

Automated analysis of security protocols with global state

Abstract—Security APIs, key servers and protocols that need to keep the status of transactions, require to maintain a global, non-monotonic state, e.g., in the form of a database or register. However, existing automated verification tools do not support the analysis of such stateful security protocols – sometimes because of fundamental reasons, such as the encoding of the protocol as Horn clauses, which are inherently monotonic. An exception is the recent tamarin prover which allows specifying protocols as multiset rewrite (MSR) rules, a formalism expressive enough to encode state. As multiset rewriting is a “low-level” specification language with no direct support for concurrent message passing, encoding protocols correctly is a difficult and error-prone process.

We propose a process calculus which is a variant of the applied pi calculus with constructs for manipulation of a global state by processes running in parallel. We show that this language can be translated to MSR rules whilst preserving all security properties expressible in a dedicated first-order logic for security properties. The translation has been implemented in a prototype tool which uses the tamarin prover as a backend. We apply the tool to several case studies among which a simplified fragment of PKCS#11, the Yubikey security token, and an optimistic contract signing protocol.

I. INTRODUCTION

Automated analysis of security protocols has been extremely successful. Using automated tools flaws have been for instance discovered in the Google Single Sign On Protocol [4], in commercial security tokens implementing the PKCS#11 standard [8], and one may also recall Lowe’s attack [18] on the Needham-Schroeder public key protocol 17 years after its publication. While efficient tools such as ProVerif [6], AVISPA [3] or Maude-NPA [12] exist, these tools fail to analyze protocols that require *non-monotonic global state*, i.e., some database, register or memory location that can be read and altered by different parallel threads. In particular ProVerif, one of the most efficient and widely used protocol analysis tools, relies on an abstraction that encodes protocols in first-order Horn clauses. This abstraction is well suited for the monotonic knowledge of an attacker (who never forgets), makes the tool extremely efficient for verifying an unbounded number of protocol sessions and allows to build on existing techniques for Horn clause resolution. However, Horn clauses are inherently monotonic: once a fact is true it cannot be set to false anymore. As a result, even though ProVerif’s input language, a variant of the applied pi calculus [1], allows a priori encodings of

a global memory, the abstractions performed by ProVerif introduce false attacks.

A prominent example where non-monotonic global state appears are security APIs, such as the RSA PKCS#11 standard [20], IBM’s CCA [9] or the trusted platform module (TPM) [23]. They have been known to be vulnerable to logical attacks for some time [17], [7] and formal analysis has shown to be a valuable tool to identify attacks and find secure configurations. One promising paradigm for analyzing security APIs is to regard them as a participant in a protocol and use existing analysis tools. However, Herzog [15] already identified not accounting for mutable global state as a major barrier to the application of security protocol analysis tools to verify security APIs. Apart from security APIs many other protocols need to maintain databases: key servers need to store the status of keys, in optimistic contract signing protocols a trusted party maintains the status of a contract, RFID protocols maintain the status of tags and more generally websites may need to store the current status of transactions.

Our contributions: We propose a tool for analyzing protocols that may involve non-monotonic global state. We designed a new process calculus that extends the applied pi calculus by defining, in addition to the usual constructs for specifying concurrent processes, constructs for explicitly manipulating global state. This calculus serves as the tool’s input language. The heart of our tool is a translation from this extended applied pi calculus to a set of multiset rewrite rules that can then be analyzed by the tamarin tool [21], [22] which we use as a backend. We prove the correctness of this translation and show that it preserves all properties expressible in a dedicated first order logic for expressing security properties. As a result, relying on the tamarin prover, we can analyze protocols without bounding the number of sessions, nor making any abstractions. Moreover it allows to model a wide range of cryptographic primitives by the means of equational theories. As the underlying verification problem is undecidable, tamarin may not terminate. However, it offers an interactive mode with a GUI which allows to manually guide the tool in its proof. Our specification language includes support for private channels, global state and locking mechanisms (which are crucial to write meaningful programs in which concurrent threads manipulate a common memory). The translation has been carefully engineered in order to favor termination by tamarin. We illustrate the tool on several case studies: a simple security API in the style of PKCS#11, a complex case study of the Yubikey security

device, as well as several examples analyzed by other tools that aim at analyzing stateful protocols. In all of these case studies we were able to avoid restrictions that were necessary in previous works.

