\mathsf{DED}(\cdot)$. We refer the reader to Appendix C for the proof of soundness of this approach and additional details on how the invariant $U_\#$ is computed.

This procedure for bi-deduction extended with inductive reasoning has been incorporated in the SQUIRREL `apply` tactic. Because it is more costly, it is not used by default, and must be turned on by using `apply ~inductive`.

Example 17. *The equivalence in Equ.* (4) *of Example 16 is proved directly using our bi-deduction procedure with inductive reasoning (see* `running-ex-deduction.sp` *in [16]).*

SPLIT-SEQ
$$\frac{\Sigma; \Theta \vdash_{\mathcal{P}, \mathcal{P}'} \sim \begin{array}{l} \vec{u}, \mathsf{seq}(\vec{\alpha} : \text{if } \phi \text{ then } t), \mathsf{seq}(\vec{\alpha} : \text{if } \neg\phi \text{ then } t) \\ \vec{v}, \mathsf{seq}(\vec{\alpha} : \text{if } \psi \text{ then } s), \mathsf{seq}(\vec{\alpha} : \text{if } \neg\psi \text{ then } s) \end{array}}{\Sigma; \Theta \vdash_{\mathcal{P}, \mathcal{P}'} \vec{u}, \mathsf{seq}(\vec{\alpha} : t) \sim \vec{v}, \mathsf{seq}(\vec{\alpha} : s)}$$

CONST-SEQ
$$\frac{\begin{array}{l} \Sigma; \Theta \vdash \left[\forall\vec{\alpha}, \bigvee_{0 \leq i \leq n} b_i \right]_{\mathcal{P}} \\ \Sigma; \Theta \vdash \left[\bigwedge_{0 \leq i \leq n} \forall\vec{\alpha}, b_i \Rightarrow t = t_i \right]_{\mathcal{P}} \\ \Sigma; \Theta \vdash \left[\bigwedge_{0 \leq i \leq n} \forall\vec{\alpha}, b_i \Rightarrow s = s_i \right]_{\mathcal{P}'} \\ \vec{\alpha} \cap \Sigma = \emptyset \quad (b_i)_{1 \leq i \leq n} \text{ pure trace formulas} \\ \Sigma; \Theta \vdash_{\mathcal{P}, \mathcal{P}'} \vec{u}, t_1, \ldots, t_n \sim \vec{v}, s_1, \ldots, s_n \end{array}}{\Sigma; \Theta \vdash_{\mathcal{P}, \mathcal{P}'} \vec{u}, \mathsf{seq}(\vec{\alpha} : t) \sim \vec{v}, \mathsf{seq}(\vec{\alpha} : s)}$$

Fig. 6. Inference rules for equivalence of sequences.

F. Sequences

We often use sequences when reasoning on protocols. *e.g.* to enrich equivalences with collections of messages that are already known to the attacker or can safely be disclosed as far as the property being proved is concerned. We give in Figure 6 two new rules that are sometimes necessary to prove equivalences between sequences.

The SPLIT-SEQ rule allow to split sequences in two, according to some arbitrary meta-formula ϕ on the left and ψ on the right. The CONST-SEQ rule can be used to simplify a sequence $\mathsf{seq}(\vec{\alpha} : t)$ into a finite collection $t_1, \ldots, t_n$ when it can be proved that all instances of t fall into this finite set of possibilities. Since we are dealing with an equivalence with sequences on both sides, this simplification must be done simultaneously on both sides, and we must make sure that the mapping for the instances of t to the constants t_i is the same as its counterpart on the other side of the equivalence: this is the role of the conditions b_i that partition the sequences. As expected, these rules are sound.

Proposition 3. *The rules in Figure 6 are sound.*

VI. CASE STUDY: YUBIKEY AND YUBIHSM

This section discusses the YubiKey and the YubiHSM protocols following the security analysis done in [19], [26] using the symbolic prover TAMARIN. The analysis of YubiHSM in TAMARIN requires the use of the interactive mode, and took the authors around one month. Doing this analysis in SQUIRREL allows to test our extensions of the tool and to provide stronger computational guarantees. Conducting this analysis in SQUIRREL required a similar amount of human effort than the TAMARIN analysis. The corresponding SQUIRREL files, `yubikey.sp` (270 LoC) and `yubihsm.sp` (720 LoC), can be found at [16].

A. Background

The YubiKey is a simple physical authentication device with a unique button. This device, manufactured by Yubico, allows users to securely authenticate to their accounts by issuing

a one-time password (OTP). In its typical configuration, the YubiKey generates a random OTP by encrypting a secret value and a counter. This message is accepted by the server only if it decrypts under the correct key to a valid secret value containing a counter whose value is larger than the last value accepted by the server for that YubiKey.

We first analyze the security of the protocol assuming that the server remains secure. Then, we consider an adversary having access to the authentication server. To provide some security guarantees in this scenario, a Hardware Security Module (HSM), called YubiHSM, is used to protect the working secret symmetric keys.

As in [19], [26], instead of having a session counter and a token counter, we consider a single counter. To conduct the security analysis in TAMARIN, the authors in [19], [26] over-approximate the behaviours of the system by allowing the adversary to instantiate the rule for any counter value that is higher than the previous one. We do not rely on this approximation in our SQUIRREL security analysis.

B. YubiKey

Yubico assigns a key k, as well as a public and a secret identifier (pid, sid) to each YubiKey. The counter inside the YubiKey is incremented whenever the YubiKey is plugged in, as well as when an OTP is generated, *i.e.* when the button of the YubiKey is pressed. This OTP is obtained by encrypting the counter value and the sid of the YubiKey with its key k.

$$\text{YUBIKEY} \rightarrow \text{SERVER} : pid, \text{senc}(\langle sid, cpt\rangle, k, r)$$

Here, r is the encryption randomness. The server accepts this message if it decrypts with a legitimate key k, and leads to a valid secret value sid. Lastly, the counter value obtained through decryption has to be larger than the current value stored in the server database. After this exchange, the server updates its counter with the value just received. We consider the same three security properties as in [19], [26]:

1) *Absence of replay attacks*: the server never accepts for the same YubiKey the same counter twice.
2) *Injective correspondence*: a successful login for the YubiKey *pid* must have been preceded by a button press on this YubiKey for the same counter value, and this counter value is not involved in another successful login.
3) *Monotonicity*: the counter values associated to successful logins are monotonically increasing in time.

As expected, modeling this protocol requires us to rely on our notion of states to store and update counter values on both the YubiKey's and the server's sides.

These three security properties are reachability properties which can be expressed using local meta-formulas, and which actually do not require the use of our new REWRITE-EQUIV inference rule. Furthermore, the security properties 1 and 3 are established relying solely on counter values, while *injective correspondence* requires the use of the INT-CTXT cryptographic assumption. Indeed, it is needed to ensure that the accepted ciphertext has been issued by a legitimate YubiKey.

C. YubiHSM

In order to provide some security guarantees even in the case where the server is compromised, Yubico has developed a specific HSM, called YubiHSM. In this setting, the working keys (those used by the YubiKey) are encrypted under a master key mk stored inside the YubiHSM. The YubiHSM lets the server use them to perform some specific cryptographic operations (through a specific API) without ever having access to them in plaintext. Therefore, the interaction between the YubiKey and the server has to go through the HSM (via a secure channel). The HSM receives the otp and the $aead = \text{senc}(\langle k, \langle pid, sid\rangle\rangle, mk, r_0)$, performs the decryption operation, and returns the value of the counter inside the otp to the server. The server, depending on the value of the counter, accepts the exchange or not.

$$
\begin{array}{rcll}
\text{YUBIKEY} & \rightarrow & \text{SERVER} : & pid, otp \\
\text{SERVER} & \xrightarrow{\text{secure}} & \text{HSM} : & pid, aead, otp \\
\text{HSM} & \xrightarrow{\text{secure}} & \text{SERVER} : & cpt \\
\text{SERVER} & \rightarrow & \text{YUBIKEY} : & accept
\end{array}
$$

The idea is that since the master key of the YubiHSM is not extractable, the attacker will never learn the value of the working keys (which are protected by mk) even if the server is under the control of the attacker.

We analyze the three same security properties. However, in this new setting, the proof of injective correspondence cannot be carried out in the same fashion, because the INT-CTXT cryptographic assumption on the working keys can no longer be applied. Indeed, these keys occur in plaintext position in, *e.g.*, $aead$. Therefore, to conduct this proof, we proceed in two steps. We first establish the equivalence between the original system and an ideal one, described below. We then prove the security property on the ideal system, using the same proof techniques as in the previous subsection. Thanks to our new rule REWRITE-EQUIV, we put together these two steps and conclude that the security property holds on the original system. In the ideal system, we replace the keys encrypted in $aead$ by dummy keys, so as to be able to rely on the INT-CTXT cryptographic assumption. To maintain the functionality on the ideal system, we replace the dummy keys by the real keys when we need to use them.

To establish the equivalence between these two systems, we first enrich the knowledge of the attacker by revealing additional terms, *e.g.* the secret identifiers and the working keys. Then, we show that the equivalence holds by induction on the length of the trace, with the help of the new BI-DEDUCE rule. Revealing additional information allows us to establish that the additional message outputted during the last step does not provide additional knowledge to the attacker.

VII. CONCLUSION

In order to reason on protocols with mutable states, we have improved both the theoretical foundations and the practical implementation of the SQUIRREL prover. In particular, we have presented a richer and more faithful semantics, an extended proof system that allows a rich interplay between

reachability and equivalence properties, and automated deduction techniques based on the notion of bi-deduction.

These improvements have made it possible to carry out new security analyses. We have been able to prove properties of several RFID protocols, overcoming some difficulties tied to a proper modeling of the random oracle model. We have finally carried out the largest SQUIRREL development so far, obtaining the first computational proofs of the YubiKey and YubiHSM protocols.

