

Towards Unconditional Soundness: Computationally Complete Symbolic Attacker

Gergei Bana¹ and Hubert Comon-Lundh²

¹ NTT Communication Science Laboratories, Atsugi, Kanagawa, Japan,
gergei.bana@lab.ntt.co.jp

² CNRS, INRIA project SecSi and LSV, ENS Cachan, France,
comon@lsv.ens-cachan.fr

Abstract. We consider the question of the adequacy of symbolic models versus computational models for the verification of security protocols. We neither try to include properties in the symbolic model that reflect the properties of the computational primitives nor add computational requirements that enforce the soundness of the symbolic model. We propose in this paper a different approach: everything is possible in the symbolic model, unless it contradicts a computational assumption. In this way, we obtain unconditional soundness almost by construction. And we do not need to assume the absence of dynamic corruption or the absence of key-cycles, which are examples of hypotheses that are always used in related works. We set the basic framework, for arbitrary cryptographic primitives and arbitrary protocols, however for trace security properties only.

The main points of this work have been published in the proceedings of POST'12 [8]. This paper largely agrees with that publication, but in Section 2.6 we present an improved computational semantics. This improvement allows us to state our main theorem for any first order formula, while in the POST publication, the possible formulas were somewhat restricted. Moreover, we introduced here Section 4, which explains the connection between our results and an earlier result of Fitting about embedding (first-order) classical logic into (first-order) S4.

1 Introduction

The automatic analysis of security protocols has been quite successful since 1990, yielding several tools [11, 18, 25]. However, when the outcome of one of these provers is “the protocol is secure”, it must be understood as “secure in our model”. Nothing guarantees that the necessary abstractions are relevant to actual implementations. For instance, consider the Needham-Schroder-Lowe protocol [22]. It has been proved secure by all the above-mentioned provers. However, there are several attacks, for instance when the encryption scheme does not guarantee the ciphertext integrity [28] or when the pairing is associative [23] or when some random number could be confused with some pairings [9].

For this reason, it is important to investigate what exactly the assumptions are, on the cryptographic primitives' implementations, that guarantee the faithfulness of the abstraction. (It is called *soundness* in the literature).

There are a lot of works providing some soundness results, typically the works initiated by Backes et al [5, 3, 6] and Abadi et al [1, 16, 14]. They essentially prove that a given symbolic model is fully abstract with respect to a computational one, assuming some properties of the security primitives. This guarantees that the security proofs that have been completed in the abstract model are also valid in a computational model.

However, these works require a very large set of assumptions, that are not always emphasized. For instance in [7] the complete list of assumptions for public-keys is listed; it is a long list of strong hypotheses, that are not fulfilled by most actual protocols. [14] make even less realistic assumptions, in order to get a stronger soundness result (which includes more security properties). All these results typically assume that no key cycles can ever be created, that bitstrings can be parsed in deterministic polynomial time into terms, that there is no dynamic corruption, that keys are certified, etc. These assumptions, as well as reasons why they are not realistic enough is discussed in [15]. Furthermore, each primitive requires a new soundness proof and each combination of primitives also requires a new soundness proof, unless much stronger properties are assumed [13]. Currently, it seems more realistic to use CRYPTOVERIF [12], completing the proofs directly in the computational model, than using a soundness result [2]. Is it really impossible to avoid manipulating computation times, probabilities, bitstring lengths... ?

In this paper, we advocate a new way of performing proofs in a symbolic, abstract, model, while keeping strong, computational guarantees establishing a general soundness result, but without establishing many specific soundness results for specific properties of primitives. Such properties can later be proven and added.

The idea is to design a symbolic setting, in which any adversarial action is possible, unless it contradicts some axiom expressing a property that must be satisfied under standard computational assumptions. In other words, computational properties, such as IND-CCA, can be (symbolically) axiomatized and added to the system in order to limit the possible adversarial moves. We do not require the axiomatization to be complete. The idea is to only list properties that we know for certain about the implementation, and allow any symbolic move consistent with those properties. In this way, either we find an attack, in which case there is at least one possible set of primitives satisfying the assumed properties and for which the security goal is violated, or the axioms were sufficient to ensure the security of the protocol, in which case any implementation fulfilling these axioms will ensure the security.

This approach has several advantages:

1. Though the proofs are performed in a symbolic setting, they are computationally valid.
2. Thanks to our result (Theorem 2), adding a new cryptographic primitive only requires to design an axiomatization of this primitive and prove it sound due to the cryptographic assumptions: the additional soundness proof is short and modular; it focuses on designated properties instead of considering whole execution models.
3. We may be able prove the security of protocols with weaker assumptions on the primitives. For instance, if we prove the security using only axioms that are sound for IND-CPA encryption, then IND-CPA will be a sufficient hypothesis for security.

4. In each security proof, all assumptions are clearly and formally stated as axioms.
5. In case an attack is found, it may be sufficient to add an axiom (expressing stronger hypotheses on the computational implementation of the primitives) ruling out the attack, then try proving again.
6. We may consider any cryptographic primitive, including XOR for instance (for which there are strong limitations of the computational soundness approach [4, 27]). Dynamic corruption, key cycles, etc. are not a priori discarded.

Related works. The most closely related works are probably those that consider a proof system that is sound w.r.t. the computational semantics, such as [17, 9]. Though these works are related, as far as the computational semantics of the logic is concerned, the overall strategy is completely different. We do not try to design a proof system working directly in the computational model: we only use first-order logic and standard inference rules in the symbolic model. Our approach is more inspired by circumscription [21], however circumscribing what is *not* possible. In other words, we do not design inference rules, we modify the transition system instead. This is similar in spirit to [26], in which any property of the hash function, that is not explicitly forbidden by some axiom, is considered as valid.

Contents of the paper. In this paper, we only state the framework of the method, prove a general soundness theorem in the case of trace properties, and prove soundness of an example axiom expressing secrecy of an IND-CCA encryption.

More precisely, protocols are identified to a formal transition system in the same spirit as CoSP [7]: we do not commit to a very particular way of specifying such a transition system. The possible transitions are, roughly, defined by a formula, that guards the transition by constraining the input message, a state move and a message that is sent out when the guard is satisfied. Such transitions can be interpreted in different models: symbolic models, in which the messages are terms and the guards are interpreted in a Herbrand model, or computational models, in which messages are bitstrings. In the symbolic models, we constrain the input messages to be *deducible* from the previous outputs and the public information. Such a deducibility condition is formalised using a deducibility predicate, whose interpretation is not fixed. This is a main difference with classical protocol verification: the attacker capabilities are not fixed, but rather they parametrize the model. Actually, we consider any attacker capability, that does not contradict the (computationally sound) axioms. On the computational side, the attacker is any probabilistic polynomial time Turing machine: the deduction capabilities are given by any such machine. These models are explained in the sections 2.2, 2.3, 2.5.

Next, we need to specify the axioms and the (trace) security properties. We consider any first-order formula, that is built on the predicate symbols, that are used in the guards, as well as the deducibility predicate symbol. We need such general formulas, since we need to constrain the symbolic models of the deducibility relations, i.e., the symbolic attacker capabilities, according to the computational assumptions on the primitives. Typically, we may consider an axiom of the form: “if a plaintext can be deduced (resp. computed) from a ciphertext and a set of messages ϕ , then the decryption key has been sent out or else the plaintext can be deduced (resp. computed) from ϕ ”, that reflects

some property of the encryption scheme. The meanings of these axioms/security properties become clear when we define a computational interpretation of such formulas, which we provide in the section 2.6.

The Section 3 is devoted to the main result, which states a general trace-mapping soundness property: independently of the primitives and their specific characteristics, if there is a computational attack, then there is a symbolic attack. Once more, the symbolic attacker has any capability, that is consistent with the axioms. So, this result, though subtle and not at all trivial to prove, is not surprising. The whole system was actually carefully designed with this aim in mind.

We also show in the Section 5 some axiom examples, that are proven sound under some standard cryptographic properties. We do not aim however at covering a large set of axioms. Further axioms will be added to a library each time they are required for the proof of a case study.

This paper aims at opening a new research direction: it seems very appealing and promising. We need however to investigate several case studies. As a “proof of concept”, we have designed a complete set of axioms (included in the Appendix) and proved the NSL protocol in our framework (available from the first-author’s web page or upon request). This sufficient set of axioms shows also that some hypotheses of earlier works are not necessary (at least for weak secrecy and authentication).

2 Symbolic and Computational models

2.1 Terms and frames

Terms are built out of a set of function symbols \mathcal{F} that contains an unbounded set of names \mathcal{N} and an unbounded set of handles \mathcal{H} . Let \mathcal{X} be an unbounded set of variables. Names and handles are zero arity function symbols. We will use names to denote items honestly generated by agents, while handles will denote inputs of the adversary. A ground term is a term without variables. *Frames* are sequences of terms together with name binders: a frame ϕ can be written $(\nu \bar{n}). p_1 \mapsto t_1, \dots, p_n \mapsto t_n$ where p_1, \dots, p_n are place holders that do not occur in t_1, \dots, t_n and \bar{n} is a sequence of names. $\text{fn}(\phi)$, the *free names* of ϕ are names occurring in some t_i and not in \bar{n} . The *variables* of ϕ are the variables of t_1, \dots, t_n .

Example 1. We typically use a randomized public-key encryption symbol: $\{m\}_{eKQ}^r$ is intended to represent the encryption of the plaintext m with the public-key of the principal Q , with a random seed r . More generally, we consider the example when there is a set of constructors $\mathcal{F}_c = \{\{-\}_-, \langle -, - \rangle, e_-, d_-, K_-\}$, and a set of destructors $\mathcal{F}_d = \{\text{dec}(-, -), \pi_1(-), \pi_2(-)\}$, and $\mathcal{F} = \mathcal{F}_c \cup \mathcal{F}_d \cup \mathcal{N} \cup \mathcal{H}$.

2.2 Formulas

Let \mathcal{P} be a set of predicate symbols over terms. \mathcal{P} is assumed to contain the equality $=$ (which is interpreted as a congruence), used as $t_1 = t_2$, and a predicate \vdash , which takes as arguments an n -tuple of terms on its left and a term on its right (and which is

intended to model the computation capabilities), that is, written as $t_1, \dots, t_n \vdash t$. (More precisely, it is an infinite sequence of predicates, with arguments $n + 1$.)

We are not interested in any specific symbolic interpretation of these predicate symbols. We wish to consider *any* possible symbolic interpretation, that satisfies some requirements; the aim is to allow anything that is not forbidden by explicit assumptions.

Example 2. $\forall x, \forall y. (\{x\}_{eK_Q}^s = \{y\}_{eK_Q}^{s'} \rightarrow x = y)$ is such a formula, the validity of which follows from the uniqueness of decryption.

Let \mathcal{M} denote then any first-order structure that interprets the function and predicate symbols of the logic. We only assume that $=$ is interpreted in \mathcal{M} as the equality in the underlying domain $D_{\mathcal{M}}$. The relation in \mathcal{M} (that is, a relation for elements in $D_{\mathcal{M}}$), interpreting the deducibility predicate \vdash is denoted as $\vdash_{\mathcal{M}}$.

Given an assignment σ of elements in $D_{\mathcal{M}}$ to the free variables of term t , we write $\llbracket t \rrbracket_{\mathcal{M}}^{\sigma}$ for the interpretation of t in \mathcal{M} ($\llbracket _ \rrbracket_{\mathcal{M}}^{\sigma}$ is the unique extension of σ into a homomorphism of \mathcal{F} -algebras).

For any first order structure \mathcal{M} over the functions \mathcal{F} and predicates \mathcal{P} , given a first order formula θ and an assignment σ of elements in the domain of \mathcal{M} to the free variables of θ , the satisfaction relation $\mathcal{M}, \sigma \models \theta$ is defined as usual in first-order logic.

Example 3. Consider the public-key encryption setting of example 1. We may use unary predicate symbols to restrict sets of data. Assume for instance that W is supposed to represent the set of agent names, and M is supposed to represent well formed terms (that are equal to a term built with symbols in \mathcal{F}_c).

$$W(\pi_1(\text{dec}(h, db))) \wedge M(\pi_2(\text{dec}(h, db)))$$

is a formula, that expresses that the handle h can be decrypted and projected into two components, one of which is an agent name.

2.3 Protocols

We do not stick to any particular syntax for the definition of protocols. We only assume that it defines a transition system as follows. Q is a set of *control states*, together with a finite set of free variables.

Definition 1. A protocol is a recursive (actually *PTIME*) set of tuples

$$(q(\bar{n}), q'(\bar{n} \cdot \bar{n}'), \langle x_1, \dots, x_k \rangle, x, \psi, s)$$

where $q, q' \in Q$, x_1, \dots, x_k, x are variables \bar{n}, \bar{n}' are finite sequences of names, ψ is a first order formula over the set of predicate symbols \mathcal{P} and function symbols \mathcal{F} and the names $\bar{n} \cup \bar{n}'$, whose free variables are in $\{x_1, \dots, x_n, x\}$ and s is a term whose free variables are in $\{x_1, \dots, x_n, x\}$.

For example, ψ can be a formula such as $\text{dec}(x, k) = n$, that checks that the current input is a ciphertext whose plaintext is a previously generated nonce n : ψ guards the transition. s is the output message, when the transition succeeds. The intended meaning

of these rules is that a transition from the state q to the state q' is possible, given the previous inputs x_1, \dots, x_n and the new input x , if the formula ψ is satisfied. In such a case, the names \bar{n}' are generated and the message s is sent.

Such a formalism is quite general; we only assume here (for simplicity) a single, public, communication channel. Typically, applied π -calculus processes can be translated into such transition rules, that are similar to the CoSP framework of [7].