Related work: The most closely related work is the StatVerif tool by Arapinis et al. [2]. They propose an extension of the applied pi calculus, similar to ours, which is translated to Horn clauses and analyzed by the ProVerif tool. Their translation is sound but allows for false attacks, limiting the scope of protocols that can be analyzed. Moreover, StatVerif can only handle a finite number of memory cells: When analyzing an optimistic contract signing protocol this appeared to be a limitation and only the status of a single contract was modeled, providing a manual proof to justify the correctness of this abstraction. Finally, StatVerif is limited to the verification of secrecy properties. As illustrated by the Yubikey case study, our work is more general and we are able to analyze complex injective correspondance properties.

Mödersheim [19] proposed a language with support for sets together with an abstraction where all objects that belong to the same sets are identified. His language, which is an extension of the low level AVISPA intermediate format, is compiled into Horn clauses that are then analyzed using ProVerif. His approach is tightly linked to this particular abstraction limiting the scope of applicability. Mödersheim also discusses the need for a more high-level specification level which we provide in this work.

There has also been work tailored to particular applications. In [11], Delaune *et al.* show by a dedicated hand proof that for analyzing PKCS#11 one may bound the message size. Their analysis still requires to artificially bound the number of keys. Similarly in spirit, Delaune *et al.* [10] give a dedicated result for analyzing protocols based on the TPM and its registers. However, the number of reboots (which reinitialize registers) needs to be limited.

Guttman [14] also extended the strand space model by adding support for state. While the protocol execution is modeled using the classical strand spaces model, state is modeled by a multiset of facts, and manipulated by multiset rewrite rules. The extended model has been used for analyzing by hand an optimistic contract signing protocol. As of now, protocol analysis in the strand space model with state has not been mechanized yet.

In the goal of relating different approaches for protocol analysis Bistarelli et al. [5] also proposed a translation from a process algebra to multiset rewriting: they do however not consider private channels, have no support for global state and assume that processes have a particular structure. These limitations significantly simplify the translation and its correctness proof. Moreover their work does not include any tool support for automated verification.

Obviously any protocol that we are able to analyze can be directly analyzed by the tamarin prover [21], [22] as

we use it as a backend. Indeed, tamarin has already been used for analyzing a model of the Yubikey device [16], the case studies presented with Mödersheim’s abstraction, as well as those presented with StatVerif. It is furthermore able to reproduce the aforementioned results on PKCS#11 [11] and the TPM [10] – moreover, it does so without bounding the number of keys, security devices, reboots, etc. Contrary to ProVerif, tamarin sometimes requires additional *typing lemmas* which are used to guide the proof. These lemmas need to be written by hand (but are proved automatically). In our case studies we also needed to provide a few such lemmas manually. In our opinion, an important disadvantage of tamarin is that protocols are modeled as a set of multiset rewrite rules. This representations is very low level and far away from actual protocol implementations, making it very difficult to model a protocol correctly. Encoding private channels, nested replications and locking mechanisms directly as multiset rewrite rules is a tricky and error prone task. As a result we observed that, in practice, the protocol models tend to be simplified. For instance, locking mechanisms are often omitted, modeling protocol steps as a single rule and making them effectively atomic. Such more abstract models may obscure issues in concurrent protocol steps and increase the risk of implicitly excluding attacks in the model that are well possible in a real implementation, e. g., race conditions. Using a more high-level specification language, such as our process calculus, arguably eases protocol specification and overcomes some of these risks.

II. PRELIMINARIES

Terms and equational theories: As usual in symbolic protocol analysis we model messages by abstract terms. Therefore we define an order-sorted term algebra with the sort *msg* and two incomparable subsorts *pub* and *fresh*. For each of these subsorts we assume a countably infinite set of names, *FN* for fresh names and *PN* for public names. Fresh names will be used to model cryptographic keys and nonces while public names model publicly known values. We furthermore assume a countably infinite set of variables for each sort *s*, \mathcal{V}_s and let \mathcal{V} be the union of the set of variables for all sorts. We write $u : s$ when the name or variable u is of sort s . Let Σ be a signature, i.e., a set of function symbols, each with an arity. We denote by \mathcal{T}_Σ the set of well-sorted terms built over Σ , *PN*, *FN* and \mathcal{V} . For a term t we denote by $names(t)$, respectively $vars(t)$ the set of names, respectively variables, appearing in t . The set of ground terms, i.e., terms without variables, is denoted by \mathcal{M}_Σ . When Σ is fixed or clear from the context we often omit it and simply write \mathcal{T} for \mathcal{T}_Σ and \mathcal{M} for \mathcal{M}_Σ .

We equip the term algebra with an equational theory E , that is a finite set of equations of the form $M = N$ where $M, N \in \mathcal{T}$. From the equational theory we define the binary relation $=_E$ on terms, which is the smallest equivalence relation containing equations in E that is closed

under application of function symbols, bijective renaming of names and substitution of variables by terms of the same sort. Furthermore, we require E to distinguish different fresh names, i. e., $\forall a, b \in FN : a \neq b \Rightarrow a \neq_E b$.