Future work will obviously include the development of more complex case studies, removing various simplifications in the present ones or tackling more complex protocols. We also plan to improve proof automation, for instance by automating common arguments such as reasonings on counters or stacks of hashes, and by leveraging general-purpose tools such as SMT solvers. We conjecture that our bi-deduction proof system subsumes the FA-DUP rule from [18], which needs to be formally established, and exploited in the implementation of the prover. More generally, we plan to study decidability and completeness issues for our bi-deduction proof system.

REFERENCES

[1] A. Armando, R. Carbone, L. Compagna, J. Cuéllar, G. Pellegrino, and A. Sorniotti, "An authentication flaw in browser-based single sign-on protocols: Impact and remediations," *Computers & Security*, vol. 33, pp. 41–58, 2013.

[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. Zanella-Béguelin, and P. Zimmermann, "Imperfect forward secrecy: How Diffie-Hellman fails in practice," in *22nd ACM Conference on Computer and Communications Security*, Oct. 2015.

[3] K. Bhargavan, B. Blanchet, and N. Kobeissi, "Verified models and reference implementations for the TLS 1.3 standard candidate," in *2017 IEEE Symposium on Security and Privacy*. IEEE, 2017, pp. 483–502.

[4] S. Meier, B. Schmidt, C. Cremers, and D. A. Basin, "The TAMARIN prover for the symbolic analysis of security protocols," in *CAV*, ser. Lecture Notes in Computer Science, vol. 8044. Springer, 2013, pp. 696–701.

[5] B. Blanchet, "Modeling and Verifying Security Protocols with the Applied Pi Calculus and ProVerif," *Foundations and Trends in Privacy and Security* , vol. 1, no. 1-2, pp. 1 – 135, 10 2016. [Online]. Available: https://hal.inria.fr/hal-01423760

[6] R. Chadha, V. Cheval, Ş. Ciobâcă, and S. Kremer, "Automated verification of equivalence properties of cryptographic protocols," *TOCL*, vol. 17, no. 4, pp. 1–32, 2016.

[7] V. Cheval, S. Kremer, and I. Rakotonirina, "The DEEPSEC prover," in *CAV (2)*, ser. Lecture Notes in Computer Science, vol. 10982. Springer, 2018, pp. 28–36.

[8] K. Bhargavan, A. Bichhawat, Q. Do, P. Hosseyni, R. Küsters, G. Schmitz, and T. Würtele, "DY*: a modular symbolic verification framework for executable cryptographic protocol code," in *EuroS&P 2021-6th IEEE European Symposium on Security and Privacy*, 2021.

[9] B. Blanchet, "A computationally sound mechanized prover for security protocols," in *IEEE Symposium on Security and Privacy*. IEEE Computer Society, 2006, pp. 140–154.

[10] G. Barthe, B. Grégoire, S. Heraud, 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. Cryptology*, vol. 33, no. 2, pp. 494–566, 2020.

[12] V. Cortier, S. Kremer, and B. Warinschi, "A survey of symbolic methods in computational analysis of cryptographic systems," *J. Autom. Reasoning*, vol. 46, no. 3-4, pp. 225–259, 2011.

[13] B. Blanchet, "An efficient cryptographic protocol verifier based on prolog rules," in *CSFW*. IEEE Computer Society, 2001, pp. 82–96.

[14] A. Armando, D. Basin, Y. Boichut, Y. Chevalier, L. Compagna, J. Cuéllar, P. H. Drielsma, P.-C. Héam, O. Kouchnarenko, J. Mantovani *et al.*, "The AVISPA tool for the automated validation of internet security protocols and applications," in *CAV*. Springer, 2005, pp. 281–285.

[15] V. Cheval, V. Cortier, and M. Turuani, "A little more conversation, a little less action, a lot more satisfaction: Global states in proverif," in *31st IEEE Computer Security Foundations Symposium, CSF 2018, Oxford, United Kingdom, July 9-12, 2018*. IEEE Computer Society, 2018, pp. 344–358. [Online]. Available: https://doi.org/10.1109/CSF.2018.00032

[16] The Squirrel Prover repository. https://github.com/squirrel-prover/squirrel-prover/.

[17] G. Bana and H. Comon-Lundh, "A computationally complete symbolic attacker for equivalence properties," in *CCS*. ACM, 2014, pp. 609–620.

[18] D. Baelde, S. Delaune, C. Jacomme, A. Koutsos, and S. Moreau, "An Interactive Prover for Protocol Verification in the Computational Model," in *SP 2021 - 42nd IEEE Symposium on Security and Privacy*, San Fransisco / Virtual, United States, May 2021, p. t.b.d.

[19] R. Künnemann and G. Steel, "Yubisecure? formal security analysis results for the yubikey and yubiHSM," in *Proc. 8th International Workshop on Security and Trust Management (STM'12)*, ser. Lecture Notes in Computer Science, vol. 7783. Springer, 2012, pp. 257–272.

[20] D. Baelde, S. Delaune, A. Koutsos, and S. Moreau, "Cracking the stateful nut – computational proofs of stateful security protocols using the SQUIRREL proof assistant," HAL report 03500056, Tech. Rep., 2022.

[21] G. Bana and H. Comon-Lundh, "Towards unconditional soundness: Computationally complete symbolic attacker," in *POST*, ser. Lecture Notes in Computer Science, vol. 7215. Springer, 2012, pp. 189–208.

[22] M. Ohkubo, K. Suzuki, S. Kinoshita *et al.*, "Cryptographic approach to "privacy-friendly" tags," in *RFID privacy workshop*, vol. 82. Cambridge, USA, 2003.

[23] O. Goldreich, S. Goldwasser, and S. Micali, "How to construct random functions," *J. ACM*, vol. 33, no. 4, pp. 792–807, 1986.

[24] T. van Deursen and S. Radomirovic, "Attacks on RFID protocols," *IACR Cryptol. ePrint Arch.*, vol. 2008, p. 310, 2008. [Online]. Available: http://eprint.iacr.org/2008/310

[25] P. Cousot and R. Cousot, "Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints," in *POPL*. ACM, 1977, pp. 238–252.

[26] R. Künnemann, "Foundations for analyzing security APIs in the symbolic and computational model. (fondations d'analyse de APIs de sécurité dans le modèle symbolique et calculatoire)," Ph.D. dissertation, École normale supérieure de Cachan, France, 2014. [Online]. Available: https://tel.archives-ouvertes.fr/tel-00942459

[27] A. Koutsos, "Preuves symboliques de propriétés d'indistinguabilité calculatoire. (symbolic proofs of computational indistinguishability)," Ph.D. dissertation, University of Paris-Saclay, France, 2019.

APPENDIX A
SEQUENT CALCULUS PROOF SYSTEM

We provide here a more complete presentation of the core rules for our reachability and general sequents, before establishing the soundness of the more complex ones.

A. Rules

Figure 7 shows classical reasoning rules for our reachability sequents. These rules are straightforward adaptations of the rules of [18]: the new context Θ is added but does not really play a role. We omit the rules for the existential quantifier, disjunction and conjunction, which are unsurprising.

Figure 8 shows classical reasoning rules for our generalized sequents. These rules (unlike the previous ones) are entirely unsurprising: global meta-formulas translate to first-order logic formulas, hence general sequents can be equipped with the usual rules of first-order classical logic. As before, the rules for the existential quantifier, conjunction and disjunction are not shown.

$$\frac{\Sigma, \alpha; \Theta : \Gamma \vdash_{\mathcal{P}} \phi}{\Sigma; \Theta : \Gamma \vdash_{\mathcal{P}} \forall \alpha.\phi} \qquad \frac{\Sigma; \Theta : \Gamma, \phi\{\tau \to T\} \vdash_{\mathcal{P}} \psi}{\Sigma; \Theta : \Gamma, \forall \tau.\phi \vdash_{\mathcal{P}} \psi}$$

$$\frac{\Sigma; \Theta : \Gamma, \psi \vdash_{\mathcal{P}} \phi}{\Sigma; \Theta : \Gamma \vdash_{\mathcal{P}} \psi \Rightarrow \phi} \qquad \frac{\Sigma; \Theta : \Gamma \vdash_{\mathcal{P}} \phi \quad \Sigma; \Theta : \Gamma, \phi' \vdash_{\mathcal{P}} \psi}{\Sigma; \Theta : \Gamma, \phi \Rightarrow \phi' \vdash_{\mathcal{P}} \psi}$$

$$\overline{\Sigma; \Theta : \Gamma, \phi \vdash_{\mathcal{P}} \phi} \qquad \overline{\Sigma; \Theta : \Gamma, \bot \vdash_{\mathcal{P}} \phi}$$

$$\frac{\Sigma; \Theta : \Gamma, \phi \Rightarrow \bot \vdash_{\mathcal{P}} \bot}{\Sigma; \Theta : \Gamma \vdash_{\mathcal{P}} \phi}$$

The right rule for $\forall$ assumes that $\alpha \notin \Sigma$. The left rule for $\forall$ has a variant for sort index but not message, since quantification over messages is not allowed in local meta-formulas.

Fig. 7. Classical reasoning for reachability sequents.

$$\frac{\Sigma, \alpha; \Theta \vdash F}{\Sigma; \Theta \vdash \tilde{\forall}\alpha.F} \qquad \frac{\Sigma; \Theta, F\{x \to t\} \vdash G}{\Sigma; \Theta, \tilde{\forall}x.F \vdash G}$$

$$\frac{\Sigma; \Theta, F \vdash G}{\Sigma; \Theta \vdash F \tilde{\Rightarrow} G} \qquad \frac{\Sigma; \Theta \vdash F \quad \Sigma; \Theta, F' \vdash G}{\Sigma; \Theta, F \tilde{\Rightarrow} F' \vdash G}$$

$$\overline{\Sigma; \Theta, F \vdash F} \qquad \overline{\Sigma; \Theta, \tilde{\bot} \vdash F} \qquad \frac{\Sigma; \Theta, F \tilde{\Rightarrow} \tilde{\bot} \vdash \tilde{\bot}}{\Sigma; \Theta \vdash F}$$

The right rule for $\tilde{\forall}$ assumes that $\alpha \notin \Sigma$. The left rule for $\tilde{\forall}$ is shown for sort message but has variants for sorts index and timestamp.