Example 4. Consider a single session of the NSL protocol. The states consist of pairs of the local states of each of the processes for A and B. Instead of listing the transitions as tuples, we write $\psi : q(\bar{n}) \xrightarrow{s} q'(\bar{n}')$ and they are displayed in the figure 1. In this version of the protocol, the responder is willing to communicate with anybody, hence only checks $W(\pi_1(dec(y, dK_B)))$; the intended meaning of W is a set of agent names. If

$$\begin{array}{l}
\top : \\
\left. \begin{array}{l}
\pi_1(dec(x, dK_A)) = B \\
\wedge \pi_1(\pi_2(dec(x, dK_A))) = n
\end{array} \right\} : \begin{array}{l}
q_0^A(n, r, r'') \\
q_1^A(n, r, r'')
\end{array} \xrightarrow[\{\pi_2(\pi_2(dec(x, dK_A)))\}_{eK_B}^{r''}]{\{\langle A, n \rangle\}_{eK_B}^r} \begin{array}{l}
q_1^A(n, r, r'') \\
q_2^A(n, r, r'')
\end{array} \\
\left. \begin{array}{l}
W(\pi_1(dec(y, dK_B))) \\
\wedge M(\pi_2(dec(y, dK_B))) \\
dec(z, dK_B) = n'
\end{array} \right\} : \begin{array}{l}
q_0^B(n', r') \\
q_1^B(n', r')
\end{array} \xrightarrow[\longrightarrow]{\{\langle B, (\pi_2(dec(y, dK_B)), n') \rangle\}_{eK_{\pi_1(dec(y, dK_B))}}^{r'}} \begin{array}{l}
q_1^B(n', r') \\
q_2^B(n', r')
\end{array}
\end{array}$$

Fig. 1. The 3 transitions of 1 session of NSL

we wish to describe an unbounded number of sessions, we need to record in the control state the states of every (opened) A -session and (opened) B -session. This yields an infinite, yet recursive, set of transition rules.

Definition 2. A symbolic state of the network consists of:

- a control state $q \in Q$ together with a sequence of names (that have been generated so far) n_1, \dots, n_k
- a sequence constants called handles h_1, \dots, h_n (recording the attacker's inputs)
- a ground frame ϕ (the agents outputs)
- a set of formulas Θ (the conditions that have to be satisfied in order to reach the state).

A symbolic transition sequence of a protocol Π is a sequence

$$(q_0(\bar{n}_0), \emptyset, \phi_0, \emptyset) \rightarrow \dots \rightarrow (q_m(\bar{n}_m), \langle h_1, \dots, h_m \rangle, \phi_m, \Theta_m)$$

if, for every $m - 1 \geq i \geq 0$, there is a transition rule

$$(q_i(\bar{\alpha}_i), q_{i+1}(\bar{\alpha}_{i+1}), \langle x_1, \dots, x_i \rangle, x, \psi, s)$$

such that $\bar{n} = \overline{\alpha_{i+1}} \setminus \overline{\alpha_i}$, $\phi_{i+1} = (\nu \bar{n}).(\phi_i \cdot p \mapsto s\rho_i \sigma_{i+1})$, $\overline{n_{i+1}} = \overline{n_i} \uplus \bar{n}$, $\Theta_{i+1} = \Theta_i \cup \{\phi_i \vdash h_{i+1}, \psi \rho_i \sigma_{i+1}\}$ where $\sigma_i = \{x_1 \mapsto h_1, \dots, x_i \mapsto h_i\}$ and ρ_i is a renaming of the sequence $\overline{\alpha_i}$ into the sequence $\overline{n_i}$. We assume a renaming that ensures the freshness of the names \bar{n} : $\bar{n} \cap \overline{n_i} = \emptyset$.

Definition 3. Given an interpretation \mathcal{M} , a transition sequence of Π

$$(q_0(\overline{n_0}), \emptyset, \phi_0, \emptyset) \rightarrow \dots \rightarrow (q_m(\overline{n_m}), \langle h_1, \dots, h_m \rangle, \phi_m, \Theta_m)$$

is valid w.r.t. \mathcal{M} if, for every $m - 1 \geq i \geq 0$,

$$\mathcal{M} \models \Theta_{i+1}$$

Example 5. We show the beginning of a possible branch in the symbolic execution of NSL.

$$(q_0, \emptyset, \phi_0, \emptyset) \quad (q_1, H_1, \phi_1, \Theta_1) \quad (q_2, H_2, \phi_2, \Theta_2) \quad (q_3, H_3, \phi_3, \Theta_3) \quad (q_4, H_4, \phi_4, \Theta_4)$$

- $\phi_0 = \nu_{K_A K_B AB}(p_0 \mapsto (A, B, eK_A, eK_B))$,
 $\Theta_0 = \emptyset$
- $H_1 = \langle h_1 \rangle$,
 ϕ_1 extends ϕ_0 with $p_1 \mapsto \{\langle A, n \rangle\}_{eK_B}^r$,
 $\Theta_1 = \{\phi_0 \vdash h_1\}$
- $H_2 = \langle h_1, h_2 \rangle$,
 ϕ_2 extends ϕ_1 with $p_2 \mapsto \{\langle B, \langle \pi_2(\text{dec}(h_2, dK_B)), n' \rangle \rangle\}_{eK_{\pi_1(\text{dec}(h_2, dK_B))}}^{r'}$,
 $\Theta_2 = \Theta_1 \cup \{\phi_1 \vdash h_2, M(\pi_2(\text{dec}(h_2, dK_B))), W(\pi_1(\text{dec}(h_2, dK_B)))\}$
- $H_3 = \langle h_1, h_2, h_3 \rangle$,
 ϕ_3 extends ϕ_2 with $p_3 \mapsto \{\pi_2(\pi_2(\text{dec}(h_3, dK_A)))\}_{eK_B}^{r''}$,
 $\Theta_3 = \Theta_2 \cup \{\phi_2 \vdash h_3, \pi_1(\pi_2(\text{dec}(h_3, dK_A))) = n, \pi_1(\text{dec}(h_3, dK_A)) = B\}$,
- $H_4 = \langle h_1, h_2, h_3, h_4 \rangle$, $\phi_4 = \phi_3$,
 $\Theta_4 = \Theta_3 \cup \{\phi_3 \vdash h_4, \text{dec}(h_4, dK_B) = n'\}$

Let \mathcal{M} be a model in which $\pi_1(\text{dec}(h_2, dK_B)) = A$ and

$$h_2 =_{\mathcal{M}} \{\langle A, n \rangle\}_{eK_B}^r, \quad h_3 =_{\mathcal{M}} \{\langle B, \langle n, n' \rangle \rangle\}_{eK_A}^{r'}, \quad h_4 =_{\mathcal{M}} \{n'\}_{eK_B}^{r''}$$

and $\vdash_{\mathcal{M}}$ is simply the classical Dolev-Yao deduction relation. Then the execution sequence is valid w.r.t. \mathcal{M} , and this corresponds to the correct execution of the NSL protocol between A and B .

There are however other models in which this transition sequence is valid. For instance let \mathcal{M}' be such that $h_2 =_{\mathcal{M}'} n$ and $\phi_1 \vdash_{\mathcal{M}'} n$ and $n =_{\mathcal{M}'} \{\langle A, n \rangle\}_{eK_B}^r$, (and h_3, h_4 as above). We get again a valid transition sequence w.r.t. \mathcal{M}' . Though, in what follows, we will discard such sequences, thanks to some axioms.

Example 6. Consider again the transitions of the example 5. Now consider a model \mathcal{M} in which $n_0, \{B, n, n'\}_{e_{K_A}}^r \vdash_{\mathcal{M}} \{B, n_0, n'\}_{e_{K_A}}^r$ for an honestly generated nonce n_0 that can be chosen by the attacker: the transition sequence of the previous example is also valid w.r.t. this model. This will yield an attack, using a malleability property of the encryption scheme, as in [28]. Discarding such attacks requires some properties of the encryption scheme (for instance IND-CCA). It can be ruled out by a non-malleability axiom like the one appearing in the Appendix.

From these examples, we see that unexpected attacks can be found when some assumption is not explicitly stated as an axiom to limit adversarial capabilities.

2.4 Axioms and security properties

For simplicity, we only consider reachability security properties. The extension to any trace property should not be very difficult: it suffices to record some values along the trace. Security properties (and, later, axioms) are first-order formulas that may contain state-dependent predicates and/or predicates that get fixed interpretation. As in the previous sections, \mathcal{M} is an arbitrary first-order structure and σ is an assignment of the free variables to elements of $D_{\mathcal{M}}$.

First, we add atomic formulas $\hat{\phi}, s_1, \dots, s_n \vdash t$, where $\hat{\phi}$ is just part of the syntax of this predicate (not an input of the predicate), which aims at ranging over frames (when interpreting the predicate) and is evaluated in every state. For t_1, \dots, t_m closed terms, $\mathcal{M}, \sigma, \langle t_1, \dots, t_m \rangle, \bar{n} \models \hat{\phi}, s_1, \dots, s_n \vdash t$ iff $\mathcal{M}, \sigma \models s_1, \dots, s_n, t_1, \dots, t_m \vdash t$

In addition, we consider the following atomic formulas, whose evaluation only depends on the state, independently of the first-order structure \mathcal{M} .

- $\text{RandGen}(s)$ (s is a ground term) expresses that s has been randomly generated:
 $\mathcal{M}, \sigma, \langle t_1, \dots, t_m \rangle, (n_1, \dots, n_k) \models \text{RandGen}(s)$ iff $s \in \{n_1, \dots, n_k\}$
- $t \sqsubseteq \hat{\phi}$ (t is a ground term) expresses that t is a subterm of the messages sent so far:
 $\mathcal{M}, \sigma, \langle t_1, \dots, t_m \rangle, \bar{n} \models t \sqsubseteq \hat{\phi}$ iff t is a subterm of some t_i .
- We also may use the derived predicate (as an abbreviation):

$$\text{fresh}(x, \hat{\phi}) = \text{RandGen}(x) \wedge x \not\sqsubseteq \hat{\phi}$$

\sqsubseteq and $\text{RandGen}()$ are *interpreted predicates* since their interpretation does not depend on \mathcal{M} . Bound variables that appear within an interpreted predicate are called *constrained variables*. As in other works on constrained logics (see for instance [20]), such variables are used to schematize several first-order formulas and are replaced with ground terms built on \mathcal{F} . Therefore, the interpretation of axioms and security properties that may involve interpreted predicates, is modified, only in case of a quantification on a constrained variable x , in which case x is replaced by any (or some, for existential quantification) ground term:

If x is a constrained variable (that is, θ has an interpreted predicate and x appears in it), then,

$$\mathcal{M}, \sigma, \langle t_1, \dots, t_n \rangle, (n_1, \dots, n_k) \models \forall x. \theta$$

iff, for every ground term t ,

$$\mathcal{M}, \sigma, \langle t_1, \dots, t_n \rangle, (n_1, \dots, n_k) \models \theta\{x \mapsto t\}$$

We have a similar definition for existential quantifications of such variables. All other cases follow the classical definition of the first-order satisfaction relation.³ This yields a satisfaction relation $\mathcal{M}, \sigma, \langle t_1, \dots, t_m \rangle, \bar{n} \models \theta$, and thus of $\mathcal{M}, \sigma, \phi, \bar{n} \models \theta$ with ϕ having the terms $\langle t_1, \dots, t_m \rangle$. When θ has no free variable, we may omit σ . Similarly, if θ does not contain atomic formulas that depend on ϕ (resp. \bar{n}), we may omit these components: we get back to the satisfaction relation of section 2.2.

We define now the satisfaction relation in a state:

$$\mathcal{M}, (q, \langle h_1, \dots, h_m \rangle, \bar{n}, \phi_m, \Theta) \models \theta \quad \text{iff} \quad \mathcal{M}, \phi_m, \bar{n} \models \theta.$$

Definition 4. A symbolic interpretation and a protocol satisfy the security property θ , written as

$$\mathcal{M}, \Pi \models \theta,$$

if for any sequence of transitions that is executable in \mathcal{M} and that yields the state $(q_m, \langle h_1, \dots, h_m \rangle, \bar{n}_m, \phi_m, \Theta_m)$,

$$\mathcal{M}, (q_m, \langle h_1, \dots, h_m \rangle, \bar{n}_m, \phi_m, \Theta_m) \models \theta.$$

Example 7. Concerning security properties, consider the NSL protocol. We may state the confidentiality of n :

$$\neg \hat{\phi} \vdash n$$

Consider now an authenticity property. We modify slightly the states of the transition system, including a commitment on the nonce on which the parties are supposed to agree. We let c_i be a special function symbol, that takes as arguments A, B, n_1, n_2 : who commits, for who and the corresponding nonces. $c_i(A, B, n_1, \pi_2(\pi_2(\text{dec}(x, dK_B))))$ is sent at the end by the initiator. For the responder, there is a similar commitment: at the end of the protocol, B emits $c_r(\pi_1(\text{dec}(x, dK_B)), B, \pi_2(\text{dec}(y, dK_B), n_2))$. We state as axioms that c_i, c_r cannot help the attacker and cannot be forged. For instance: $\forall x, y, z, w. \hat{\phi}, c_i(x, y, z, w) \not\vdash z, w$ and $\forall x, y, z, w. \hat{\phi} \vdash c_i(x, y, z, w) \rightarrow c_i(x, y, z, w) \sqsubseteq \hat{\phi}$. The agreement property (on n) may then be stated (for instance) as:

$$\forall x, y, z, w. c_r(x, y, z, w) \sqsubseteq \hat{\phi} \rightarrow \exists x' z' w' (c_i(x', y, z', w') \sqsubseteq \hat{\phi} \wedge x = x' \wedge z = z' \wedge w = w')$$

That is: x 's view of z, w is the same as y 's view of z, w .