Fig. 8. Classical reasoning for general sequents.

$$\frac{\Sigma; \Theta : \Gamma \vdash_{\mathcal{P}} \forall \tau. (\forall \tau'. \tau' < \tau \Rightarrow \phi) \Rightarrow \phi}{\Sigma; \Theta : \Gamma \vdash_{\mathcal{P}} \forall \tau.\phi}$$

$$\frac{\Sigma; \Theta \vdash \tilde{\forall}\tau. (\tilde{\forall}\tau'. [\tau' < \tau]_{\mathcal{P}} \tilde{\Rightarrow} F) \tilde{\Rightarrow} F}{\Sigma; \Theta \vdash \tilde{\forall}\tau.F}$$

Fig. 9. Induction rules.

$$\frac{\Sigma; \Theta, [\phi]_{\mathcal{P}}, [\phi']_{\mathcal{P}} \vdash F}{\Sigma; \Theta, [\phi \wedge \phi']_{\mathcal{P}} \vdash F} \qquad \frac{\Sigma; \Theta, [\phi\{x \to t\}]_{\mathcal{P}} \vdash F}{\Sigma; \Theta, [\forall x.\phi]_{\mathcal{P}} \vdash F}$$

$$\frac{\Sigma; \Theta \vdash [\phi]_{\mathcal{P}}}{\Sigma; \Theta, [\neg\phi]_{\mathcal{P}} \vdash F}$$

$$\frac{\Sigma; \Theta, [\phi]_{\mathcal{P}} \vdash F \quad \Sigma; \Theta, [\phi']_{\mathcal{P}} \vdash F}{\Sigma; \Theta, [\phi \vee \phi']_{\mathcal{P}} \vdash F}$$

$$\frac{\Sigma; \Theta \vdash [\phi]_{\mathcal{P}} \quad \Sigma; \Theta, [\phi']_{\mathcal{P}} \vdash F}{\Sigma; \Theta, [\phi \Rightarrow \phi']_{\mathcal{P}} \vdash F}$$

$$\frac{\Sigma; \Theta, [\phi]_{\mathcal{P}} \vdash F \quad \Sigma; \Theta, [\phi']_{\mathcal{P}} \vdash F}{\Sigma; \Theta, [\phi \vee \phi']_{\mathcal{P}} \vdash F}$$

$$\frac{\Sigma, x; \Theta, [\phi]_{\mathcal{P}} \vdash F}{\Sigma; \Theta, [\exists x.\phi]_{\mathcal{P}} \vdash F}$$

In the disjunction and implication rules, one of ϕ or ϕ' must be a pure trace formula. In the existential quantifier rule, ϕ must be a pure trace formula.

Fig. 10. Rules for reasoning on local meta-formula atoms in general sequents.

Figure 9 shows two induction rules: the first one, for reachability sequents, is an immediate adaptation of the corresponding rule in [18]; the second rule, for generalized sequents, significantly generalizes the induction rule for equivalence sequents of [18] since it may be used with a global meta-formula F that is not restricted to an equivalence atom.

We finally provide in Figure 10 some extra rules which allow reasoning on local meta-formula atoms within general sequents. These rules make a crucial use of the notion of pure trace formula, as discussed in Section V. We only consider left rules here as such atoms on the right are simply handled using LOCAL-GLOBAL and the classical rules for reachability sequents.

B. Soundness

In this section we show the soundness of our most complex rules, *i.e.* the rules of Figure 6 for reasoning on sequences. To this end, we first define base logic sequents and rules.

A base logic sequent is an element of the form:

$$\Sigma_b; \Theta_b \vdash_b \phi$$

where ϕ is a base formula, Θ_b is a sequence of base formulas and Σ_b is a sequence of variables of sort message. We require that base logic sequents are closed, *i.e.* Σ_b must bind all variables in Θ_b, ϕ. Quite naturally, $\Sigma_b; \Theta_b \vdash_b \phi$ is valid iff $\forall \Sigma_b. \Theta_b \Rightarrow \phi$ is valid.

We shall consider the rules of Figure 11 for base logic sequents. We omit the proof of soundness for these simple rules: the first two rules are generalizations of the corresponding rules in [27]; and the soundness of the last rule is straightforward exploiting the properties of the pair.

We are now ready to establish the soundness of our rules for sequences.

Proposition 3. *The rules in Figure 6 are sound.*

Proof. A) The SPLIT-SEQ *rule:* Assume that the following

15

$$
\frac{
\begin{array}{c}
\textsc{Equ} \\
\Sigma_b;\ \Theta_b \vdash_b \vec{u}_1, t_1 \sim \vec{u}_2, v_2 \\
\Sigma_b;\ \Theta_b \vdash_b \mathsf{EQ}(v_1, t_1) \sim \mathsf{true}
\end{array}
}{
\Sigma_b;\ \Theta_b \vdash_b \vec{u}_1, v_1 \sim \vec{u}_2, v_2
}
$$

$$
\frac{
\begin{array}{c}
\textsc{FaB} \\
\Sigma_b;\ \Theta_b \vdash_b \vec{u}_1, \vec{v}_1 \sim \vec{u}_2, \vec{v}_2
\end{array}
}{
\Sigma_b;\ \Theta_b \vdash_b \vec{u}_1, f(\vec{v}_1) \sim \vec{u}_2, f(\vec{v}_2)
}
$$

$$
\frac{
\begin{array}{c}
\textsc{Inline-Pair} \\
\Sigma_b;\ \Theta_b \vdash_b \vec{u}_0, \langle s_0, t_0 \rangle \sim \vec{u}_1, \langle s_1, t_1 \rangle
\end{array}
}{
\Sigma_b;\ \Theta_b \vdash_b \vec{u}_0, s_0, t_0 \sim \vec{u}_1, s_1, t_1
}
$$

Fig. 11. Rules for base logic sequents.

sequent is valid:

$$
\Sigma;\ \Theta \vdash_{\mathcal{P},\mathcal{P}'} \sim \begin{array}{l} \vec{u}, \mathsf{seq}(\vec{\alpha} : \mathsf{if}\ \phi\ \mathsf{then}\ t), \mathsf{seq}(\vec{\alpha} : \mathsf{if}\ \neg\phi\ \mathsf{then}\ t) \\ \vec{v}, \mathsf{seq}(\vec{\alpha} : \mathsf{if}\ \psi\ \mathsf{then}\ s), \mathsf{seq}(\vec{\alpha} : \mathsf{if}\ \neg\psi\ \mathsf{then}\ s) \end{array} \quad (5)
$$

We have to show that the following sequent is valid:

$$
\Sigma;\ \Theta \vdash_{\mathcal{P},\mathcal{P}'} \vec{u}, \mathsf{seq}(\vec{\alpha} : t) \sim \vec{v}, \mathsf{seq}(\vec{\alpha} : s) \quad (6)
$$

First, for any base formula ϕ_b and base term x, if we let $(x_1 \overset{\text{def}}{=} \mathsf{if}\ \phi_b\ \mathsf{then}\ x)$ and $(x_2 \overset{\text{def}}{=} \mathsf{if}\ \neg\phi_b\ \mathsf{then}\ x)$, then the base logic rule below is sound:

$$
\frac{}{
x, x_1, x_2;\ \Theta_b \vdash_b \begin{array}{c} \mathsf{EQ}(x, \mathsf{if}\ \mathsf{EQ}(x_2, \mathsf{empty})\ \mathsf{then}\ x_1\ \mathsf{else}\ x_2) \\ \sim \mathsf{true} \end{array}
} \quad (7)
$$

Indeed, the term on the left of the indistinguishability predicate is always true, regardless of probabilistic samplings.

To prove that the validity of Equ. (6) follows from the validity of Equ. (5), we essentially proceed as follows:

- we rewrite all terms in the sequence of the sequent in Equ. (6) using Equ. (7);
- we get rid of the if _ then _ else _ and empty function symbols using the the FaB rule;
- we conclude by observing that all the resulting terms appear in the sequences in Equ. (5).

We now detail this proof. We let $\Sigma = \Sigma_m \uplus \Sigma_{it}$ where Σ_m is a sequence of variables of sort message, and Σ_{it} is a sequence of variables of sort index or timestamp. Then the meta-logic sequent in Equ. (6) is valid iff for any trace model $\mathbb{T}$, the base logic sequent

$$
\Sigma_m;\ (\Theta)^{\mathbb{T}} \vdash_b (\vec{u})^{\mathbb{T}}_{\mathcal{P}}, (\mathsf{seq}(\vec{\alpha} : t))^{\mathbb{T}}_{\mathcal{P}} \sim (\vec{v})^{\mathbb{T}}_{\mathcal{P}'}, (\mathsf{seq}(\vec{\alpha} : s))^{\mathbb{T}}_{\mathcal{P}'}
$$

is valid.