With such a definition, for any security property and any protocol there will (almost) always be an interpretation for which the property is violated. Hence we restrict the class of symbolic interpretations, ruling out the interpretation whose all computational

³ It would in fact be possible to avoid the notion of constrained variables if we defined $D_{\mathcal{M}}$ to be a free \mathcal{F} -algebra, and $=$ a congruence relation on it (as opposed to the equality of $D_{\mathcal{M}}$), and later parts of the paper could be adjusted accordingly. However, since constrained variables are more convenient for automatic verification, the authors decided to present the theory utilizing them.

counterparts would violate some security assumption on the primitives. More precisely, we consider a set of *axioms* \mathcal{A} , which is a set of first-order formulas in the same format as the security properties. We restrict our attention to symbolic interpretations that satisfy \mathcal{A} .

Example 8. – For instance we could include in \mathcal{A} a formula

$$\text{fresh}(k, \hat{\phi}) \rightarrow \neg(\hat{\phi} \vdash k)$$

that states that an attacker cannot guess (except with negligible probability) a randomly generated name. Adding such an axiom in \mathcal{A} rules out symbolic interpretations in which this deduction is possible.

– If the computational implementation is such (e.g. they are tagged), we may include,

$$\forall x, y, z, A, r. \langle x, y \rangle \neq \{z\}_{K_A}^r$$

stating that pairs and ciphertexts cannot be confused.

We will see more examples in Section 5.

We may assume w.l.o.g that the axioms and security properties are just (universally quantified) clauses.

2.5 Computational interpretation

The computational interpretations are just a special case of interpretation of our formulas, when they do not depend on the state of the transition system. We define them again here, since we wish to introduce some additional notions. Also, the computational executions of the protocols rely on a concrete adversary, given by a Turing machine, while in general, the interpretation of functions and predicates need not to be computable.

We consider a family computational algebras, parametrized by a security parameter η , in which each function symbol is interpreted as a polynomially computable function on bitstrings (that may return an error message). Given then a sample τ of names (for every name n , its interpretation is a bitstring $\tau(n)$), every ground term t is interpreted as a bitstring $\llbracket t \rrbracket_\tau$ in such a way that $\llbracket - \rrbracket_\tau$ is a homomorphism of \mathcal{F} -algebras. More generally, if σ is an assignment of the variables of t to bitstrings $\llbracket t \rrbracket_\tau^\sigma$ is the (unique) extension of τ (on names) and σ (on variables) as a homomorphism of \mathcal{F} -algebras.

Similarly, all predicate symbols are interpreted as polynomially computable functions on bitstrings. The equality predicate is interpreted as a strict equality on bitstrings: $\tau \models^c t_1 = t_2$ if $\llbracket t_1 \rrbracket_\tau$ is not an error, $\llbracket t_2 \rrbracket_\tau$ is not an error and $\llbracket t_1 \rrbracket_\tau = \llbracket t_2 \rrbracket_\tau$.

This interpretation is extended to arbitrary closed formulas whose atomic formulas do not depend on the state. This yields the satisfaction relation $\tau \models^c \theta$. We will define later the computational interpretation of arbitrary formulas in a given state.

We now define computational executions.

Definition 5. *Given a set of transition rules, a computational state consists of*

- A symbolic state s (that is itself a tuple $q(\bar{n}, \bar{h}, \phi, \Theta)$)
- a sequence of bitstrings $\langle b_1, \dots, b_m \rangle$ (the attacker's outputs)

- A sequence $\langle b'_1, \dots, b'_n \rangle$ of bitstrings (the agents outputs)
- The configuration γ of the attacker.

Definition 6. Given a PPT interactive Turing machine \mathcal{M} and a sample τ , a sequence of transitions

$$(s_0, \emptyset, \mathbf{b}'_0, \gamma_0) \rightarrow \dots \rightarrow (s_m, \langle b_1, \dots, b_m \rangle, \langle b'_1, \dots, b'_m \rangle, \gamma_m)$$

is (computationally) valid w.r.t. \mathcal{M} and τ if

- $s_0 \rightarrow \dots \rightarrow s_m$ is a transition sequence of the protocol
- for every $i = 0, \dots, m-1$, $s_i = (q_i(\bar{n}_i), \bar{h}_i, \phi_i, \Theta_i)$, $\phi_{i+1} = (\nu \bar{n}).\phi_i \cdot u_i$, $\llbracket u_i \rrbracket_\tau = b'_{i+1}$
- for every $i = 0, \dots, m-1$, there is a configuration γ'_i of the machine \mathcal{M} such that $\gamma_i \vdash_M^* \gamma'_i \vdash_M^* \gamma_{i+1}$ and γ'_i is in a sending state, the sending tape containing b_{i+1} , γ_{i+1} is in a receiving state, the receiving tape containing b'_{i+1}
- for every $i = 0, \dots, m-1$, $\tau, \{x_1 \mapsto b_1, \dots, x \mapsto b_{i+1}\} \models^c \Theta_{i+1}$.

Intuitively, \mathbf{b}'_0 is the attacker's initial knowledge and we simply replaced symbolic deductions/symbolic models of the section 2.3 with computations/computational models.

2.6 Computational validity of security properties and axioms

We already considered the computational satisfaction of formulas, except for formulas that depend on the states. Given a PT Turing machine \mathcal{A} , we define then

$$\mathcal{A}, \tau \models^c t_1, \dots, t_n \vdash t \quad \text{iff} \quad \mathcal{A}(\llbracket t_1 \rrbracket_\tau, \dots, \llbracket t_n \rrbracket_\tau) = \llbracket t \rrbracket_\tau$$

The difficulty now is that we do not want to define $\mathcal{A}, \tau \models^c \hat{\phi} \vdash t_1 \rightarrow \hat{\phi} \vdash t_2$ as $\mathcal{A}, \tau \models \hat{\phi} \vdash t_2$ or $\mathcal{A}, \tau \not\models \hat{\phi} \vdash t_1$. In order to understand this, consider for instance the formula

$$\theta : \quad \forall_{t,K,R} (\hat{\phi}, \{t\}_{eK}^R \vdash t \rightarrow \{t\}_{eK}^R \sqsubseteq \hat{\phi} \vee dK \sqsubseteq \hat{\phi} \vee \hat{\phi} \vdash t)$$

We want (intuitively) IND-CCA encryption schemes to satisfy this formula. However, consider an instance of this axiom in which $\hat{\phi}$ is the pair $\phi = \nu n_1 n_2. \langle n_1, n_2 \rangle$, and t is n_1 . Now, let \mathcal{A} be a machine which, on input $\llbracket \langle n_1, n_2 \rangle \rrbracket_\tau$, $\llbracket \{n_1\}_{eK}^r \rrbracket_\tau$ returns n_1 and, on input $\llbracket \langle n_1, n_2 \rangle \rrbracket_\tau$ only, returns $\llbracket n_2 \rrbracket_\tau$. For every τ , $\mathcal{A}, \tau \not\models^c \theta$. Hence, whatever security is provided by the encryption scheme, there is an attack on the property.

This paradox comes from the deterministic interpretation of the deducibility relation: while, symbolically, it is a relation, it must be a function in the computational setting since we cannot consider non-deterministic machines. The intended interpretation therefore involves several machines: roughly, for any machine that can compute $\llbracket t \rrbracket_\tau$ from $\llbracket \phi \rrbracket_\tau$, $\llbracket \{t\}_{eK}^r \rrbracket_\tau$, either there is a machine that can compute $\llbracket t \rrbracket_\tau$ from $\llbracket \phi \rrbracket_\tau$ or else the actual frame contains either dK or $\{t\}_{eK}^r$. These two machines need of course to be independent of τ . This is the definition that we formalize now for arbitrary security properties.

Let \mathcal{M} be an interactive PPT Turing machine with a special challenge control state q_c . We may regard this machine as an attacker, who moves to the state q_c when (s)he thinks that (s)he is ready to break the security property.

In what follows, S is any (polynomial time) non-negligible set of interpretations of names, and S_\top is the set of all name interpretations. $\mathcal{M}, \Pi \models^c \theta$ iff $\mathcal{M}, \Pi, S_\top \models^c \theta$ and $\Pi \models^c \theta$ if $\mathcal{M}, \Pi \models^c \theta$ for every \mathcal{M} with q_c .

We introduce machines that compute witnesses for the unconstrained quantified variables.

- $\mathcal{M}, \Pi, S \models^c \exists x.\theta$ iff for any $S' \subseteq S$ non-negligible, there is a $S'' \subseteq S$ non-negligible and a PT machine \mathcal{A}_x such that $\mathcal{M}, \Pi, S'', \mathcal{A}_x \models^c \theta$
- $\mathcal{M}, \Pi, S, \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_n} \models^c \forall x.\theta$ iff for any probabilistic polynomial time machine $\mathcal{A}_x, \mathcal{M}, \Pi, S, \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_n}, \mathcal{A}_x \models^c \theta$

If x is a constrained variable, the interpretation of $\forall x.\theta$ is analogous to the symbolic case: $\mathcal{M}, \Pi, S, \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_n} \models^c \forall x.\theta$ if and only if for every ground term t , the satisfaction $\mathcal{M}, \Pi, S, \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_n} \models^c \theta\{x \mapsto t\}$ holds (and similarly for existential quantification). If σ is a sequence of machines, one for each free variable x of θ ,

- $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1 \wedge \theta_2$ iff $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1$ and $\mathcal{M}, \Pi, S, \sigma \models^c \theta_2$.
- $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1 \vee \theta_2$ iff for any $S' \subseteq S$ non-negligible, there is a $S'' \subseteq S$ non-negligible such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$.
- $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1 \rightarrow \theta_2$ iff for any $S' \subseteq S$ non-negligible, $\mathcal{M}, \Pi, S', \sigma \models^c \theta_1$ implies $\mathcal{M}, \Pi, S', \sigma \models^c \theta_2$
- $\mathcal{M}, \Pi, S, \sigma \models^c \neg\theta$ iff for any S' , if $\mathcal{M}, \Pi, S', \sigma \models^c \theta$, then $S \cap S'$ is negligible
- in the case of an atomic formula $P(t_1, \dots, t_n)$ where $P \notin \{\vdash, \sqsubseteq, \text{RandGen}()\}$, $\mathcal{M}, \Pi, S, \sigma \models^c P(t_1, \dots, t_n)$ if there is an overwhelming subset S' of S such that the following holds. For any $\tau \in S'$, consider the unique valid computation (if there is one) of Π with respect to \mathcal{M}, τ that yields a configuration of \mathcal{M} , which is in the control state q_c with a bitstring b on the output tape. Let $q(\bar{n})$ be the control state reached at this point and c be the restriction of τ to \bar{n} . Let $b_x = \mathcal{A}_x(b, c)$ for every $\mathcal{A}_x \in \sigma$, and α be the sequence of assignments $x \mapsto b_x$. Then $(\llbracket t_1 \rrbracket_\tau^\alpha, \dots, \llbracket t_n \rrbracket_\tau^\alpha) \in \llbracket P \rrbracket$.
- For the deducibility predicate, $\mathcal{M}, \Pi, S, \sigma \models^c \hat{\phi}, t_1, \dots, t_n \vdash t$ if for all non-negligible $S' \subseteq S$, there is a non-negligible $S'' \subseteq S'$ such that there is a PT Turing machine \mathcal{A}_D such that, for all $\tau \in S''$, the unique valid computation (if there is one) of Π with respect to \mathcal{M}, τ that yields a configuration of \mathcal{M} , which is in the control state q_c with a bitstring b on the output tape, an actual frame ϕ_m and such that, letting $b_x = \mathcal{A}_x(b, \bar{c})$ where $\bar{c} = \tau(\bar{n})$ for the names \bar{n} in the current state, for every $\mathcal{A}_x \in \sigma$, and α be the sequence of assignments $x \mapsto b_x$, $\mathcal{A}_D(\llbracket \phi_m \rrbracket_\tau, \llbracket t_1 \rrbracket_\tau^\alpha, \dots, \llbracket t_n \rrbracket_\tau^\alpha, b) = \llbracket t \rrbracket_\tau^\alpha$.
- $\mathcal{M}, \Pi, S, \sigma \models^c t_1, \dots, t_n \vdash t$ is defined exactly as above, however removing ϕ .
- If P is an interpreted predicate, $\mathcal{M}, \Pi, S, \sigma \models P(t_1, \dots, t_n)$ iff there is an overwhelming subset $S' \subseteq S$ such that, for any $\tau \in S'$, the unique valid computation of Π with respect to \mathcal{M}, τ that yields a computational state $(q(\bar{n}, \bar{h}, \phi, \Theta), \bar{b}, \bar{b}', \gamma)$ in the control state q_c such that $P(t_1, \dots, t_n)$ is true in $q(\bar{n}, \bar{h}, \phi, \Theta)$. (Remember that the evaluation of such predicates do not depend on the model).

$\mathcal{M}, \Pi \models^{nmp} \theta$ (read “ \mathcal{M}, Π satisfies θ with non negligible probability”) if there is a non-negligible set S and a PPT machine \mathcal{A} such that $\mathcal{A}(n_1, \dots, n_k, b_1, \dots, b_k)$ returns 1 iff there is a $\tau \in S$ such that, for all i , $\tau(n_i) = b_i$ and $\mathcal{M}, \Pi, S \models^c \theta$.

Lemma 1. *If $\mathcal{M}, \Pi, S, \sigma \models^c \theta$ and $S' \subseteq S$, then we also have $\mathcal{M}, \Pi, S', \sigma \models^c \theta$*

Proof. We proceed by induction on θ . Since \mathcal{M}, Π are fixed, we sometimes omit these components.

- If θ is a atomic formula $P(t_1, \dots, t_n)$ and $P \notin \{\vdash, \text{fresh}(), \sqsubseteq\}$, then, by definition, there is an overwhelming subset $S_1 \subseteq S$ such that, for any $\tau \in S_1$, $(\llbracket t_1 \rrbracket_\tau^\alpha, \dots, \llbracket t_n \rrbracket_\tau^\alpha) \in \llbracket P \rrbracket$. If $S' \subseteq S$, we choose $S'_1 = S' \cap S_1$. It is an overwhelming subset of S' and, for any $\tau \in S'_1$, $(\llbracket t_1 \rrbracket_\tau^\alpha, \dots, \llbracket t_n \rrbracket_\tau^\alpha) \in \llbracket P \rrbracket$.
- If θ is a formula $\hat{\phi}, t_1, \dots, t_n \vdash t$, then for any non negligible $S_1 \subseteq S$, there is a non-negligible $S_2 \subseteq S_1$ and there is a machine \mathcal{A} , such that, for any $\tau \in S_2$, $\mathcal{A}(\llbracket \phi \rrbracket_\tau, \llbracket t_1 \rrbracket_\tau^\alpha, \dots, \llbracket t_n \rrbracket_\tau^\alpha) = \llbracket t \rrbracket_\tau^\alpha$. If $S' \subseteq S$ is non negligible, then any non-negligible $S'_1 \subseteq S'$ is also a non-negligible $S'_1 \subseteq S$, hence the result.
- Other atomic formulas with $=$ and \sqsubseteq are rather trivial.
- If $\theta = \neg\theta_1$, $\mathcal{M}, \Pi, S, \sigma \models^c \theta$ iff for any $S_1 \subseteq S$ such that $\mathcal{M}, \Pi, S_1, \sigma \models^c \theta_1$, $S_1 \cap S$ is negligible. In that case, for any $S' \subseteq S$, $S' \cap S_1$ is also negligible, hence the result (we do not use here the induction hypothesis).
- Let now $\theta = \theta_1 \vee \theta_2$, $S, \sigma \models^c \theta_1 \vee \theta_2$ and $S' \subseteq S$. Take any $S'' \subseteq S'$. Clearly, $S'' \subseteq S$, so by $S, \sigma \models^c \theta_1 \vee \theta_2$ there is a subset $S''' \subseteq S''$ such that $S''', \sigma \models^c \theta_1$ or $S''', \sigma \models^c \theta_2$. That exactly means $S', \sigma \models^c \theta_1 \vee \theta_2$.
- If $\theta = \theta_1 \wedge \theta_2$, we simply use the induction hypothesis for θ_1 and θ_2 , with the same $S' \subseteq S$.
- If $\theta = \theta_1 \rightarrow \theta_2$, then $\mathcal{M}, \Pi, S, \sigma \models^c \theta$ means exactly that for every $S_1 \subseteq S$, if $\mathcal{M}, \Pi, S_1, \sigma \models^c \theta_1$ then $\mathcal{M}, \Pi, S_1, \sigma \models^c \theta_2$. Clearly, the same will be true if $S_1 \subseteq S' \subseteq S$. Hence $\mathcal{M}, \Pi, S', \sigma \models^c \theta$ also holds.
- For $\theta = \exists x.\theta_1$, the proof is analogous to \vee and \rightarrow .
- For $\theta = \forall x.\theta_1$, the proof uses the induction hypothesis that the statement is true for θ . Since for all x , θ is satisfied on all subsets, $\forall x\theta$ is also satisfied.

Lemma 2. *With the above definitions, $\mathcal{M}, \Pi, S, \sigma \models^c \theta$ if and only if for all $S' \subseteq S$ non-negligible, there is a $S'' \subseteq S$ non-negligible such that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$*

Proof. The only if part is clear from the previous lemma that satisfaction carries over to subsets. The if proof again goes by induction.

- Suppose for all $S' \subseteq S$ non-negligible, there is a $S'' \subseteq S'$ non-negligible such that $\mathcal{M}, \Pi, S'', \sigma \models^c \neg\theta$. That means, for all $S' \subseteq S$, $\mathcal{M}, \Pi, S', \sigma \not\models^c \theta$, because, by Lemma 1, $\mathcal{M}, \Pi, S', \sigma \models^c \theta$ would imply $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$. So, we have $\mathcal{M}, \Pi, S, \sigma \models^c \neg\theta$ by the definition of the satisfaction of negation.
- Suppose for all $S' \subseteq S$ non-negligible, there is a $S'' \subseteq S'$ non-negligible such that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1 \vee \theta_2$. Then, by the definition of satisfaction of $\theta_1 \vee \theta_2$, there is a subset $S''' \subseteq S''$ such that either $\mathcal{M}, \Pi, S''', \sigma \models^c \theta_1$ or $\mathcal{M}, \Pi, S''', \sigma \models^c \theta_2$. Hence each S' has such a subset, which exactly means by our definition that $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1 \vee \theta_2$.

- Suppose for all $S' \subseteq S$ non-negligible, there is a $S'' \subseteq S'$ non-negligible such that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1 \wedge \theta_2$. Then, by the definition of satisfaction of $\theta_1 \wedge \theta_2$, both $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$ and $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$. So, using our induction assumption, we get that $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1$ and $\mathcal{M}, \Pi, S, \sigma \models^c \theta_2$, which exactly means that $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1 \wedge \theta_2$.
- Now suppose for all $S' \subseteq S$ non-negligible, there is a $S'' \subseteq S'$ non-negligible such that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1 \rightarrow \theta_2$. Take an arbitrary S_1 , and suppose $\mathcal{M}, \Pi, S_1, \sigma \models^c \theta_1$. We have to show that $\mathcal{M}, \Pi, S_1, \sigma \models^c \theta_2$. So for any $S' \subseteq S_1$ non-negligible, there is a $S'' \subseteq S'$ non-negligible such that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1 \rightarrow \theta_2$. By the definition of the satisfaction of entailment, this latter means $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$, because we had $\mathcal{M}, \Pi, S_1, \sigma \models^c \theta_1$ and satisfaction is preserved under taking subsets. So we have that for any $S' \subseteq S_1$ non-negligible, there is a $S'' \subseteq S'$ non-negligible such that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$. By the induction hypothesis on θ_2 , $\mathcal{M}, \Pi, S_1, \sigma \models^c \theta_2$.
- The case of $\forall x\theta$ is easy, because the statement holds for all θ , it also holds for $\forall x\theta$.
- Suppose for all $S' \subseteq S$ non-negligible, there is a $S'' \subseteq S'$ non-negligible such that $\mathcal{M}, \Pi, S'', \sigma \models^c \exists x\theta$. Take an arbitrary $S_1 \subseteq S$. We have to show that there is a $S_2 \subseteq S_1$ and there is an \mathcal{A}_x such that $\mathcal{M}, \Pi, S_2, \sigma, \mathcal{A}_x \models^c \theta$. We know that there is a $S'' \subseteq S_1$ with $\mathcal{M}, \Pi, S'', \sigma \models^c \exists x\theta$. That implies by the semantics of $\exists x\theta$ that there is a $S_2 \subseteq S''$ and there is an \mathcal{A}_x such that $\mathcal{M}, \Pi, S_2, \sigma, \mathcal{A}_x \models^c \theta$. Since $S_2 \subseteq S'' \subseteq S_1$, this completes the proof.

Corollary 1. $\mathcal{M}, \Pi, S, \sigma \models^c \theta$ if and only if for all $S' \subseteq S$, $\mathcal{M}, \Pi, S', \sigma \models^c \theta$ holds. $\mathcal{M}, \Pi, S, \sigma \not\models^c \theta$ if and only if there is a $S' \subseteq S$ such that $\mathcal{M}, \Pi, S', \sigma \not\models^c \theta$.

Lemma 3. *The De-Morgan identities hold for conjunction and disjunction as well as for exist and forall. Moreover, $\neg\neg\theta$ is satisfied iff θ is satisfied, $\neg\theta \wedge \theta$ is satisfied only on negligible sets, $\neg\theta \vee \theta$ is valid, $\theta_1 \rightarrow \theta_2$ is satisfied iff $\neg\theta_1 \vee \theta_2$ is satisfied. Moreover, the distributivity rules also hold: $\theta \wedge (\theta_1 \vee \theta_2)$ is satisfied iff $(\theta \wedge \theta_1) \vee (\theta \wedge \theta_2)$ is satisfied and $\theta \vee (\theta_1 \wedge \theta_2)$ is satisfied iff $(\theta \vee \theta_1) \wedge (\theta \vee \theta_2)$ is satisfied.*

Proof. The proof is based on Lemma 2. The following points are sufficient to prove.

- $\mathcal{M}, \Pi, S, \sigma \models^c \neg\neg\theta$ holds by the semantics of negation iff for all $S' \subseteq S$, we have $\mathcal{M}, \Pi, S', \sigma \not\models^c \neg\theta$. That again by the semantics of negation is equivalent with saying that iff for all $S' \subseteq S$ there is a $S'' \subseteq S'$ such that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$. That in turn is equivalent with $\mathcal{M}, \Pi, S, \sigma \models^c \theta$ by Lemma 2.
- It is clear that $\mathcal{M}, \Pi, S, \sigma \models^c \neg\theta \wedge \theta$ never holds, because $\mathcal{M}, \Pi, S, \sigma \models^c \neg\theta_1$ implies $\mathcal{M}, \Pi, S, \sigma \not\models^c \theta$, a contradiction. We only allowed contradictions on negligible sets.
- $\mathcal{M}, \Pi, S, \sigma \models^c \neg(\theta_1 \wedge \theta_2)$ holds iff for all $S' \subseteq S$, we have $\mathcal{M}, \Pi, S', \sigma \not\models^c \theta_1 \wedge \theta_2$, which means that for all $S' \subseteq S$, either $\mathcal{M}, \Pi, S', \sigma \not\models^c \theta_1$ or $\mathcal{M}, \Pi, S', \sigma \not\models^c \theta_2$. By Lemma 2, this is equivalent with that for all $S' \subseteq S$, there is a $S'' \subseteq S'$ such that either for all $S''' \subseteq S''$, $\mathcal{M}, \Pi, S''', \sigma \not\models^c \theta_1$, or for all $S''' \subseteq S''$, $\mathcal{M}, \Pi, S''', \sigma \not\models^c \theta_2$. This means exactly that for all $S' \subseteq S$, there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \neg\theta_1$, or $\mathcal{M}, \Pi, S'', \sigma \models^c \neg\theta_2$, which is the same as $\mathcal{M}, \Pi, S, \sigma \models^c \neg\theta_1 \vee \neg\theta_2$.

- For the equivalence of $\theta_1 \rightarrow \theta_2$ and $\neg\theta_1 \vee \theta_2$ we do the two directions separately.
 - Let now $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1 \rightarrow \theta_2$ hold. This implies that for all $S' \subseteq S$, if $\mathcal{M}, \Pi, S', \sigma \models^c \theta_1$ holds, then $\mathcal{M}, \Pi, S', \sigma \models^c \theta_2$ also holds. That is, either $\mathcal{M}, \Pi, S', \sigma \not\models^c \theta_1$, or $\mathcal{M}, \Pi, S', \sigma \models^c \theta_2$. By Lemma 2, it follows that for all $S' \subseteq S$, either there is a $S'' \subseteq S'$ with $\mathcal{M}, \Pi, S'', \sigma \models^c \neg\theta_1$, or $\mathcal{M}, \Pi, S', \sigma \models^c \theta_2$. This latter condition is equivalent with saying that for all $S'' \subseteq S'$, $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$, which clearly implies that there is a $S'' \subseteq S'$ with $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$. Putting these together, we have that for all $S' \subseteq S$, there is a $S'' \subseteq S'$ with either $\mathcal{M}, \Pi, S'', \sigma \models^c \neg\theta_1$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$.
 - In the other direction, suppose $\mathcal{M}, \Pi, S, \sigma \models^c \neg\theta_1 \vee \theta_2$ holds. To show that $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1 \rightarrow \theta_2$ also holds, take an arbitrary $S' \subseteq S$, and suppose $\mathcal{M}, \Pi, S', \sigma \models^c \theta_1$. Now take an arbitrary $S'' \subseteq S'$. By Lemma 2, for all $S''' \subseteq S''$, $\mathcal{M}, \Pi, S''', \sigma \models^c \theta_1$. But, $\mathcal{M}, \Pi, S, \sigma \models^c \neg\theta_1 \vee \theta_2$ implies, that there is a $S''' \subseteq S''$ such that either $\mathcal{M}, \Pi, S''', \sigma \models^c \neg\theta_1$ or $\mathcal{M}, \Pi, S''', \sigma \models^c \neg\theta_2$. So, by the foregoing, the latter must be the case. Since we showed that there is such a S''' for any S'' , we have by Lemma 2 that $\mathcal{M}, \Pi, S', \sigma \models^c \neg\theta_2$.
- Suppose $\mathcal{M}, \Pi, S, \sigma \models^c \neg\forall x\theta$, which is equivalent with saying that for all $S' \subseteq S$, $\mathcal{M}, \Pi, S', \sigma \not\models^c \forall x\theta$, which in turn is equivalent with saying that for all $S' \subseteq S$, there is an \mathcal{A}_x such that $\mathcal{M}, \Pi, S', \sigma, \mathcal{A}_x \not\models^c \theta$. This, by Lemma 2 is equivalent with that for all $S' \subseteq S$, there is an \mathcal{A}_x and an S'' such that $\mathcal{M}, \Pi, S'', \sigma, \mathcal{A}_x \models^c \neg\theta$. This is by the definition of the semantics of the existential quantification, equivalent with $\mathcal{M}, \Pi, S, \sigma \models^c \exists x\neg\theta$.
- For the equivalence of $\mathcal{M}, \Pi, S, \sigma \models^c \theta \wedge (\theta_1 \vee \theta_2)$ and $\mathcal{M}, \Pi, S, \sigma \models^c (\theta \wedge \theta_1) \vee (\theta \wedge \theta_2)$ we again do the two directions separately.
 - $\mathcal{M}, \Pi, S, \sigma \models^c \theta \wedge (\theta_1 \vee \theta_2)$ implies that $\mathcal{M}, \Pi, S, \sigma \models^c \theta$ and $\mathcal{M}, \Pi, S, \sigma \models^c \theta_1 \vee \theta_2$ hold, which in turn means by the semantics of the disjunction that $\mathcal{M}, \Pi, S, \sigma \models^c \theta$ and for every $S' \subseteq S$, there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$. By Corollary 1, we have that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$, so we have that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ and $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$, or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ and $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$, which in turn is equivalent with $\mathcal{M}, \Pi, S, \sigma \models^c (\theta \wedge \theta_1) \vee (\theta \wedge \theta_2)$.
 - For the other direction, assume now that $\mathcal{M}, \Pi, S, \sigma \models^c (\theta \wedge \theta_1) \vee (\theta \wedge \theta_2)$. This means that for every $S' \subseteq S$ there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ and $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$, or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ and $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$. Since $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ holds in both cases, this, on one hand implies that for every $S' \subseteq S$ there is a $S'' \subseteq S'$ such that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$, which by Lemma 2 means that $\mathcal{M}, \Pi, S, \sigma \models^c \theta$. On the other hand, we also have that for every $S' \subseteq S$ there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$, or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$, implying $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1 \vee \theta_2$. Hence, we arrive at $\mathcal{M}, \Pi, S, \sigma \models^c \theta \wedge (\theta_1 \vee \theta_2)$.
- Finally, for the equivalence of $\mathcal{M}, \Pi, S, \sigma \models^c \theta \vee (\theta_1 \wedge \theta_2)$ and $\mathcal{M}, \Pi, S, \sigma \models^c (\theta \vee \theta_1) \wedge (\theta \vee \theta_2)$ we again do the two directions separately.
 - If $\mathcal{M}, \Pi, S, \sigma \models^c \theta \vee (\theta_1 \wedge \theta_2)$ holds, then for every $S' \subseteq S$, there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1 \wedge \theta_2$. But this is equivalent with saying that for every $S' \subseteq S$, there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$, or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$ and $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$.