Let $\vec{l}_1, \ldots, \vec{l}_n$ be an enumeration of all possible values of $\vec{\alpha}$ in $\mathbb{T}$. Let $\mathbb{T}_i = \mathbb{T}[\vec{\alpha} \mapsto \vec{l}_i]$. Then, from the definition of the interpretation of a sequence, the sequent above is equal to:

$$
\Sigma_m;\ (\Theta)^{\mathbb{T}} \vdash_b \begin{array}{l} (\vec{u})^{\mathbb{T}}_{\mathcal{P}}, \langle (t)^{\mathbb{T}_1}_{\mathcal{P}}, \langle \ldots, (t)^{\mathbb{T}_n}_{\mathcal{P}} \rangle \ldots \rangle \sim \\ (\vec{v})^{\mathbb{T}}_{\mathcal{P}'}, \langle (s)^{\mathbb{T}_1}_{\mathcal{P}'}, \langle \ldots, (s)^{\mathbb{T}_n}_{\mathcal{P}'} \rangle \ldots \rangle \end{array}
$$

Using the base logic rule FaB several times on the pair symbol, it is sufficient to prove that:

$$
\Sigma_m;\ (\Theta)^{\mathbb{T}} \vdash_b \begin{array}{l} (\vec{u})^{\mathbb{T}}_{\mathcal{P}}, (t)^{\mathbb{T}_1}_{\mathcal{P}}, \ldots, (t)^{\mathbb{T}_n}_{\mathcal{P}} \sim \\ (\vec{v})^{\mathbb{T}}_{\mathcal{P}'}, (s)^{\mathbb{T}_1}_{\mathcal{P}'}, \ldots, (s)^{\mathbb{T}_n}_{\mathcal{P}'} \end{array} \quad (\star)
$$

For all $1 \le i \le n$, we let:

$$
t^i_1 \overset{\text{def}}{=} \mathsf{if}\ (\phi)^{\mathbb{T}_i}_{\mathcal{P}}\ \mathsf{then}\ (t)^{\mathbb{T}_i}_{\mathcal{P}} \qquad t^i_2 \overset{\text{def}}{=} \mathsf{if}\ \neg(\phi)^{\mathbb{T}_i}_{\mathcal{P}}\ \mathsf{then}\ (t)^{\mathbb{T}_i}_{\mathcal{P}}
$$

and:

$$
t^i \overset{\text{def}}{=} \mathsf{if}\ \mathsf{EQ}(t^i_2, \mathsf{empty})\ \mathsf{then}\ t^i_1\ \mathsf{else}\ t^i_2
$$

We define similarly s^i_1, s^i_2 and s^i using term s, condition ψ and protocol $\mathcal{P}'$. Then, using the Equ rule and Equ. (7) n times on each side, we only have to prove that the following sequent is valid:

$$
\Sigma_m;\ (\Theta)^{\mathbb{T}} \vdash_b (\vec{u})^{\mathbb{T}}_{\mathcal{P}}, t^1, \ldots, t^n \sim (\vec{v})^{\mathbb{T}}_{\mathcal{P}'}, s^1, \ldots, s^n
$$

Using the FaB rule on the function symbols if _ then _ else _, EQ and empty, and the Dup rule to merge duplicates, it is sufficient to prove that:

$$
\Sigma_m;\ (\Theta)^{\mathbb{T}} \vdash_b \begin{array}{l} (\vec{u})^{\mathbb{T}}_{\mathcal{P}}, t^1_1, \ldots, t^n_1, t^1_2, \ldots, t^n_2 \sim \\ (\vec{v})^{\mathbb{T}}_{\mathcal{P}'}, s^1_1, \ldots, s^n_1, s^1_2, \ldots, s^n_2 \end{array}
$$

Using Inline-Pair, we get:

$$
\Sigma_m;\ (\Theta)^{\mathbb{T}} \vdash_b \begin{array}{l} (\vec{u})^{\mathbb{T}}_{\mathcal{P}}, \langle t^1_1, \langle \ldots, t^n_1 \rangle \ldots \rangle, \langle t^1_2, \langle \ldots, t^n_2 \rangle \ldots \rangle \sim \\ (\vec{v})^{\mathbb{T}}_{\mathcal{P}'}, \langle s^1_1, \langle \ldots, s^n_1 \rangle \ldots \rangle, \langle s^1_2, \langle \ldots, s^n_2 \rangle \ldots \rangle \end{array}
$$

We conclude by observing that the sequent above is entailed by the translation in $\mathbb{T}$ of the sequent in Equ. (5), which is valid by hypothesis.

B) The Const-Seq *rule:* Assume that the following sequents are valid:

$$
\Sigma;\ \Theta \vdash \left[\forall \vec{\alpha}.\ \bigvee_{0 \le i \le l} b_i\right]_{\mathcal{P}} \quad (8)
$$

$$
\Sigma;\ \Theta \vdash \left[\bigwedge_{0 \le i \le l} \forall \vec{\alpha}.\ b_i \Rightarrow t = t_i\right]_{\mathcal{P}} \quad (9)
$$

$$
\Sigma;\ \Theta \vdash \left[\bigwedge_{0 \le i \le l} \forall \vec{\alpha}.\ b_i \Rightarrow s = s_i\right]_{\mathcal{P}'} \quad (10)
$$

$$
\Sigma;\ \Theta \vdash_{\mathcal{P},\mathcal{P}'} \vec{u}, t_1, \ldots, t_l \sim \vec{v}, s_1, \ldots, s_l \quad (11)
$$

where $\vec{\alpha}$ are variables of sort index or timestamp, and that:

$$
(b_i)_{1 \le i \le l} \text{ are pure trace formulas} \quad (12)
$$

We have to show that the following sequent is valid:

$$
\Sigma;\ \Theta \vdash_{\mathcal{P},\mathcal{P}'} \vec{u}, \mathsf{seq}(\vec{\alpha} : t) \sim \vec{v}, \mathsf{seq}(\vec{\alpha} : s) \quad (13)
$$

Starting our proof exactly like for the Split-Seq rule until Equ. $(\star)$, and using the same notations, it is sufficient to prove that for any trace model $\mathbb{T}$:

$$
\Sigma_m;\ (\Theta)^{\mathbb{T}} \vdash_b \begin{array}{l} (\vec{u})^{\mathbb{T}}_{\mathcal{P}}, (t)^{\mathbb{T}_1}_{\mathcal{P}}, \ldots, (t)^{\mathbb{T}_n}_{\mathcal{P}} \sim \\ (\vec{v})^{\mathbb{T}}_{\mathcal{P}'}, (s)^{\mathbb{T}_1}_{\mathcal{P}'}, \ldots, (s)^{\mathbb{T}_n}_{\mathcal{P}'} \end{array} \quad (14)
$$

For every $1 \le i \le n$, we let $t^i \overset{\text{def}}{=} (t)^{\mathbb{T}_i}_{\mathcal{P}}$ and

$$
t'^i \overset{\text{def}}{=} \mathsf{if}\ (b_1)^{\mathbb{T}_i}_{\mathcal{P}}\ \mathsf{then}\ (t_1)^{\mathbb{T}_i}_{\mathcal{P}}\ \mathsf{else}
$$

$$
\ldots
$$

$$
\mathsf{if}\ (b_l)^{\mathbb{T}_i}_{\mathcal{P}}\ \mathsf{then}\ (t_l)^{\mathbb{T}_i}_{\mathcal{P}}\ \mathsf{else}\ \mathsf{empty}
$$

16

We define similarly s^i and s'^i. Using Equ. (8) and Equ. (9), we can show that:

$$\Sigma_m; (\Theta)^{\mathbb{T}} \vdash_b \mathsf{EQ}(t^i, t'^i) \sim \mathsf{true} \tag{15}$$

Since $(b_i)_{1 \le i \le l}$ are pure trace formulas (Equ. (12)), we know that their interpretation in a trace model is independent of the protocol. Hence, from Equ. (8), we get that:

$$\Sigma; \Theta \vdash \left[\forall \vec{\alpha}, \bigvee_{0 \le i \le l} b_i \right]_{\mathcal{P}'} \tag{16}$$

Hence, similarly to the proof of Equ. (15), we can show, using Equ. (16) and Equ. (10), that:

$$\Sigma_m; (\Theta)^{\mathbb{T}} \vdash_b \mathsf{EQ}(s^i, s'^i) \sim \mathsf{true}$$

Hence, using the EQU rule several times, the validity of Equ. (14) follows from that of:

$$\Sigma_m; (\Theta)^{\mathbb{T}} \vdash_b (\vec{u})_{\mathcal{P}}^{\mathbb{T}}, t'^1 \dots, t'^n \sim (\vec{v})_{\mathcal{P}'}^{\mathbb{T}}, s'^1 \dots, s'^n \tag{17}$$

We can assume that the variables $\vec{\alpha}$ are distinct from other variables, without loss of generality since we can always rename bound variables:

$$\left(\bigcup_{1 \le j \le l} \mathsf{fv}(t_j) \right) \cap \vec{\alpha} = \emptyset \tag{18}$$

Hence, since for every $1 \le j \le l$, t_j's free variables are disjoint from $\vec{\alpha}$, for any i such that $1 \le i \le n$, we have that:

$$(t_j)_{\mathcal{P}}^{\mathbb{T}_i} \stackrel{\mathsf{def}}{=} (t_j)_{\mathcal{P}}^{\mathbb{T}[\vec{\alpha} \mapsto \vec{l_i}]} = (t_j)_{\mathcal{P}}^{\mathbb{T}} \tag{19}$$

Similarly, we show that:

$$(s_j)_{\mathcal{P}}^{\mathbb{T}_i} \stackrel{\mathsf{def}}{=} (s_j)_{\mathcal{P}}^{\mathbb{T}[\vec{\alpha} \mapsto \vec{l_i}]} = (s_j)_{\mathcal{P}}^{\mathbb{T}} \tag{20}$$

Starting from Equ. (17), and applying the FAB rule to get rid of the if _ then _ else _ and empty function symbols, it is sufficient to show that:

$$\Sigma_m; (\Theta)^{\mathbb{T}} \vdash_b (\vec{u})_{\mathcal{P}}^{\mathbb{T}}, \left((t_j)_{\mathcal{P}}^{\mathbb{T}_i}, (b_j)_{\mathcal{P}}^{\mathbb{T}_i} \right)_{1 \le j \le l, 1 \le i \le n} \sim$$
$$(\vec{v})_{\mathcal{P}'}^{\mathbb{T}}, \left((s_j)_{\mathcal{P}'}^{\mathbb{T}_i}, (b_j)_{\mathcal{P}'}^{\mathbb{T}_i} \right)_{1 \le j \le l, 1 \le i \le n}$$

Using Equ. (19) and Equ. (20), and the DUP to get rid of duplicates, we get:

$$\Sigma_m; (\Theta)^{\mathbb{T}} \vdash_b (\vec{u})_{\mathcal{P}}^{\mathbb{T}}, (t_1)_{\mathcal{P}}^{\mathbb{T}}, \dots, (t_l)_{\mathcal{P}}^{\mathbb{T}}, \left((b_j)_{\mathcal{P}}^{\mathbb{T}_i} \right)_{1 \le j \le l, 1 \le i \le n}$$
$$\sim$$
$$(\vec{v})_{\mathcal{P}'}^{\mathbb{T}}, (s_1)_{\mathcal{P}'}^{\mathbb{T}}, \dots, (s_l)_{\mathcal{P}'}^{\mathbb{T}}, \left((b_j)_{\mathcal{P}'}^{\mathbb{T}_i} \right)_{1 \le j \le l, 1 \le i \le n}$$

Since, for every $1 \le j \le l$, b_j is a pure trace formula, we know that for every $1 \le i \le n$, $(b_j)_{\mathcal{P}}^{\mathbb{T}_i} = (b_j)_{\mathcal{P}'}^{\mathbb{T}_i}$, and does not contain any names. Therefore, we can remove them using FAB:

$$\Sigma_m; (\Theta)^{\mathbb{T}} \vdash_b (\vec{u})_{\mathcal{P}}^{\mathbb{T}}, (t_1)_{\mathcal{P}}^{\mathbb{T}}, \dots, (t_l)_{\mathcal{P}}^{\mathbb{T}} \sim$$
$$(\vec{v})_{\mathcal{P}'}^{\mathbb{T}}, (s_1)_{\mathcal{P}'}^{\mathbb{T}}, \dots, (s_l)_{\mathcal{P}'}^{\mathbb{T}}$$

This concludes the proof since the base logic sequent above is equal to the translation in $\mathbb{T}$ of the sequent in Equ. (11), which is valid by hypothesis. $\square$

Using the notations of Figure 2:

$$\mathsf{unfold}_{\mathcal{P}}(m, a[\vec{i}_0]) = (m_{a[\vec{i}]})\{\vec{i} \mapsto \vec{i}_0\}$$
$$\mathsf{unfold}_{\mathcal{P}}(m, \mathsf{init}) = m_{\mathsf{init}}$$
$$\mathsf{unfold}_{\mathcal{P}}(s[\vec{j}_0], a[\vec{i}_0]) = (u_{a[\vec{i}], s[\vec{j}]})\{\vec{i} \mapsto \vec{i}_0, \vec{j} \mapsto \vec{j}_0\}$$
$$\mathsf{unfold}_{\mathcal{P}}(s[\vec{j}_0], \mathsf{init}) = (u_{0, s[\vec{j}]})\{\vec{j} \mapsto \vec{j}_0\}$$

Fig. 12. Interpretation of the partial macro evaluation function $\mathsf{unfold}_{\mathcal{P}}(\cdot, \cdot)$.

APPENDIX B
BI-DEDUCTION

We extend bi-deduction to bi-formulas in a straightforward manner. For any sequence Σ of variables of any sort, protocols $\mathcal{P}_1, \mathcal{P}_2$, set Θ of global meta-formulas, sequences of bi-terms $\#(\vec{u_1}, \vec{u_2})$, bi-formulas $\#(\lambda_1, \lambda_2)$ and $\#(\phi_1, \phi_2)$, we have a bi-deduction judgement:

$$\Sigma; \Theta; \#(\vec{u_1}, \vec{u_2}) \triangleright_{\mathcal{P}_1, \mathcal{P}_2} \{\#(\lambda_1, \lambda_2) \mid \#(\phi_1, \phi_2)\}.$$

The validity of this judgement is defined in the same way as the validity of a bi-deduction judgement for a bi-term (Equ. (2)): the machine $\mathcal{B}_c$ must compute ϕ_i and $\mathcal{B}_o$ must compute (if ϕ_i then λ_i), i.e. $(\phi_i \Rightarrow \lambda_i)$.

a) Unfold: Before giving our rules, we define the partial macro evaluation function $\mathsf{unfold}_{\mathcal{P}}(\cdot, \cdot)$ in Figure 12.

b) Rules: We now present the bi-deduction rules of Figure 13, which are grouped in three categories, depending on whether they act on the left, right, or both sides of the $\triangleright$ symbol.

We start by describing left rules. The WEAK-L rule states that if $v_{\#}$ can be bi-deduced from $\vec{u}_{\#}^1$, then it can bi-deduced from $\vec{u}_{\#}^1, \vec{u}_{\#}^2$. SEQ-DED allows to bi-deduce a bi-term $\#(v_1, v_2)$, whenever $\#(\phi_1, \phi_2)$ holds, from $\vec{u}_{\#}, \mathsf{seq}(\vec{\alpha} : \#(u_1, u_2) \mid \#(\psi_1, \psi_2))$ if: i) $\#(\phi_1, \phi_2)$ can be bi-deduced using the allowed bi-terms; and ii), the bi-term $\#(v_1, v_2)$ appears in the sequence *simultaneously* on both sides: there must exists $\vec{\alpha}$ such that, for any $i \in \{1, 2\}$, if ϕ_i holds then $u_i = v_i$ and the sequence condition ψ_i holds. PAIR-PROJ allows to exploit the fact that a machine can recover a pair's components. Using the UNFOLD-L-I rules, we can replace a macro by its definition, as long as it is at a timestamp that happens. The FRAME rule states that frame@T contains all frame@T' for any $T' \le T$.

We now describe right rules. The rule SPLIT-COND allows to weaken the condition on the right-hand side when it is bi-deducible. The rule FA says that the image of a function $f(\vec{v}_{\#})$ is bi-deducible whenever its arguments $\vec{v}_{\#}$ are bi-deducible. In case f is of arity zero, we must also check that the condition $\phi_{\#}$ is bi-deducible — see rule FA$_0$. We then have many variants of the FA rule, for other logical constructs that are translated, in the base logic, to function symbol applications: FA-ITE for conditionals; FA-FIND for lookups; FA-$\mathcal{Q}$ for quantification

Left rules

<center>SEQ-DED</center>

$$\Sigma; \Theta \vdash \tilde{\exists} \vec{\alpha}. \left(\begin{array}{l} [\phi_1 \Rightarrow (\psi_1 \wedge u_1 = v_1)]_{\mathcal{P}_1} \tilde{\wedge} \\ [\phi_2 \Rightarrow (\psi_2 \wedge u_2 = v_2)]_{\mathcal{P}_2} \end{array} \right)$$

<center>WEAK-L</center>

$$\Sigma; \Theta; \vec{u}_\#^1 \rhd \{v_\# \mid \phi_\#\}$$

$$\frac{\Sigma; \Theta; \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: \#(u_1, u_2) \mid \#(\psi_1, \psi_2)) \rhd \#(\phi_1, \phi_2)}{\Sigma; \Theta; \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: \#(u_1, u_2) \mid \#(\psi_1, \psi_2)) \rhd \{\#(v_1, v_2) \mid \#(\phi_1, \phi_2)\}}$$

$$\frac{\Sigma; \Theta; \vec{u}_\#^1 \rhd \{v_\# \mid \phi_\#\}}{\Sigma; \Theta; \vec{u}_\#^1, \vec{u}_\#^2 \rhd \{v_\# \mid \phi_\#\}}$$

<center>PAIR-PROJ</center>

$$\frac{\Sigma; \Theta; \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: u_\#^1 \mid \psi_\#), \mathsf{seq}(\vec{\alpha}: u_\#^2 \mid \psi_\#) \rhd \{v_\# \mid \phi_\#\}}{\Sigma; \Theta; \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: \langle u_\#^1, u_\#^2 \rangle \mid \psi_\#) \rhd \{v_\# \mid \phi_\#\}}$$

<center>UNFOLD-L-1</center>

$$\Sigma; \Theta; \vec{u}_\#, \#(\mathsf{seq}(\vec{\alpha}: \mathsf{unfold}_{\mathcal{P}_1}(m, \mathsf{a}[\vec{i}]) \mid \psi), u_2) \rhd \{v_\# \mid \phi_\#\}$$

$$\frac{\Sigma, \vec{\alpha}; \Theta \vdash [\psi \Rightarrow \mathsf{happens}(\mathsf{a}[\vec{i}])]_{\mathcal{P}_1}}{\Sigma; \Theta; \vec{u}_\#, \#(\mathsf{seq}(\vec{\alpha}: m@(\mathsf{a}[\vec{i}]) \mid \psi), u_2) \rhd \{v_\# \mid \phi_\#\}}$$

<center>UNFOLD-L-2</center>

$$\Sigma; \Theta; \vec{u}_\#, \#(u_1, \mathsf{seq}(\vec{\alpha}: \mathsf{unfold}_{\mathcal{P}_2}(m, \mathsf{a}[\vec{i}]) \mid \psi)) \rhd \{v_\# \mid \phi_\#\}$$

<center>FRAME</center>

$$\frac{\Sigma, \vec{\alpha}; \Theta \vdash [\psi \Rightarrow \mathsf{happens}(\mathsf{a}[\vec{i}])]_{\mathcal{P}_2}}{\Sigma; \Theta; \vec{u}_\#, \#(u_1, \mathsf{seq}(\vec{\alpha}: m@(\mathsf{a}[\vec{i}]) \mid \psi)) \rhd \{v_\# \mid \phi_\#\}}$$

$$\frac{\Sigma; \Theta; \vec{u}_\#, \mathsf{seq}(\vec{\alpha}, T': \mathsf{frame}@T' \mid \psi_\# \wedge T' \leq T) \rhd \{v_\# \mid \phi_\#\}}{\Sigma; \Theta; \vec{u}_\#, \mathsf{seq}(\vec{\alpha}: \mathsf{frame}@T \mid \psi_\#) \rhd \{v_\# \mid \phi_\#\}}$$

Right rules

<center>SPLIT-COND</center>

$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{v_\# \mid \psi_\#\} \qquad \Sigma; \Theta; \vec{u}_\# \rhd \{\phi_\# \mid \psi_\#\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{v_\# \mid \phi_\# \wedge \psi_\#\}}$$

<center>FA</center>

$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{\vec{v}_\# \mid \phi_\#\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{f(\vec{v}_\#) \mid \phi_\#\}}$$

<center>FA$_0$</center>

$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \phi_\#}{\Sigma; \Theta; \vec{u}_\# \rhd \{f() \mid \phi_\#\}}$$