That in turn is equivalent with that for every $S' \subseteq S$, there is a $S'' \subseteq S'$ such that both either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$, and either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$. This implies that on one hand for every $S' \subseteq S$, there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$, and, on the other hand, for every $S' \subseteq S$, there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$. This is equivalent with $\mathcal{M}, \Pi, S, \sigma \models^c (\theta \vee \theta_1) \wedge (\theta \vee \theta_2)$.

- For the other direction, suppose now that we have $\mathcal{M}, \Pi, S, \sigma \models^c (\theta \vee \theta_1) \wedge (\theta \vee \theta_2)$, which is equivalent with that on one hand for every $S' \subseteq S$, there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$, and, on the other hand, for every $S' \subseteq S$, there is a $S'' \subseteq S'$ such that either $\mathcal{M}, \Pi, S'', \sigma \models^c \theta$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$. Fix an arbitrary $S' \subseteq S$. By the foregoing, there is a $S''' \subseteq S'$ such that either $\mathcal{M}, \Pi, S''', \sigma \models^c \theta$ or $\mathcal{M}, \Pi, S''', \sigma \models^c \theta_1$. If $\mathcal{M}, \Pi, S''', \sigma \models^c \theta$, then set $S'' = S'$. If $\mathcal{M}, \Pi, S''', \sigma \not\models^c \theta$, then we have $\mathcal{M}, \Pi, S''', \sigma \models^c \theta_1$, and by the foregoing, there is a $S'' \subseteq S'''$ such that $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_2$. For Lemma 1, $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1$, so $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1 \wedge \theta_2$. Hence, in either case, we have an $S'' \subseteq S'$ such that $\mathcal{M}, \Pi, S'', \sigma \not\models^c \theta$ or $\mathcal{M}, \Pi, S'', \sigma \models^c \theta_1 \wedge \theta_2$. Then, by the semantics of the disjunction, we have $\mathcal{M}, \Pi, S, \sigma \models^c \theta \vee (\theta_1 \wedge \theta_2)$.

3 Computational soundness

Definition 7 (Almost Covering Family of Decreasing Chains of Polynomial-Time Non-Negligible Sets of Coins). For a set of coins S , we say that a family \mathfrak{C} of chains of the form $S_1 \supseteq S_2 \supseteq \dots$, each non-negligible, PT computable subsets of S , almost covers S , if for any $S' \subseteq S$, for all i , there is a $\{S_j\}_{j=1}^\infty \in \mathfrak{C}$ such that the intersection of S_i and S' is non-negligible.

Let \mathfrak{C}^1 and \mathfrak{C}^2 be two such decreasing families, taking all elements $\{S_i^1\}_{i=1}^\infty \in \mathfrak{C}^1$ and all elements $\{S_i^2\}_{i=1}^\infty \in \mathfrak{C}^2$, the set of non-negligible decreasing chains of the form $\{S_i^1 \cap S_i^2\}_{i=1}^\infty$ is also an almost covering family. It is indeed decreasing, because $S_i^1 \cap S_i^2 \subseteq S_i^1 \subseteq S_{i-1}^1$ and $S_i^1 \cap S_i^2 \subseteq S_i^2 \subseteq S_{i-1}^2$, so $S_i^1 \cap S_i^2 \subseteq S_i^1 \subseteq S_{i-1}^1 \cap S_{i-1}^2$. Computability is preserved by intersection. For the almost covering property, consider any non-negligible subset $S' \subseteq S$. Since \mathfrak{C}^1 is almost covering, for all i there is an $\{S_j^1\}_{j=1}^\infty \in \mathfrak{C}^1$ such that $S' \cap S_i^1$ is non-negligible. Then, since \mathfrak{C}^2 is almost covering, there is an $\{S_j^2\}_{j=1}^\infty \in \mathfrak{C}^2$ such that $S' \cap S_i^1 \cap S_i^2$ is non-negligible.

Definition 8 (Conjunction of Almost Covering Family of Chains). We call the family of chains \mathfrak{C} created from two families of chains \mathfrak{C}^1 and \mathfrak{C}^2 via intersections in the above manner, the conjunction of \mathfrak{C}^1 and \mathfrak{C}^2 .

Definition 9 (Subchain). For two decreasing chains of PT computable, non-negligible sets of coins, $S_1 \supseteq S_2 \supseteq \dots$, and $S'_1 \supseteq S'_2 \supseteq \dots$, we say that $\{S_i\}_{i=1}^\infty$ is a subchain of $\{S'_i\}_{i=1}^\infty$ if for all i , $S'_i \supseteq S_i$.

Theorem 1. Let Π be a protocol, $s_1 \rightarrow \dots \rightarrow s_m$ be a symbolic transition sequence of Π and \mathcal{M} be a probabilistic polynomial time interactive Turing machine. If there is a

non-negligible set of coins S such that, for any $\tau \in S$, there is a sequence of transitions $(s_0, \mathbf{b}_0, \mathbf{b}'_0, \gamma_0) \rightarrow \dots \rightarrow (s_m, \mathbf{b}_m, \mathbf{b}'_m, \gamma_m)$ that is computationally valid w.r.t. \mathcal{M} , τ and γ_m is in the challenge state q_c , then for any formula θ , $\mathcal{M}, \Pi, S \models^c \theta$ implies there is a symbolic model \mathcal{S} such that $s_0 \rightarrow \dots \rightarrow s_m$ is a valid symbolic execution w.r.t. S and $\mathcal{S} \models \theta$.

Proof. We assume in this proof that there are only two predicate symbols: $=$ and \vdash . The extension to other predicate symbols is straightforward.

For any term t with free variables x_1, \dots, x_n and machines $\mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_n}$, and any sample $\tau \in S$, let $\llbracket t \rrbracket_{\tau, \sigma}$ be the computational interpretation of t , in which each variable x_i is interpreted according to $\sigma(\tau)(x_i) = \mathcal{A}_{x_i}(b_\tau, \tau(\mathbf{n}))$ if b_τ is the bitstring on the output tape of γ_m , and \mathbf{n} is the set of names in the state s_m , for the execution corresponding to τ .

Given a decreasing chain of non-negligible sets of coins $S \supseteq S_1 \supseteq S_2 \supseteq \dots$, we define a first-order structure $\mathcal{M}_{S_1 \supseteq S_2 \supseteq \dots}$ as follows. The domain of $\mathcal{M}_{S_1 \supseteq S_2 \supseteq \dots}$ is the set of terms built on the function symbols, the names and the additional constants \mathcal{A} for any PT machine \mathcal{A} . The interpretation of the predicate symbols is given, for any assignment σ of the variables x_1, \dots, x_n to machines $\mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_n}$ by:

- $\mathcal{M}_{S_1 \supseteq S_2 \supseteq \dots}, \sigma \models t = u$ iff there is an i such that $\forall \tau \in S_i, \llbracket t \rrbracket_{\tau, \sigma(\tau)} = \llbracket u \rrbracket_{\tau, \sigma(\tau)}$
- $\mathcal{M}_{S_1 \supseteq S_2 \supseteq \dots}, \sigma \models t_1, \dots, t_n \vdash t$ iff there is a PT algorithm \mathcal{A} , an i such that $\forall \tau \in S_i, \mathcal{A}(\llbracket t_1 \rrbracket_{\tau, \sigma(\tau)}, \dots, \llbracket t_n \rrbracket_{\tau, \sigma(\tau)}) = \llbracket t \rrbracket_{\tau, \sigma(\tau)}$.

Let

$$\mathcal{M}_S = \mathcal{M}_{S \supseteq S \supseteq \dots}$$

Remark 1. Notice, that the definition is such that for $S_1 \supseteq S_2 \supseteq \dots$ and $S'_1 \supseteq S'_2 \supseteq \dots$, if for some $m \in \mathbb{N}$, $S'_i = S_i$ for all $i > m$, then $\mathcal{M}_{S_1 \supseteq S_2 \supseteq \dots}, \sigma \models \theta$ if and only if $\mathcal{M}_{S'_1 \supseteq S'_2 \supseteq \dots}, \sigma \models \theta$. This is rather trivially true for θ atomic formula, and hence true for any formula.

Let θ be a formula with free variables x_1, \dots, x_n such that only atomic formulas are negated. We prove, by induction on θ that, if on a non-negligible set of coins S , $\mathcal{M}, \Pi, S, \sigma \models^c \theta$, then there is an almost covering family of chains \mathcal{C} such that for any decreasing chain of non-negligible subsets $S \supseteq S_1 \supseteq S_2 \supseteq \dots$ that is a subchain of some chain in \mathcal{C} , there is a decreasing chain of non-negligible subsets $S'_1 \supseteq S'_2 \supseteq \dots$ such that $S'_i \subseteq S_i$ for all $i = 1, 2, \dots$, and for any decreasing chain of non-negligible subsets $S''_1 \supseteq S''_2 \supseteq \dots$ with $S''_i \subseteq S'_i$ for all $i = 1, 2, \dots$, $\mathcal{M}_{S''_1 \supseteq S''_2 \supseteq \dots}, \sigma \models \theta$.

- Suppose $\theta \equiv t = u$. We know from Lemma 1 that $\mathcal{M}, \Pi, S, \sigma \models^c \theta$ implies $\mathcal{M}, \Pi, S', \sigma \models^c \theta$ for any subset $S' \subseteq S$. Hence, given any decreasing chain of non-negligible subsets $S \supseteq S_1 \supseteq S_2 \supseteq \dots$, it suffices to choose $S'_i = S_i$ for every i , and \mathcal{C} can be defined to have a single element, $\{S\}_{i=1}^\infty$.
- If $\theta \equiv t \neq u$ and $S \supseteq S_1 \supseteq S_2 \supseteq \dots$ is any decreasing sequence of non-negligible sets, let $S'_i = S_i$ for every i , and let \mathcal{C} have a single element, $\{S\}_{i=1}^\infty$. For any decreasing sequence of non-negligible sets $S''_i \subseteq S'_i$, for all i , since $S''_i, \sigma \models^c t \neq u$ by lemma 1, $\{\tau \in S''_i : \llbracket t \rrbracket_{\tau, \sigma} = \llbracket u \rrbracket_{\tau, \sigma}\}$ is negligible. Hence there is at least one $\tau \in S''_i$ such that $\llbracket t \rrbracket_{\tau, \sigma} \neq \llbracket u \rrbracket_{\tau, \sigma}$. Hence $\mathcal{M}_{S''_1 \supseteq S''_2 \supseteq \dots}, \sigma \not\models t = u$.

- For $\theta \equiv t_1, \dots, t_n \vdash t$, again given any decreasing chain of non-negligible subsets $S \supseteq S_1 \supseteq S_2 \supseteq \dots$, it suffices to choose $S'_i = S_i$, and \mathfrak{C} can be defined to have a single element, $\{S\}_{i=1}^\infty$.
- If $\theta \equiv \hat{\phi}, u_1, \dots, u_k \vdash t$, we may replace $\hat{\phi}$ with the frame ϕ_m of the symbolic state s_m (this is because for any $\tau \in S$, we reach the same symbolic state s_m), hence we are back to the previous case.
- If $\theta \equiv t_1, \dots, t_n \not\vdash t$, given any decreasing chain of non-negligible subsets $S \supseteq S_1 \supseteq S_2 \supseteq \dots$, it suffices to choose $S'_i = S_i$, and \mathfrak{C} to have a single element, $\{S\}_{i=1}^\infty$, as $\mathcal{M}, \Pi, S', \sigma \models^c t_1, \dots, t_n \vdash t$ is not true on any non-negligible S' .
- If $\theta \equiv \hat{\phi}, t_1, \dots, t_n \not\vdash t$, as before, we may replace $\hat{\phi}$ with the frame in s_m and we are back to the previous case.
- If $\theta \equiv \theta_1 \vee \theta_2$, then consider the family of sets \mathfrak{S} that consists of all sets $S'' \subseteq S$ such that either $S'', \sigma \models^c \theta_1$, or $S'', \sigma \models^c \theta_2$. Note, that $S, \sigma \models^c \theta_1 \vee \theta_2$ means by definition that any S' has a subset in \mathfrak{S} . We define \mathfrak{C} the following way. For a S'' in \mathfrak{S} , since either θ_1 or θ_2 is satisfied on it, the induction hypothesis gives a family of chains $\mathfrak{C}^{S''}$ that almost covers S'' . Take $\mathfrak{C} = \bigcup_{S'' \in \mathfrak{S}} \mathfrak{C}^{S''}$. The family we receive this way is almost covering S , because taking any non-negligible, PT computable set $S' \subseteq S$, there is an $S'' \in \mathfrak{S}$ with $S'' \subseteq S'$ non-negligible, and the chains of $\mathfrak{C}^{S''}$ all run in S'' , and hence in S' , so their intersections with S' are non-negligible. Take now any decreasing chain of non-negligible subsets $S \supseteq S_1 \supseteq S_2 \supseteq \dots$ that is a subchain of some chain in \mathfrak{C} . By construction, there is an S'' such that $\{S_i\}_{i=1}^\infty$ is a subchain of some chain in $\mathfrak{C}^{S''}$. But $\mathfrak{C}^{S''}$ was given by the induction hypothesis for the satisfaction of either θ_1 or θ_2 on S'' . Suppose the first, θ_1 . Then again by the induction hypothesis for θ_1 , there is a chain $S \supseteq S' \supseteq S'_1 \supseteq S'_2 \supseteq \dots$ such that $S'_i \subseteq S_i$ and for any non-negligible decreasing chain $S''_1 \supseteq S''_2 \supseteq \dots$ with $S''_i \subseteq S'_i$, $\mathcal{M}_{S'_1 \supseteq S''_2 \supseteq \dots}, \sigma \models \theta_1$. Then $\mathcal{M}_{S'_1 \supseteq S''_2 \supseteq \dots}, \sigma \models \theta$. That is, there is a chain $S \supseteq S' \supseteq S'_1 \supseteq S'_2 \supseteq \dots$ such that $S'_i \subseteq S_i$ and for any non-negligible decreasing chain $S''_1 \supseteq S''_2 \supseteq \dots$ with $S''_i \subseteq S'_i$, $\mathcal{M}_{S'_1 \supseteq S''_2 \supseteq \dots}, \sigma \models \theta$. The same procedure holds when on S_2 , θ_2 is satisfied.
- If $\theta \equiv \theta_1 \wedge \theta_2$, by definition, $S, \sigma \models^c \theta_1$ and $S, \sigma \models^c \theta_2$. Denoting by \mathfrak{C}^1 and \mathfrak{C}^2 the almost covering families given by the induction hypothesis for θ_1 and θ_2 respectively, take \mathfrak{C} to be the conjunction of them. This is also an almost covering family as we discussed earlier. Given any decreasing chain of non-negligible subsets $S \supseteq S_1 \supseteq S_2 \supseteq \dots$ that is a subchain of some chain in \mathfrak{C} , it is also a subchain of some chain in \mathfrak{C}^1 because all chains in \mathfrak{C} are subchains of chains in \mathfrak{C}^1 by construction. By the induction hypothesis for θ_1 , there is a chain $S \supseteq S'_{11} \supseteq S'_{12} \supseteq \dots$ with $S'_{1i} \subseteq S_i$, such that, for any non-negligible decreasing chain $S''_1 \supseteq S''_2 \supseteq \dots$ with $S''_i \subseteq S'_{1i}$ for all i , $\mathcal{M}_{S'_{11} \supseteq S''_2 \supseteq \dots}, \sigma \models \theta_1$. By induction hypothesis for θ_2 , since $\{S'_{1i}\}_{i=1}^\infty$ is a subchain of a chain in \mathfrak{C}^2 , there is a chain $S \supseteq S'_{21} \supseteq S'_{22} \supseteq \dots$ with $S'_{2i} \subseteq S'_{1i}$ such that, for any non-negligible decreasing chain $S''_1 \supseteq S''_2 \supseteq \dots$ with $S''_i \subseteq S'_{2i}$ for all i , $\mathcal{M}_{S'_{21} \supseteq S''_2 \supseteq \dots}, \sigma \models \theta_2$. Since $S'_{2i} \subseteq S'_{1i}$, by the choice of S'_{1i} , we also have that for any non-negligible decreasing chain $S''_1 \supseteq S''_2 \supseteq \dots$ with $S''_i \subseteq S'_{2i}$ for all i , $\mathcal{M}_{S'_{21} \supseteq S''_2 \supseteq \dots}, \sigma \models \theta_1$. Hence, for any non-negligible decreasing chain $S''_1 \supseteq S''_2 \supseteq \dots$ with $S''_i \subseteq S'_{2i}$ for all i , $\mathcal{M}_{S'_{21} \supseteq S''_2 \supseteq \dots}, \sigma \models \theta_1 \wedge \theta_2$. Thus, taking $S'_i = S'_{2i}$ works.

- If $\theta \equiv \exists x.\theta_1$, then consider the family of sets \mathfrak{S} that consists of all sets $S'' \subseteq S$ such that there is an \mathcal{A}_x , with $S'', \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_k}, \mathcal{A}_x \models^c \theta$. Since $S, \sigma \models^c \exists x.\theta_1$, by definition of the semantics, any S' has a subset in \mathfrak{S} . We define \mathfrak{C} the following way. For an S'' in \mathfrak{S} , since θ_1 is satisfied on it by some \mathcal{A}_x , the induction hypothesis gives an almost covering family $\mathfrak{C}^{S''}$. Take $\mathfrak{C} = \bigcup_{S'' \in \mathfrak{S}} \mathfrak{C}^{S''}$. The family we receive this way is an almost covering family, for the same reason as in the $\theta_1 \vee \theta_2$ case before. Take now any decreasing chain of non-negligible subsets $S \supseteq S_1 \supseteq S_2 \supseteq \dots$ that is a subchain of some chain in \mathfrak{C} . By construction, there is an S'' such that $\{S_i\}_{i=1}^\infty$ is a subchain of some chain in $\mathfrak{C}^{S''}$. But $\mathfrak{C}^{S''}$ was given by the induction hypothesis for the satisfaction of θ_1 on S'' . Then again by the induction hypothesis for θ_1 , there is a chain $S \supseteq S' \supseteq S'_1 \supseteq S'_2 \supseteq \dots$ such that $S'_i \subseteq S_i$ and for any non-negligible decreasing chain $S''_1 \supseteq S''_2 \supseteq \dots$ with $S''_i \subseteq S'_i$, $\mathcal{M}_{S''_1 \supseteq S''_2 \supseteq \dots}, \sigma, x \mapsto \mathcal{A}_x \models \theta_1$. But then this implies $\mathcal{M}_{S'_1 \supseteq S'_2 \supseteq \dots}, \sigma \models \exists x.\theta_1$. So the same $S'_1 \supseteq S'_2 \supseteq \dots$ works.
- If $\theta \equiv \forall x.\theta_1$, then for all $\mathcal{A}_x, S, \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_k}, \mathcal{A}_x \models^c \theta_1$. Enumerate all possible algorithms for \mathcal{A}_x : $\mathcal{A}_1, \mathcal{A}_2, \dots$. First consider that for $\mathcal{A}_1, S, \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_k}, \mathcal{A}_1 \models^c \theta_1$ holds. By induction hypothesis, there is an almost covering family of chains \mathfrak{C}^1 such that for all $S \supseteq S_1 \supseteq S_2 \supseteq \dots$ that is a subchain of some chain in \mathfrak{C}^1 , there is a chain $S'_{11} \supseteq S'_{12} \supseteq S'_{13} \supseteq \dots$ with $S'_{1i} \subseteq S_i$, such that, for any non-negligible $S''_1 \supseteq S''_2 \supseteq \dots$, if $S''_i \subseteq S'_{1i}$ for all i , then $\mathcal{M}_{S''_1 \supseteq S''_2 \supseteq \dots}, \sigma, x \mapsto \mathcal{A}_1 \models \theta_1$. Take now \mathcal{A}_2 . Then $S, \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_k}, \mathcal{A}_2 \models^c \theta_1$ holds. By the induction hypothesis, there is an almost covering family of chains \mathfrak{C}^2 with the usual properties. And so on, for all \mathcal{A}_j , we have \mathfrak{C}^j . We construct \mathfrak{C} the following way. We take all decreasing chains of the form

$$S_1^1 \supseteq S_2^1 \cap S_2^2 \supseteq S_3^1 \cap S_3^2 \cap S_3^3 \supseteq \dots \supseteq \bigcap_{j=1}^i S_i^j \supseteq \dots$$

where for each j , $\{S_i^j\}_{i=1}^\infty \in \mathfrak{C}^j$. That is,

$$\mathfrak{C} := \left\{ x \mid x = \left\{ \bigcap_{j=1}^i S_i^j \right\}_{i=1}^\infty, \{S_i^j\}_{i=1}^\infty \in \mathfrak{C}^j \right\}$$

Let's now fix $S_1 \supseteq S_2 \supseteq \dots$ such that it is a subchain of a chain in \mathfrak{C} . By construction, it is also a subchain of a chain in \mathfrak{C}^1 . Let $\{\bigcap_{j=1}^i S_i^j\}_{i=1}^\infty$ denote that chain. By the induction hypothesis for θ_1 , since $S, \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_k}, \mathcal{A}_1 \models^c \theta_1$, there is a chain $S'_{11} \supseteq S'_{12} \supseteq S'_{13} \supseteq \dots$ with $S'_{1i} \subseteq S_i$, such that, for any non-negligible $S''_1 \supseteq S''_2 \supseteq \dots$, if $S''_i \subseteq S'_{1i}$ for all i , then $\mathcal{M}_{S''_1 \supseteq S''_2 \supseteq \dots}, \sigma, x \mapsto \mathcal{A}_1 \models \theta_1$. Note now that the chain $S'_{11} \cap S_1^2 \supseteq S'_{12} \supseteq S'_{13} \supseteq \dots$ is in \mathfrak{C}^2 . Here, $S'_{11} \cap S_1^2 \supseteq S'_{12}$ follows from that $S_2^2 \supseteq S_2 \supseteq S'_{12}$ by the choice of S_2 . By the induction hypothesis, since $S, \mathcal{A}_{x_1}, \dots, \mathcal{A}_{x_k}, \mathcal{A}_2 \models^c \theta_1$, there is a chain $S'_{21} \supseteq S'_{22} \supseteq \dots$ with $S'_{21} \subseteq S'_{11} \cap S_1^2$ and $S'_{2i} \subseteq S'_{1i}$ for $i > 1$, such that, for any non-negligible chain $S''_1 \supseteq S''_2 \supseteq \dots$ such that $S''_i \subseteq S'_{2i}$ for all i , $\mathcal{M}_{S''_1 \supseteq S''_2 \supseteq \dots}, \sigma, x \mapsto \mathcal{A}_2 \models \theta_1$. But, because of Remark 1, it is also true that for the chain $S'_{11} \supseteq S'_{22} \supseteq S'_{23} \supseteq \dots$, for any non-negligible $S''_1 \supseteq S''_2 \supseteq \dots$, with $S''_1 \subseteq S'_{11}$, and $S''_i \subseteq S'_{2i}$ for $i = 2, 3, \dots$,

$\mathcal{M}_{S_1'' \supseteq S_2'' \supseteq \dots}, \sigma, x \mapsto \mathcal{A}_2 \models \theta_1$, as it does not matter what the first set is. Furthermore, since $S_{2i}' \subseteq S_{1i}'$ holds, we also have $\mathcal{M}_{S_1'' \supseteq S_2'' \supseteq \dots}, \sigma, x \mapsto \mathcal{A}_1 \models \theta_1$. Continuing in this manner, we get a chain $S_{11}' \supseteq S_{22}' \supseteq S_{33}' \supseteq \dots$. Then, take any chain $S_1'' \supseteq S_2'' \supseteq \dots$, with $S_{ii}'' \subseteq S_{ii}'$. Clearly, because of the construction, $S_{ii}'' \subseteq S_{1i}'$ also holds (as $S_{ii}' \subseteq S_{1i}'$). Hence we have $\mathcal{M}_{S_1'' \supseteq S_2'' \supseteq \dots}, \sigma, x \mapsto \mathcal{A}_1 \models \theta_1$. Further, since $S_{ii}'' \subseteq S_{2i}'$ for $i = 2, 3, \dots$, and $S_{11}'' \subseteq S_{11}'$, we also have $\mathcal{M}_{S_1'' \supseteq S_2'' \supseteq \dots}, \sigma, x \mapsto \mathcal{A}_2 \models \theta_1$. And so on, we have for all j , $\mathcal{M}_{S_1'' \supseteq S_2'' \supseteq \dots}, \sigma, x \mapsto \mathcal{A}_j \models \theta_1$. Now, if v is any term in the domain of our models, $\mathcal{M}_{S_1'' \supseteq S_2'' \supseteq \dots}, \sigma, x \mapsto v \models \theta_1$. Indeed, let v' be the term v , in which any \mathcal{A}_i occurring in v is replaced with a variable x_i and σ' be $x_i \mapsto \mathcal{A}_i$. The algorithm, that computes, for every $\tau \in S$, $\llbracket v' \rrbracket_{\tau, \sigma'}$ can be constructed from the \mathcal{A}_i and is PT. Hence there is an index n such that, for any $\tau \in S$, \mathcal{A}_n outputs $\llbracket v' \rrbracket_{\tau, \sigma'}$. Therefore, we also have $\mathcal{M}_{S_1'' \supseteq S_2'' \supseteq \dots}, \sigma \models \forall x \theta_1$, and that is what we wanted to prove. \square

The above result can be applied to a formula θ that is the conjunction of

- the intermediate conditions (that are part of the symbolic states) Θ
- finitely many computationally valid axioms A
- a formula that expresses the existence of an attack. **NotSec**

Then it can be read as follows: if there is a computational attack, corresponding to a symbolic trace $s_1 \rightarrow \dots \rightarrow s_m$, then this symbolic trace is valid in a model, which is also a model of A and **NotSec**.