<center>FA-ITE</center>

$$\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\# \mid \phi_\#\}$$
$$\Sigma; \Theta; \vec{u}_\# \rhd \{t_\# \mid \phi_\# \wedge \psi_\#\}$$
$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{t'_\# \mid \phi_\# \wedge \neg\psi_\#\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\mathsf{if}\ \psi_\#\ \mathsf{then}\ t_\#\ \mathsf{else}\ t'_\# \mid \phi_\#\}}$$

<center>FA-FIND</center>

$$\Sigma, \vec{i}; \Theta; \vec{u}_\# \rhd \{\psi_\# \mid \phi_\#\}$$
$$\Sigma, \vec{i}; \Theta; \vec{u}_\# \rhd \{t_\# \mid \phi_\# \wedge \psi_\#\}$$
$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{t'_\# \mid \phi_\# \wedge \forall \vec{i}, \neg\psi_\#\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\mathsf{find}\ \vec{i}\ \mathsf{suchthat}\ \psi_\#\ \mathsf{in}\ t_\#\ \mathsf{else}\ t'_\# \mid \phi_\#\}}$$

<center>FA-$\mathcal{Q}$</center>

$$\frac{\Sigma, \alpha; \Theta; \vec{u}_\# \rhd \{v_\# \mid \phi_\#\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\mathcal{Q}\alpha.v_\# \mid \phi_\#\}} \ \mathcal{Q} \in \{\exists, \forall\}$$

<center>FA-$\vee$-L</center>

$$\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^0 \mid \phi_\#\}$$
$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^1 \mid \phi_\# \wedge \neg\psi_\#^0\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^0 \vee \psi_\#^1 \mid \phi_\#\}}$$

<center>FA-$\vee$-R</center>

$$\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^1 \mid \phi_\#\}$$
$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^0 \mid \phi_\# \wedge \neg\psi_\#^1\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^0 \vee \psi_\#^1 \mid \phi_\#\}}$$

<center>FA-$\wedge$-L</center>

$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^1 \mid \phi_\# \wedge \psi_\#^0\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^0 \wedge \psi_\#^1 \mid \phi_\#\}}$$

<center>FA-$\wedge$-R</center>

$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^0 \mid \phi_\# \wedge \psi_\#^1\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^0 \wedge \psi_\#^1 \mid \phi_\#\}}$$

<center>FA-$\Rightarrow$</center>

$$\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^0 \mid \phi_\#\}$$
$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^1 \mid \phi_\# \wedge \psi_\#^0\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\#^0 \Rightarrow \psi_\#^1 \mid \phi_\#\}}$$

<center>FA-$\neg$</center>

$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{\psi_\# \mid \phi_\#\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\neg\psi_\# \mid \phi_\#\}}$$

<center>FA-†</center>

$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \phi_\#}{\Sigma; \Theta; \vec{u}_\# \rhd \{\dagger \mid \phi_\#\}} \ \dagger \in \{\top, \bot\}$$

<center>FA-SEQ</center>

$$\frac{\Sigma, \vec{\alpha}; \Theta; \vec{u}_\# \rhd \{v_\# \mid \phi_\# \wedge \psi_\#\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\mathsf{seq}(\vec{\alpha}: v_\# \mid \psi_\#) \mid \phi_\#\}}$$

<center>PURE-DED</center>

$$\frac{v, \phi \text{ pure trace formula}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\#(v, v) \mid \#(\phi, \phi)\}}$$

<center>UNFOLD-R-1</center>

$$\Sigma; \Theta; \vec{u}_\# \rhd \{\#(\mathsf{unfold}_{\mathcal{P}_1}(m, \mathsf{a}[\vec{i}]), v_2) \mid \phi_\#\}$$
$$\frac{\Sigma; \Theta \vdash [\phi_1 \Rightarrow \mathsf{happens}(\mathsf{a}[\vec{i}])]_{\mathcal{P}_1}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\#(m@(\mathsf{a}[\vec{i}]), v_2) \mid \#(\phi_1, \phi_2)\}}$$

<center>UNFOLD-R-2</center>

$$\Sigma; \Theta; \vec{u}_\# \rhd \{\#(v_1, \mathsf{unfold}_{\mathcal{P}_2}(m, \mathsf{a}[\vec{i}])) \mid \phi_\#\}$$
$$\frac{\Sigma; \Theta \vdash [\phi_2 \Rightarrow \mathsf{happens}(\mathsf{a}[\vec{i}])]_{\mathcal{P}_2}}{\Sigma; \Theta; \vec{u}_\# \rhd \{\#(v_1, m@(\mathsf{a}[\vec{i}])) \mid \#(\phi_1, \phi_2)\}}$$

<center>UNFOLD-R-SYM</center>

$$\frac{\Sigma; \Theta; \vec{u}_\# \rhd \{\#(\mathsf{unfold}_{\mathcal{P}_1}(m, \mathsf{a}[\vec{i}]), \mathsf{unfold}_{\mathcal{P}_2}(m, \mathsf{a}[\vec{i}])) \mid \phi_\#\}}{\Sigma; \Theta; \vec{u}_\# \rhd \{m@(\mathsf{a}[\vec{i}]) \mid \phi_\#\}}$$

Symmetric rule

<center>CASE</center>

$$\frac{\left(\Sigma, \vec{i}; \Theta\{\tau \mapsto \mathsf{a}[\vec{i}]\}; \vec{u}_\#\{\tau \mapsto \mathsf{a}[\vec{i}]\} \rhd \{v_\#\{\tau \mapsto \mathsf{a}[\vec{i}]\} \mid \phi_\#\{\tau \mapsto \mathsf{a}[\vec{i}]\}\} \right)_{\mathsf{a} \in \mathcal{A}}}{\Sigma, \tau; \Theta; \vec{u}_\# \rhd \{v_\# \mid \phi_\#\}}$$

Conventions: $\mathcal{A}$ is the set of action names of $\mathcal{P}_1, \mathcal{P}_2$. When necessary, we use the fact that $u = \mathsf{seq}(\epsilon: u \mid \top)$ to apply a rule. We assume that the environment Σ is only extended by fresh variables (*e.g.* in FA-$\mathcal{Q}$, we assume that $\alpha \notin \Sigma$). This is always possible through alpha-renaming. The FA rule only applies when f is of non-zero arity.

<center>Fig. 13. Bi-deduction rules for protocols $\mathcal{P}_1, \mathcal{P}_2$.</center>

over indices or timestamps; FA-∨, FA-∧, FA-⇒, FA-¬ and FA-† for boolean connectives; and FA-SEQ for sequences. Remark that, whenever possible, we restrict the set of terms that must be deduced by adding new conditions. *E.g.*, in the FA-ITE rule, we know that we must deduce the then branch only if the condition ψ holds. PURE-DED states that pure trace formulas, *i.e.* formulas that only talk about the trace model, are always bi-deducible. Indeed, once the trace model is fixed, they have a constant value (either always true or always false). The UNFOLD-R-I rules allow to replace, on the right-hand side, a macro $m@\mathsf{a}[\vec{i}]$ by its definition, as long as it is at a timestamp that happens. We have another unfolding rule on the right-hand side. The UNFOLD-R-SYM rules allow to replace a macro $m@\mathsf{a}[\vec{i}]$ by its definition of *both* sides of the bi-term. Remark that we do not need to check that $\mathsf{a}[\vec{i}]$ happens: indeed, if it does not happen, then by definition $m@\mathsf{a}[\vec{i}] = \mathsf{empty}$ (simultaneously on both sides of the bi-term), which is a constant, and is therefore always bi-deducible.

Finally, the only symmetric rule is the CASE rule, which allows to do a case analysis.

A. Soundness of the bi-deduction rules

Proposition 2. *The* BI-DEDUCE *rule and the rules in Figure 5 (as well as the complete version of the rules given in Figure 13 in Appendix) are sound.*

Proof sketch. We sketch the proof of soundness of our rules.

- WEAK-L. Trivial: if you can bi-deduce $v_\#$ from $\vec{u}_\#^1$, then you can do it from $\vec{u}_\#^1, \vec{u}_\#^2$.
- SEQ-DED. Assume that:
$$\tilde{\exists}\vec{\alpha}.\ \left(\begin{matrix} [\phi_1 \Rightarrow (\psi_1 \wedge u_1 = v_1)]_{\mathcal{P}_1} \tilde{\wedge} \\ [\phi_2 \Rightarrow (\psi_2 \wedge u_2 = v_2)]_{\mathcal{P}_2} \end{matrix} \right) \quad (21)$$
and that we can bi-deduce $\#(\phi_1, \phi_2)$. We must bi-deduce $\#(v_1, v_2)$ whenever $\#(\phi_1, \phi_2)$ holds.
Take a trace model $\mathbb{T}$. If $\#(\phi_1, \phi_2)$ is false in $\mathbb{T}$, we return empty. Otherwise, from Equ. (21), we know that there exists some valuation $\vec{l}$ of $\vec{\alpha}$ such that
$$(\psi_1 \wedge u_1 = v_1)_{\mathcal{P}_1}^{\mathbb{T}\{\vec{\alpha} \mapsto \vec{l}\}} \quad \text{and} \quad (\psi_2 \wedge u_2 = v_2)_{\mathcal{P}_2}^{\mathbb{T}\{\vec{\alpha} \mapsto \vec{l}\}}$$
hold. Hence we can extract from the sequence
$$\mathsf{seq}(\vec{\alpha} : \#(u_1, u_2) \mid \#(\psi_1, \psi_2))$$
an element (corresponding to $\{\vec{\alpha} \mapsto \vec{l}\}$) such that $u_1 = v_1$ for $\mathcal{P}_1$ and $u_2 = v_2$ for $\mathcal{P}_2$ (simultaneously on both sides of the bi-term we must deduce). This concludes this case.
- PURE-DED. Since ϕ is a pure trace formula, we know that ϕ is independent of the protocol used to interpret it in a trace model. Hence, for any trace model $\mathbb{T}$, $(\phi)_{\mathcal{P}_1}^{\mathbb{T}}$ and $(\phi)_{\mathcal{P}_2}^{\mathbb{T}}$ are exactly the same terms. Moreover, since ϕ is a pure trace formula, $(\phi)^{\mathbb{T}}$ is either equal to true or false (simultaneously for both $\mathcal{P}_1$ and $\mathcal{P}_2$). Hence $\#(\phi, \phi)$ is bi-deducible. Similarly, $\#(v, v)$ is bi-deducible. Consequently, we can bi-deduce $\#(v, v)$ whenever $\#(\phi, \phi)$.
- PAIR-PROJ. The soundness of this rule is straightforward, as we can extract from a pair the value of its components.

- FRAME. The rule exploits the fact that, given frame@T as input, it is possible to extract all its predecessors frame@T' for any $T' \leq T$. Remark that this holds even if T does not happens: indeed, in that case, $T' \leq T$ is false for every T', and frame@T is empty.
- SPLIT-COND. From the premises, we know that we can bi-deduce $\phi_\#$, and that we can bi-deduce $\psi_\#$ whenever $\phi_\#$ holds. Taking the conjunction of the two machines, we can thus bi-deduce:
$$\phi_\# \wedge \psi_\#$$
Moreover, we know that we can bi-deduce $v_\#$ whenever $\phi_\#$ holds. Hence, using the same machine, we can bi-deduce $v_\#$ whenever $\phi_\# \wedge \psi_\#$ holds, since (obviously) $\phi_\# \wedge \psi_\#$ holds whenever $\phi_\#$ does.
- FA-ITE. From the premises, there exist machines: $\mathcal{B}_{\psi_\#}$ bi-deducing $\psi_\#$ whenever $\phi_\#$; $\mathcal{B}_{t_\#}$ bi-deducing $t_\#$ whenever $\phi_\# \wedge \psi_\#$; and $\mathcal{B}_{t'_\#}$ bi-deducing $t'_\#$ whenever $\phi_\# \wedge \psi_\#$. Then:
$$\text{if } \mathcal{B}_{\psi_\#} \text{ then } \mathcal{B}_{t_\#} \text{ else } \mathcal{B}_{t'_\#}$$
bi-deduces
$$\text{if } \psi_\# \text{ then } t_\# \text{ else } t'_\#$$
whenever $\phi_\#$ holds.
- FA-∨-L, FA-∨-R, and FA-⇒. These rules correspond to lazy evaluation strategies for the disjunction and implication operators.
We only sketch the FA-∨-L rule (the other rules are similar). Assume that there exist machines: $\mathcal{B}_{\psi_\#^0}$ bi-deducing $\psi_\#^0$ whenever $\phi_\#$ holds; and $\mathcal{B}_{\psi_\#^1}$ bi-deducing $\psi_\#^1$ whenever $\phi_\# \wedge \neg\psi_\#^0$ holds. Then:
$$\mathcal{B}_{\psi_\#^0} \vee \mathcal{B}_{\psi_\#^1}$$
bi-deduces $\psi_\#^0 \vee \psi_\#^1$ whenever $\phi_\#$ holds.
- FA, FA$_0$, FA-$\mathcal{Q}$, FA-FIND, FA-SEQ, FA-¬, FA-†. These rules rely on the fact that if the arguments of a function can be bi-deduced, then so can its image.
- FA-∧-L, FA-∧-R. The meta-formulas to be computed in the premise and the conclusion of these rules are equivalent, hence computing the value of one amounts to computing the value of the other.
- UNFOLD-L-1, UNFOLD-L-2, UNFOLD-R-1 and UNFOLD-R-2. We only sketch UNFOLD-L-1 (the same arguments applies to the other rules). The soundness of the rule follows from the fact that $m@\mathsf{a}[\vec{i}]$ is equal to its unfolding $\mathsf{unfold}_{\mathcal{P}_1}(m, \mathsf{a}[\vec{i}])$ whenever $\mathsf{a}[\vec{i}]$ happens (remark that the two meta-terms may not be equal if $\mathsf{a}[\vec{i}]$ does not happen).
- UNFOLD-R-SYM. Assume that we can bi-deduce
$$t_h \overset{\text{def}}{=} \text{if } \#(\mathsf{unfold}_{\mathcal{P}_1}(m, \mathsf{a}[\vec{i}]), \mathsf{unfold}_{\mathcal{P}_2}(m, \mathsf{a}[\vec{i}])).$$
Then we can bi-deduce
$$t \overset{\text{def}}{=} m@\mathsf{a}[\vec{i}]$$
since (if $\mathsf{happens}(\mathsf{a}[\vec{i}])$ then t_h) is always equal to t, for both sides of the bi-terms (we show this by taking an

$$\dfrac{\overline{\tau;\;;\;\_ \rhd_{\mathcal{P}_1,\mathcal{P}_2} \mathsf{init} < \tau}\;\text{PURE-DED}}{\dfrac{\tau;\;;\;\_ \rhd_{\mathcal{P}_1,\mathcal{P}_2} \{\mathsf{empty}\mid\mathsf{init}<\tau\}}{\tau;\;;\;\_ \rhd_{\mathcal{P}_1,\mathcal{P}_2}\{\mathsf{input}@\mathsf{init}\mid\mathsf{init}<\tau\}}\;\text{FA}_0}\quad\dfrac{\Pi_a}{\tau,\vec{i};\;;\;\mathsf{frame}@\tau\rhd_{\mathcal{P}_1,\mathcal{P}_2}\{\mathsf{input}@\mathsf{a}[\vec{i}]\mid\mathsf{a}[\vec{i}]<\tau\}\;\;\forall\mathsf{a}\in\mathcal{A}\backslash\{\mathsf{init}\}}$$

UNFOLD-R-SYM · CASE

$$\dfrac{\dfrac{\dfrac{\tau,\tau_1;\;;\;\mathsf{frame}@\tau\rhd_{\mathcal{P}_1,\mathcal{P}_2}\{\mathsf{input}@\tau_1\mid\tau_1<\tau\}}{\tau;\;;\;\mathsf{frame}@\tau\rhd_{\mathcal{P}_1,\mathcal{P}_2}\mathsf{seq}(\tau_1:\mathsf{input}@\tau_1\mid\tau_1<\tau)}\;\text{FA-SEQ}}{}}{\tau;\begin{bmatrix}\mathsf{frame}@\tau\sim\\\mathsf{frame}@\tau\end{bmatrix}_{\mathcal{P}_1,\mathcal{P}_2}\vdash\begin{bmatrix}\mathsf{seq}(\tau_1:\mathsf{input}@\tau_1\mid\tau_1<\tau)\sim\\\mathsf{seq}(\tau_1:\mathsf{input}@\tau_1\mid\tau_1<\tau)\end{bmatrix}_{\mathcal{P}_1,\mathcal{P}_2}}\;\text{BI-DEDUCE}$$

where, for every $a \in \mathcal{A}\backslash\{\mathsf{init}\}$, Π_a is the derivation:

$$\dfrac{\dfrac{\dfrac{\overline{\tau,\vec{i};\;;\;\_\rhd_{\mathcal{P}_1,\mathcal{P}_2}\mathsf{a}[\vec{i}]<\tau}\;\text{PURE-DED}\quad\overline{\tau,\vec{i};\vdash\Big[\exists\tau_0.\,\mathsf{a}[\vec{i}]<\tau\to(\tau_0\le\tau\wedge\mathsf{frame}@\tau_0=\mathsf{frame}@\mathsf{pred}(\mathsf{a}[\vec{i}]))\Big]_{\mathcal{P}_1,\mathcal{P}_2}}}{\tau,\vec{i};\;;\;\mathsf{seq}(\tau_0:\mathsf{frame}@\tau_0\mid\tau_0\le\tau)\rhd_{\mathcal{P}_1,\mathcal{P}_2}\{\mathsf{frame}@\mathsf{pred}(\mathsf{a}[\vec{i}])\mid\mathsf{a}[\vec{i}]<\tau\}}\;\text{SEQ-DED}}{\dfrac{\tau,\vec{i};\;;\;\mathsf{frame}@\tau\rhd_{\mathcal{P}_1,\mathcal{P}_2}\{\mathsf{frame}@\mathsf{pred}(\mathsf{a}[\vec{i}])\mid\mathsf{a}[\vec{i}]<\tau\}}{\tau,\vec{i};\;;\;\mathsf{frame}@\tau\rhd_{\mathcal{P}_1,\mathcal{P}_2}\{\mathsf{att}(\mathsf{frame}@\mathsf{pred}(\mathsf{a}[\vec{i}]))\mid\mathsf{a}[\vec{i}]<\tau\}}\;\text{FRAME}}}{\tau,\vec{i};\;;\;\mathsf{frame}@\tau\rhd_{\mathcal{P}_1,\mathcal{P}_2}\{\mathsf{input}@\mathsf{a}[\vec{i}]\mid\mathsf{a}[\vec{i}]<\tau\}}\;\text{FA · UNFOLD-R-SYM}$$

Convention: let $\mathcal{A}$ be the set of action names of $\mathcal{P}_1,\mathcal{P}_2$.