Consider then a symbolic procedure, that discards only symbolic states, in which $\Theta \wedge A$ is inconsistent. Then the symbolic procedure will not miss any attack. More precisely, we get:

Theorem 2. *For a bounded number of sessions, if there is a computational attack, there is also a symbolic attack.*

In other words, if the protocol is symbolically secure, then it is also computationally secure.

It might be true for an unbounded number of sessions as well, but we need the boundedness assumption if we wish to derive the theorem from the theorem 1: The trick is, that in the bounded case, if there is a computational attack, there is also a computational attack corresponding to a fixed sequence of symbolic states. This is simply because the bounded number of sessions ensures that there are only finitely many possible sequences, and if there is a computational attack, that is, the property expressing the attack is satisfied on some non-negligible set, then it must be satisfied non-negligibly on one of the possible sequences.

4 The S4 Connection

Soon after we have completed the proofs of our Lemmas and the main Theorem, we realized that it is possible to give a purely logical presentation of our definitions and propositions. Observe first that if we think of non-negligible sets as possible worlds, and the subset relation as accessibility, then our computational semantics resembles to

Kripke semantics of modal logic. For any first-order formula θ , consider the transformation $\theta \rightarrow \theta^*$, where θ^* is a formula of first-order S4, and is defined recursively as follows:

- For any atomic formula θ , let $\theta^* \equiv \Box \Diamond \theta$.
- $(\neg \theta)^* \equiv \Box \neg \theta^*$
- $(\theta_1 \rightarrow \theta_2)^* \equiv \Box (\theta_1^* \rightarrow \theta_2^*)$
- $(\theta_1 \wedge \theta_2)^* \equiv (\theta_1^* \wedge \theta_2^*)$
- $(\theta_1 \vee \theta_2)^* \equiv \Box \Diamond (\theta_1^* \vee \theta_2^*)$
- $(\forall x \theta)^* \equiv \forall x \theta^*$
- $(\exists x \theta)^* \equiv \Box \Diamond \exists x \theta^*$

With these definitions, considering only S4 with the Barcan formula $(\forall x \Box \theta \rightarrow \Box \forall x \theta)$, Lemma 1 shows that for any θ , $\theta^* \rightarrow \Box \theta^*$. Lemma 2 shows that $\theta^* \leftrightarrow \Box \Diamond \theta^*$. Lemma 3 shows that our definitions are consistent with tautologies. Finally, once the proof of Theorem 1 is reduced to a purely logical one, the meaning of it becomes that any formula θ is derivable in first-order logic if and only if θ^* is derivable in first order S4 with the Barcan formula. Indeed, after some search of logic literature, we found a publication of Fitting [19] with exactly this result, using forcing techniques introduced by Cohen. If we don't want to use the Barcan formula, then the above is still all true, but $(\forall x \theta)^*$ has to be defined as $\Box \Diamond \forall x \theta^*$.

5 Examples of axioms

5.1 Examples of axioms that are computationally valid

- Increasing capabilities: $\hat{\phi} \vdash y \rightarrow \hat{\phi}, x \vdash y$
- Function of derivable items: $\hat{\phi} \vdash t_1 \wedge \hat{\phi} \vdash t_2 \wedge \dots \wedge \hat{\phi} \vdash t_n \rightarrow \hat{\phi} \vdash f(t_1, t_2, \dots, t_n)$
- Self derivability: $\hat{\phi}, t \vdash t$

The validity of these axioms is straightforward. We also include the following:

$$\text{No telepathy: } \text{fresh}(x, \hat{\phi}) \rightarrow \hat{\phi} \not\vdash x$$

whose computational soundness follows from the polynomial bound on the machines that interpret the deducibility relation on the one hand and the exponential number of interpretations of any names, on the other hand.

5.2 Secrecy Axiom

The intuitive meaning of the following axiom is that the adversary cannot derive the plaintext of a freshly generated encryption, unless its decryption key has been sent out, or the plaintext could be derived earlier.

Proposition 1. *If the encryption scheme is IND-CCA2, then the following axiom*

$$\theta = \forall_{tKR} \left(\text{RandGen}(K) \wedge \text{fresh}(R, \hat{\phi}) \wedge \hat{\phi}, \{t\}_{eK}^R \vdash t \rightarrow dK \sqsubseteq \hat{\phi} \vee \hat{\phi} \vdash t \right)$$

is computationally valid.

Proof. Suppose that it is not computationally valid. That is, there is a computational structure (\mathcal{M}, Π, S) , with $\mathcal{M}, \Pi, S \not\models \theta$. There are PPT machines $\mathcal{A} = (\mathcal{A}_t, \mathcal{A}_K, \mathcal{A}_R)$ such that $\mathcal{M}, \Pi, S, \mathcal{A} \not\models \text{fresh}(R, \hat{\phi}) \wedge \hat{\phi}, \{t\}_{e_K}^R \vdash t \rightarrow dK \sqsubseteq \hat{\phi} \vee \hat{\phi} \vdash t$. Therefore, there is a $S_1 \subseteq S$ non-negligible such that $\mathcal{M}, \Pi, S_1, \mathcal{A} \models \text{fresh}(R, \hat{\phi}) \wedge \hat{\phi}, \{t\}_{e_K}^R \vdash t$ and $\mathcal{M}, \Pi, S_1, \mathcal{A} \not\models dK \sqsubseteq \hat{\phi} \vee \hat{\phi} \vdash t$. We claim that the second implies that there is a non-negligible subset S_2 of S_1 such that $\mathcal{M}, \Pi, S_2, \mathcal{A} \models \neg(dK \sqsubseteq \hat{\phi})$ and $\mathcal{M}, \Pi, S_2, \mathcal{A} \not\models \hat{\phi} \vdash t$. To see this, consider the following:

- Take $S_2 = S_1 \setminus \{\tau \mid \text{the computation of } \mathcal{A} \text{ on } \tau \text{ yields a state } q \text{ such that } q \models dK \sqsubseteq \hat{\phi}\}$. Clearly, $\mathcal{M}, \Pi, S_2, \mathcal{A} \models \neg(dK \sqsubseteq \hat{\phi})$, and $\mathcal{M}, \Pi, S_1 \setminus S_2, \mathcal{A} \models dK \sqsubseteq \hat{\phi}$
- Since $\mathcal{M}, \Pi, S_1 \setminus S_2, \mathcal{A} \models dK \sqsubseteq \hat{\phi}$, we have $\mathcal{M}, \Pi, S_2, \mathcal{A} \not\models \hat{\phi} \vdash t$, because otherwise we would have $\mathcal{M}, \Pi, S_1, \mathcal{A} \models dK \sqsubseteq \hat{\phi} \vee \hat{\phi} \vdash t$ contradicting $\mathcal{M}, \Pi, S_1, \mathcal{A} \not\models dK \sqsubseteq \hat{\phi} \vee \hat{\phi} \vdash t$.

Since $\mathcal{M}, \Pi, S_2, \mathcal{A} \not\models \hat{\phi} \vdash t$, by the definition of the computational semantics of the derivability predicate, there is a subset S_4 of S_2 such that on all subsets of S_4 , there is no PT algorithm that computes the interpretation of t from the computational frame. Then it is straightforward to check that $\mathcal{M}, \Pi, S_4, \mathcal{A} \models \neg(\hat{\phi} \vdash t)$:

- Suppose it is not true, that is, $\mathcal{M}, \Pi, S_4, \mathcal{A} \not\models \neg(\hat{\phi} \vdash t)$.
- Then there is an $S_5 \subseteq S_4$ such that $\mathcal{M}, \Pi, S_5, \mathcal{A} \models \hat{\phi} \vdash t$.
- That implies that S_5 has a subset on which there is an algorithm that computes the interpretation t from the computational frame, a contradiction.

Since $S_4 \subseteq S_2$, we also have that $\mathcal{M}, \Pi, S_4, \mathcal{A} \models \neg(dK \sqsubseteq \hat{\phi})$, and since $S_4 \subseteq S_1$, we also have $\mathcal{M}, \Pi, S_4, \mathcal{A} \models \text{fresh}(R, \hat{\phi}) \wedge \hat{\phi}, \{t\}_{e_K}^R \vdash t$. That is, $\mathcal{M}, \Pi, S_4, \mathcal{A} \models \hat{\phi}, \{t\}_{e_K}^R \vdash t$ and $\mathcal{M}, \Pi, S_4, \mathcal{A} \models \text{fresh}(R, \hat{\phi})$ and $\mathcal{M}, \Pi, S_4, \mathcal{A} \models \neg(dK \sqsubseteq \hat{\phi})$ and $\mathcal{M}, \Pi, S_4, \mathcal{A} \models \neg(\hat{\phi} \vdash t)$. We have to create an adversary $\mathcal{A}_{\text{CCA2}}$ that wins the CCA2 game. Let $x = \{t\}_{e_K}^R$.

Since $\mathcal{M}, \Pi, S_4, \mathcal{A} \models \hat{\phi}, \{t\}_{e_K}^R \vdash t$ holds, there is an $S_5 \subseteq S_4$ and an algorithm \mathcal{C} that computes the interpretation of t from the interpretation of $\hat{\phi}$ and $\{t\}_{e_K}^R$ on S_5 . Clearly, $\mathcal{M}, \Pi, S_5, \mathcal{A} \models \text{fresh}(R, \hat{\phi})$ and $\mathcal{M}, \Pi, S_5, \mathcal{A} \models \neg(dK \sqsubseteq \hat{\phi})$ and $\mathcal{M}, \Pi, S_5, \mathcal{A} \models \neg(\hat{\phi} \vdash t)$. It may be the case that the S_5 we have chosen depends on evaluations of τ that are determined after \mathcal{M} reaches the challenge state q_c . However, clearly, if we include all possible future evaluations, the set that we receive this way, S' will still be such that there is an algorithm \mathcal{C} that computes the interpretation of t from the frame at the challenge state q_c and $\{t\}_{e_K}^R$ on S' . Moreover, it is easy to see that $\mathcal{M}, \Pi, S', \mathcal{A} \models \text{fresh}(R, \hat{\phi})$ and $\mathcal{M}, \Pi, S', \mathcal{A} \models \neg(dK \sqsubseteq \hat{\phi})$ and $\mathcal{M}, \Pi, S', \mathcal{A} \models \neg(\hat{\phi} \vdash t)$ because these are properties that depend only on conditions in the challenge stated, and not later ones.

Since $\mathcal{M}, \Pi, S', \mathcal{A} \models dK \not\sqsubseteq \hat{\phi}$, the decryption key has never been sent.

We show that we can construct an algorithm $\mathcal{A}_{\text{CCA2}}$ that breaks CCA2 security.

Let \mathcal{A}_{Π} mean the protocol adversary.

- $\mathcal{A}_{\text{CCA2}}$ generates computational keys that \mathcal{A}_{Π} uses, except for the one corresponding to K .