Fig. 14. Bi-deduction derivation example

TRANS-$\rhd$
$$\dfrac{\Sigma;\Theta;\vec{u}_{\#}\rhd\{\vec{w}_{\#}\mid\phi_{\#}\}\qquad\Sigma;\Theta;\vec{w}_{\#}\rhd\{\vec{v}_{\#}\mid\phi_{\#}\}}{\Sigma;\Theta;\vec{u}_{\#}\rhd\{\vec{v}_{\#}\mid\phi_{\#}\}}$$

SEQ-SYM
$$\dfrac{\Sigma,\vec{\alpha};\Theta;\vec{u}_{\#}^0,\vec{u}_{\#}\rhd\vec{v}_{\#}\qquad\Sigma,\vec{\alpha};\Theta;\vec{u}_{\#}^0\rhd\phi_{\#}}{\Sigma;\Theta;\vec{u}_{\#}^0,\mathsf{seq}(\vec{\alpha}:\vec{u}_{\#}\mid\phi_{\#})\rhd\mathsf{seq}(\vec{\alpha}:\vec{v}_{\#}\mid\phi_{\#})}$$

SIMPL
$$\dfrac{\Sigma;\Theta;\vec{u}_{\#}\rhd\{\vec{v}_{\#}\mid\phi_{\#}\}\qquad\mathsf{fv}(\vec{u}_{\#},\vec{v}_{\#},\phi_{\#})\cap\vec{\alpha}=\emptyset}{\Sigma,\vec{\alpha};\Theta;\vec{u}_{\#}\rhd\{\vec{v}_{\#}\mid\phi_{\#}\}}$$

Fig. 15. Additional bi-deduction rules for protocols $\mathcal{P}_1,\mathcal{P}_2$.

arbitrary trace model $\mathbb{T}$ and doing a case analysis over $\mathsf{happens}(\mathsf{a}[\vec{i}])$), and since $\mathsf{happens}(\mathsf{a}[\vec{i}])$ is a pure trace formula, hence is always bi-deducible.

- CASE. This is a straightforward case analysis. $\qquad\square$

B. Bi-deduction example

We give in Figure 14 an example of a proof by bi-deduction, for the following judgement:

$$\tau;\begin{bmatrix}\mathsf{frame}@\tau\sim\\\mathsf{frame}@\tau\end{bmatrix}_{\mathcal{P}_1,\mathcal{P}_2}\vdash\begin{bmatrix}\mathsf{seq}(\tau_1:\mathsf{input}@\tau_1\mid\tau_1<\tau)\sim\\\mathsf{seq}(\tau_1:\mathsf{input}@\tau_1\mid\tau_1<\tau)\end{bmatrix}_{\mathcal{P}_1,\mathcal{P}_2}$$

APPENDIX C
AUTOMATIC INFERENCE OF BI-DEDUCTION INVARIANTS

Assume we want to prove the bi-deduction judgement $\Sigma;\Theta;u_{\#}\rhd v_{\#}$. We give a high-level description of how we infer the invariant $U_{\#}$ of Section V-E.

Consider a fixed set of invariant candidates $\mathcal{I}$ of the form $U_{\#}^i\stackrel{\text{def}}{=}\lambda\tau.u_{\#}^i[\tau]$, where $U_{\#}^i$ free variables are bound by Σ. Our

high-level description is agnostic in the precise set of invariant candidates $\mathcal{I}$, as long as it supports some operations (discussed below). Still, to help the reader, we give an example of such a set which can be used to instantiate our algorithm.

Example 18. *A state slice is an element of the form:*

$$\lambda\tau.\,\mathsf{seq}(\vec{j}:s[\vec{i}]@\tau)$$

where $\vec{i}\subseteq\vec{j}\cup\Sigma$. Then, a state invariant candidate is any finite union of state slices, and we let $\mathcal{I}_S$ be the set of state invariant candidates.

Our inductive invariant procedure computes a sequence of elements $(U_{\#}^i)_{i\le N}\in\mathcal{I}^{N+1}$ such that:

(A) for all i, $U_{\#}^{i+1}[\tau]$ can be deduced from the initial knowledge and the known terms at previous steps:

$$\Sigma,\tau;\Theta;u_{\#},\mathsf{seq}(\tau':U_{\#}^i[\tau']\mid\tau'<\tau)\rhd U_{\#}^{i+1}[\tau]$$

(B) the candidate invariants are monotonically decreasing:

$$\Sigma,\tau;\Theta;U_{\#}^i[\tau]\rhd U_{\#}^{i+1}[\tau]\qquad\text{(for all }i\text{)}$$

Our procedure stops as soon as it computes a candidate invariant $U_{\#}^N$ which is smaller than $U_{\#}^{N-1}$ w.r.t. $\rhd$. More precisely:

(C) the procedure stops at step $N>0$ if it reached a post-fixpoint:

$$\Sigma,\tau;\Theta;u_{\#},U_{\#}^N[\tau]\rhd U_{\#}^{N-1}[\tau]$$

We can prove that if these conditions hold, then $U_{\#}^N$ is an inductive invariant.

Proposition 4. *If a sequence $(U_{\#}^i)_{i\le N}\in\mathcal{I}^{N+1}$ satisfies conditions (A) to (C), then $U_{\#}^N$ is an invariant of the protocol:*

$$\Sigma,\tau;\Theta;u_{\#},\mathsf{seq}(\tau':U_{\#}^N[\tau']\mid\tau'<\tau)\rhd U_{\#}^N[\tau]\qquad(\star)$$

Bi-deduction derivation used in Proposition 4:

$$
\dfrac{
\dfrac{
\dfrac{
\dfrac{\rule{2.2cm}{0.4pt}}{\Sigma, \tau'; \Theta; u_\#, U_\#^N[\tau'] \triangleright U_\#^{N-1}[\tau']}\ (C)
}{\Sigma, \tau, \tau'; \Theta; u_\#, U_\#^N[\tau'] \triangleright U_\#^{N-1}[\tau']}\ \text{SIMPL}
\quad
\dfrac{\rule{2.2cm}{0.4pt}}{\Sigma, \tau, \tau'; \Theta; \ldots \triangleright \tau' < \tau}\ \text{PURE-DED}
}{\Sigma, \tau; \Theta; u_\#, \mathsf{seq}(\tau' : U_\#^N[\tau'] \mid \tau' < \tau) \triangleright \mathsf{seq}(\tau' : U_\#^{N-1}[\tau'] \mid \tau' < \tau)}\ \text{SEQ-SYM}
}{\rule{0pt}{0pt}}
$$

$$
\dfrac{
\begin{array}{c}
\vdots \ \text{(remove } u_\# \text{ on the rhs)}\\[4pt]
\Sigma, \tau; \Theta; u_\#, \mathsf{seq}(\tau' : U_\#^N[\tau'] \mid \tau' < \tau) \triangleright u_\#, \mathsf{seq}(\tau' : U_\#^{N-1}[\tau'] \mid \tau' < \tau)
\end{array}
\quad
\dfrac{\rule{4cm}{0.4pt}}{\Sigma, \tau; \Theta; u_\#, \mathsf{seq}(\tau' : U_\#^{N-1}[\tau'] \mid \tau' < \tau) \triangleright U_\#^N[\tau]}\ (A)
}{\Sigma, \tau; \Theta; u_\#, \mathsf{seq}(\tau' : U_\#^N[\tau'] \mid \tau' < \tau) \triangleright U_\#^N[\tau]}\ \text{TRANS-}\triangleright
$$

Bi-deduction derivation used in Lemma 1:

$$
\dfrac{
\dfrac{\rule{4cm}{0.4pt}}{\Sigma, \tau; \Theta; u_\#, \mathsf{seq}(\tau' : U_\#^N[\tau'] \mid \tau' < \tau) \triangleright U_\#^N[\tau]}\ (\star)
}{\Sigma; \Theta; u_\# \triangleright \mathsf{seq}(\tau : U_\#^N[\tau])}\ \text{IND-}\triangleright
$$

$$
\dfrac{
\begin{array}{c}
\vdots \ \text{(remove } u_\# \text{ on the rhs)}\\[4pt]
\Sigma; \Theta; u_\# \triangleright u_\#, \mathsf{seq}(\tau : U_\#^N[\tau])
\end{array}
\quad
\dfrac{\rule{4cm}{0.4pt}}{\Sigma; \Theta; u_\#, \mathsf{seq}(\tau : U_\#^N[\tau]) \triangleright v_\#}\ (\diamond)
}{\Sigma; \Theta; u_\# \triangleright \{v_\# \mid \phi_\#\}}\ \text{TRANS-}\triangleright
$$

Fig. 16. Bi-deduction derivations used in the proof of Proposition 4 and Lemma 1.

Proof. In this proof, we use some additional rules for bi-deduction, which are given in Figure 15. Their soundness is straightforward to show, we omit the details.

The proof of Equ. $(\star)$ relies on conditions (A) and (C), and is given in Figure 16. $\qquad\square$

We now put everything together.

Lemma 1. *If a sequence* $(U_\#^i)_{i \le N} \in \mathcal{I}^{N+1}$ *satisfies conditions (A) to (C), and if:*

$$\Sigma; \Theta; u_\#, \mathsf{seq}(\tau : U_\#[\tau]) \triangleright v_\# \qquad (\diamond)$$

then $\Sigma; \Theta; u_\# \triangleright v_\#$.

Proof. The proof relies on Equ. $(\star)$ of Proposition 4, and is given in Figure 16. $\qquad\square$

Instantiating $\mathcal{I}$. To instantiate this procedure with some set of candidate invariants $\mathcal{I}$, we require the following.

- An effective procedure that can compute, for any initial knowledge $u_\#$ and for any $U_\#$ in $\mathcal{I}$, a candidate invariant $V_\# \in \mathcal{I}$ such that

$$\Sigma, \tau; \Theta; u_\#, \mathsf{seq}(\tau' : U_\#[\tau'] \mid \tau' < \tau) \triangleright V_\#[\tau]$$

  and $\Sigma, \tau; \Theta; U_\#[\tau] \triangleright V_\#[\tau]$. This procedure allows to compute the decreasing sequence of candidate invariants $(U_\#^i)_{i \le N}$ satisfying conditions (A) and (B).
- A sound (but not necessarily complete) procedure that can check, for any $U_\#, V_\#$ in $\mathcal{I}$, if:

$$\Sigma, \tau; \Theta; U_\#[\tau] \triangleright V_\#[\tau]$$

  This is used to verify the termination condition (C).
- To ensure termination, any sequence $(U_\#^i)_i \in \mathcal{I}^\infty$ such that conditions (A) and (B) holds must be stationnary.[1]

We have designed and implemented such procedures for the set $\mathcal{I}_S$ of state invariants introduced in Example 18. We do not detail them here.

---

[1] Sets of invariant candidates $\mathcal{I}$ which do not have this property could still be used, by adapting the procedure we presented so far with a widening operator [25].