- The encryption oracle generates a random bit b .
- The encryption oracle generates a computational key and publishes its public part. \mathcal{A}_{CCA2} encrypts with this key for encryptions with K , except for t .
- \mathcal{A}_{CCA2} simulates both the agents and \mathcal{A}_Π : It computes all messages that the agents output according to their algorithm, and computes all messages that \mathcal{A}_Π outputs according to its algorithm. This way it builds up $\hat{\phi}$ and the bit strings corresponding to them as well as the equations.
- Whenever a decryption with dK has to be computed, there are two possibilities:
 - If the ciphertext was created by \mathcal{A}_{CCA2} using the encryption algorithm, then it knows the plaintext, so it can use it without decryption.
 - If the ciphertext was created in some other way, the decryption oracle is used. This can be freely done until x occurs.
- When \mathcal{A} reaches the challenge state q_c , using \mathcal{A}_t , \mathcal{A}_{CCA2} computes the bit string for t , and submits it to the encryption oracle as well as a random bit string that has the same length as the plaintext.
- According to our definition of satisfaction the computation by \mathcal{C} is based on the frame at the challenge state. We had $\mathcal{M}, \Pi, S', \mathcal{A} \models \text{fresh}(R, \hat{\phi})$, which means that R is independent of the items in ϕ . Further, since we included all future random choices in S' , R is also independent of S' . Hence having it encrypted by the encryption oracle will not lose any information as long as the oracle encrypts the correct bit.
- The encryption oracle encrypts the interpretation of t if $b = 0$, and encrypts the random bit string if $b = 1$. It gives the result c back to \mathcal{A}_{CCA2} .
- Run \mathcal{C} on the bit string c returned by the oracle and on the bit strings of ϕ_n .
- If
 - \mathcal{A}_{CCA2} receives the value for t back using c and if the execution is in S' , then \mathcal{A}_{CCA2} returns $b_{\mathcal{A}_{CCA2}} = 0$.
 - Otherwise \mathcal{A}_{CCA2} throws a fair coin and stores $b_{\mathcal{A}_{CCA2}} = 0$ or $b_{\mathcal{A}_{CCA2}} = 1$ with probability $1/2$.
- We have $\mathbf{Prob}\{b_{\mathcal{A}_{CCA2}} = b \mid S' \wedge b = 0\}$ (the conditional probability of $b_{\mathcal{A}_{CCA2}} = b$ given S' and $b = 0$) is negligibly different from 1 because in this case the oracle encrypts the correct string, and \mathcal{C} 's computations are employed on the correct bit string, and so it gives the interpretation of t . Note, we also use here that S' and the interpretation of R do not correlate.
- On the other hand, observe that $\mathbf{Prob}\{b_{\mathcal{A}_{CCA2}} = b \mid S' \wedge b = 1\} - 1/2$ is negligible. The reason is that when $b = 1$, the encryption oracle computes something that has nothing to do with the protocol and t . So the probability of computing t with or without the encryption in this case, is the same. But, remember, we had that $\mathcal{M}, \Pi, S', \mathcal{A} \models \hat{\phi} \not\models t$. This means that t cannot be computed without the encryption anywhere and therefore the adversary's computation on the fake encryption cannot give good result by more than negligible probability. So the adversary will end up throwing a coin in this case.
- Putting the previous two points together, we have $\mathbf{Prob}\{b_{\mathcal{A}_{CCA2}} = b \mid S'\} - \frac{1}{2}$ is non-negligible. Then, since outside S' , \mathcal{A}_{CCA2} throws a coin, $\mathbf{Prob}\{b_{\mathcal{A}_{CCA2}} = b\} - \frac{1}{2}$ is non-negligible, which means CCA2 security is broken. \square

6 Conclusions

We have shown a technique to define symbolic adversaries that are at least as strong as computational adversaries. The basic idea is that, instead of listing all manipulations the symbolic adversary is allowed to do, we allow the symbolic adversary to do anything unless it contradicts some axioms, which are derived from the limitations of the computational adversary. In a rather involved theorem, we showed that at least when only bounded number of protocol sessions are allowed, to any computational attack there is a corresponding symbolic attack. Further, we have shown a few axioms that arise from the limitations of computational adversaries, and which are to limit the symbolic adversary. Besides some rather trivially valid axioms, we showed the validity of a "secrecy axiom", that relies on IND-CCA2 security.

From our method, we can derive a verification procedure, simulating the (symbolic) protocol rules, and checking at each computation step the consistency of the formulas expressing that transitions are enabled, together with the axioms and the negation of the security properties. In order to automate this process we mainly need a (hopefully efficient) procedure checking the consistency of such a set of constrained formulas. This is future work. We are however optimistic, because the examples of axioms that we considered yield a saturated set of constrained formulas (as defined in [24]). On the other hand, as shown in [10], the consistency of ground clauses, together with a saturated set of clauses, can be performed in polynomial time.

We carried out a proof of a two sessions NSL, showing what are the minimal assumptions that guarantee its correctness, but we need to design an automated tool, in order to carry out further experiments. Also extensions of the results to indistinguishability properties could be investigated.

References

1. M. Abadi and P. Rogaway. Reconciling two views of cryptography: the computational soundness of formal encryption. In *Proc. 1st IFIP International Conference on Theoretical Computer Science*, volume 1872 of *Lecture Notes in Computer Science*, Sendai, Japan, 2000.
2. Martín Abadi, Bruno Blanchet, and Hubert Comon-Lundh. Models and proofs of protocol security: A progress report. In *Computer Aided Verification (CAV)*, volume 5643 of *Lecture Notes in Computer Science*, pages 35–49. Springer, 2009.
3. M. Backes and B. Pfitzmann. Symmetric encryption in a simulatable dolev-yao style cryptographic library. In *Proc. IEEE Computer Security Foundations workshop*, 2004.
4. M. Backes and B. Pfitzmann. Limits of the cryptographic realization of Dolev-Yao style XOR. In *Proc. 10th European Symposium on Research in Computer Security (ESORICS)*, 2005.
5. M. Backes, B. Pfitzmann, and M. Waidner. A composable cryptographic library with nested operations. In *Proc. 10th ACM Conference on Computer and Communications Security (CCS'03)*, 2003.
6. M. Backes, B. Pfitzmann, and M. Waidner. The reactive simulatability (rsim) framework for asynchronous systems. *Information and Computation*, 205(12), 2007.
7. Michael Backes, Dennis Hofheinz, and Dominique Unruh. Cosp: A general framework for computational soundness proofs. In *ACM CCS 2009*, pages 66–78, November 2009. Preprint on IACR ePrint 2009/080.

8. G. Bana and H. Comon-Lundh. Towards unconditional soundness: Computationally complete symbolic attacker. In *Proceedings of POST'12*, Lecture Notes in Computer Science, 2012.
9. G. Bana, K. Hasebe, and M. Okada. Secrecy-oriented first-order logical analysis of cryptographic protocols, 2010. <http://eprint.iacr.org/2010/080>.
10. David Basin and Harald Ganzinger. Automated complexity analysis based on ordered resolution. *Journal of the Association of Computing Machinery*, 48(1):70–109, January 2001.
11. Bruno Blanchet. An automatic security protocol verifier based on resolution theorem proving (invited tutorial). In *20th International Conference on Automated Deduction (CADE-20)*, Tallinn, Estonia, July 2005.
12. Bruno Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Transactions on Dependable and Secure Computing*, 5(4):193–207, 2008. Special issue IEEE Symposium on Security and Privacy 2006.
13. Ran Canetti and Tal Rabin. Universal composition with joint state. In *Proceedings of the 23rd Annual International Cryptology Conference (CRYPTO)*, volume 2729 of *Lecture Notes in Computer Science*, pages 265–281. Springer, 2003.
14. Hubert Comon-Lundh and Véronique Cortier. Computational soundness of observational equivalence. In *Proc. ACM Conf. Computer and Communication Security (CCS)*, 2008.
15. Hubert Comon-Lundh and Véronique Cortier. How to prove security of communication protocols? a discussion on the soundness of formal models w.r.t. computational ones. In Christoph Dürr and Thomas Schwentick, editors, *Proceedings of the 28th Annual Symposium on Theoretical Aspects of Computer Science (STACS'11)*, volume 9 of *Leibniz International Proceedings in Informatics*, pages 29–44, Dortmund, Germany, March 2011. Leibniz-Zentrum für Informatik.
16. V. Cortier and B. Warinschi. Computationally sound, automated proofs for security protocols. In *Proc. 14th European Symposium on Programming (ESOP'05)*, volume 3444 of *Lecture Notes in Computer Science*, pages 157–171, 2005.
17. Anupam Datta, Ante Derek, John C. Mitchell, Vitaly Shmatikov, and Mathieu Turuani. Probabilistic polynomial-time semantics for a protocol security logic. In *Proceedings of 32nd International Colloquium on Automata, Languages and Programming (ICALP)*, volume 3580 of *Lecture Notes in Computer Science*, pages 16–29. Springer, 2005.
18. A. Armando et al. The AVISPA Tool for the automated validation of internet security protocols and applications. In *Computer Aided Verification, (CAV'05)*, volume 3576 of *Lecture Notes in Computer Science*, pages 281–285, 2005.
19. Melvin Fitting. An embedding of classical logic in s4. *The Journal of Symbolic Logic*, 35(4):529–534, 1970.
20. Harald Ganzinger and Robert Nieuwenhuis. Constraints and theorem proving. In *Constraints in Computational Logics: Theory and Applications*, volume 2002 of *Lecture Notes in Computer Science*, pages 159–201, 2001.
21. Vladimir Lifschitz. Closed-world databases and circumscription. *Artif. Intell.*, 27(2):229–235, 1985.
22. Gavin Lowe. Breaking and fixing the needham-schroeder public-key protocol using *fd*. In *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 1055 of *Lecture Notes in Computer Science*, pages 147–166, 1996.
23. J. Millen and V. Shmatikov. Constraint solving for bounded-process cryptographic protocol analysis. In *Proc. 8th ACM Conference on Computer and Communications Security*, 2001.
24. Robert Nieuwenhuis and Albert Rubio. Paramodulation-based theorem proving. In *Handbook of Automated Reasoning*, pages 371–443. Elsevier and MIT Press, 2001.
25. Peter Ryan, Steve Schneider, Michael Goldsmith, Gavin Lowe, and Bill Roscoe. *The Modelling and Analysis of Security Protocols*. Addison Wesley, 2000.

26. Dominique Unruh. Computational soundness of hash functions. Presented at the 6th workshop on Formal and Computational Cryptography (FCC), July 2010.
27. Dominique Unruh. The impossibility of computationally sound xor, July 2010. Preprint on IACR ePrint 2010/389.
28. Bogdan Warinschi. A computational analysis of the needham-schroeder protocol. In *16th Computer security foundation workshop (CSFW)*, pages 248–262. IEEE, 2003.

A Axioms for the Needham-Schroeder-Lowe protocol

In our NSL proof, we assumed that A executes initiator sessions only, and B executes responder sessions only. We further assumed that they don't initiate sessions with themselves. We also assumed that all agents other than A and B are corrupted, so their keys are available to the adversary. Clearly, if we prove security in this case, it also holds with other honest agents. We used the convention $\langle x, y, z \rangle \equiv \langle x, \langle y, z \rangle \rangle$. We showed that in a symbolic execution, violating the secrecy of nonces N_1 and N_2 , as well as violating the agreement and authentication properties are inconsistent with our axioms.

Computationally Sound Axioms Used

- $x = x$, and the substitutability (congruence) property of equal terms holds for $=, \vdash$ predicates.
- Self derivability: $\hat{\phi}, \mathbf{x}, x \vdash x$
- Increasing capabilities: $\hat{\phi}, \mathbf{x} \vdash y \longrightarrow \hat{\phi}, \mathbf{x}, x \vdash y$
- Commutativity: If \mathbf{x}' is a permutation of \mathbf{x} , then $\hat{\phi}, \mathbf{x} \vdash y \longrightarrow \hat{\phi}, \mathbf{x}' \vdash y$
- Transitivity of derivability: $\hat{\phi}, \mathbf{x} \vdash \mathbf{y} \wedge \hat{\phi}, \mathbf{x}, \mathbf{y} \vdash \mathbf{z} \longrightarrow \hat{\phi}, \mathbf{x} \vdash \mathbf{z}$
- Functions are derivable: $\hat{\phi}, \mathbf{x} \vdash f(\mathbf{x})$
- No telepathy: $\text{fresh}(x; \hat{\phi}) \longrightarrow \hat{\phi} \not\vdash x$
- Fresh items are independent:

$$\text{fresh}(x; \hat{\phi}, \mathbf{x}) \wedge \text{RandGen}(N) \wedge \mathbf{x} \preceq \hat{\phi} \wedge \hat{\phi}, \mathbf{x}, x \vdash N \longrightarrow \hat{\phi}, \mathbf{x} \vdash N \vee x = N$$

- Special to IND-CCA encryption:

- Secrecy:

$$\begin{aligned} & \text{RandGen}(K) \wedge eK \sqsubseteq \hat{\phi} \wedge \text{fresh}(R; \hat{\phi}, \mathbf{x}, x) \wedge \mathbf{x} \preceq \hat{\phi} \wedge x \preceq \hat{\phi} \\ & \wedge \hat{\phi}, \mathbf{x}, \{x\}_{eK}^R \vdash x \\ & \longrightarrow dK \sqsubseteq \hat{\phi}, \mathbf{x}, x \vee \hat{\phi}, \mathbf{x} \vdash x \end{aligned}$$

- Non-malleability (assuming there is only one kind of encryption and pairing):

$$\begin{aligned} & \text{RandGen}(N) \wedge \text{RandGen}(K) \wedge eK \sqsubseteq \hat{\phi} \wedge N \sqsubseteq \hat{\phi} \wedge \mathbf{x} \preceq \hat{\phi} \\ & \wedge \hat{\phi}, \mathbf{x} \vdash y \wedge \hat{\phi}, \mathbf{x}, \text{dec}(y, dK) \vdash N \wedge \forall xR(y = \{x\}_{eK}^R \rightarrow \{x\}_{eK}^R \not\sqsubseteq \hat{\phi}) \\ & \longrightarrow dK \sqsubseteq \hat{\phi}, \mathbf{x} \vee \hat{\phi}, \mathbf{x} \vdash N \end{aligned}$$

Here, new constraints are

- $t \sqsubseteq \hat{\phi}, s_1, \dots, s_n$, where s_1, \dots, s_n and t are closed terms:
 $\mathcal{M}, \sigma, \langle t_1, \dots, t_m \rangle, \bar{n} \models t \sqsubseteq \hat{\phi}, s_1, \dots, s_n$ if t is a subterm of some t_i or some s_i
- $t \preceq \hat{\phi}$, where t is closed:
 $\mathcal{M}, \sigma, \langle t_1, \dots, t_m \rangle, \bar{n} \models t \preceq \hat{\phi}$ if for every handle h of t , $\hat{\phi} \vdash h$.
- And

$$\text{fresh}(x; \hat{\phi}, \mathbf{x}) = \text{RandGen}(x) \wedge x \not\sqsubseteq \hat{\phi}, \mathbf{x}$$

Further Needed Axiom (The implementation needs to satisfy this too)

For this protocol, we need an additional axiom, namely that for an honestly generated nonce N ,

$$\neg W(\pi_2(N)).$$

That is, the second projection of a nonce can never be a name (by overwhelming probability on a non-negligible set).