Example. Symmetric encryption can be modelled using a signature

$$\Sigma = \{ \text{senc}/2, \text{sdec}/2, \text{encSucc}/2, \text{true}/0 \}$$

and an equational theory defined by

$$\text{sdec}(\text{senc}(m, k), k) = m \quad \text{encSucc}(\text{senc}(x, y), y) = \text{true}$$

The last equation allows to check whether a term can be decrypted with a certain key.

For the rest of the paper we assume that E refers to some fixed equational theory and that the signature and equational theory always contain symbols and equations for pairing and projection, i. e., $\{\langle \cdot, \cdot \rangle, \text{fst}, \text{snd}\} \subseteq \Sigma$ and equations $\text{fst}(\langle x, y \rangle) = x$ and $\text{snd}(\langle x, y \rangle) = y$ are in E . We will sometimes use $\langle x_1, x_2, \dots, x_n \rangle$ as a shortcut for $\langle x_1, \langle x_2, \langle \dots, \langle x_{n-1}, x_n \rangle \dots \rangle \rangle$.

We also use the usual notion of positions for terms. A position p is a sequence of positive integers and $t|_p$ denotes the subterm of t at position p .

Facts: We also assume an unsorted signature Σ_{fact} , disjoint from Σ . The set of *facts* is defined as

$$\mathcal{F} := \{F(t_1, \dots, t_k) \mid t_i \in \mathcal{T}_\Sigma, F \in \Sigma_{\text{fact}} \text{ of arity } k\}.$$

Facts will be used both to annotate protocols, by the means of events, and for defining multiset rewrite rules. We partition the signature Σ_{fact} into linear and persistent fact symbols. We suppose that Σ_{fact} always contains a unary, persistent symbol $!K$ and a linear, unary symbol Fr . Given a sequence or set of facts S we denote by $\text{lfacts}(S)$ the multiset of all linear facts in S and $\text{pfacts}(S)$ the set of all persistent facts in S . By notational convention facts whose identifier starts with ‘!’ will be persistent. \mathcal{G} denotes the set of ground facts, i. e., the set of facts that does not contain variables. For a fact f we denote by $\text{ginsts}(f)$ the set of ground instances of f . This notation is also lifted to sequences and sets of facts as expected.

Substitutions: A substitution σ is a partial function from variables to terms. We suppose that substitutions are well-typed, i. e., they only map variables of sort s to terms of sort s , or of a subsort of s . We denote by $\sigma = \{t^1/x_1, \dots, t^n/x_n\}$ the substitution whose domain is $\mathbf{D}(\sigma) = \{x_1, \dots, x_n\}$ and which maps x_i to t_i . As usual we homomorphically extend σ to apply to terms and facts and use a postfix notation to denote its application, e. g., we write $t\sigma$ for the application of σ to the term t . A substitution σ is grounding for a term t if $t\sigma$ is ground. Given function g we define the function $f := g[a \mapsto b]$ with $\text{dom}(f) = \text{dom}(g) \cup \{a\}$ as $f(a) := b$ and $f(x) := g(x)$ for $x \neq a$.

$$\begin{aligned} \langle M, N \rangle ::= & x, y, z \in \mathcal{V} \\ & \mid p \in PN \\ & \mid n \in FN \\ & \mid f(M_1, \dots, M_n) \quad (f \in \Sigma \text{ of arity } n) \\ \langle P, Q \rangle ::= & 0 \\ & \mid P \mid Q \\ & \mid !P \\ & \mid \nu n; P \\ & \mid \text{out}(M, N); P \\ & \mid \text{in}(M, N); P \\ & \mid \text{if } M=N \text{ then } P \text{ [else } Q] \\ & \mid \text{event } F ; P \quad (F \in \mathcal{F}) \\ & \mid \text{insert } M, N; P \\ & \mid \text{delete } M; P \\ & \mid \text{lookup } M \text{ as } x \text{ in } P \text{ [else } Q] \\ & \mid \text{lock } M; P \\ & \mid \text{unlock } M; P \\ & \mid [L] \text{ } \neg[A] \rightarrow [R]; P \quad (L, R, A \in \mathcal{F}^*) \end{aligned}$$

Figure 1. Syntax

Sets, sequences and multisets: We write \mathbb{N}_n for the set $\{1, \dots, n\}$. Given a set S we denote by S^* the set of finite sequences of elements from S and by $S^\#$ the set of finite multisets of elements from S . We use the superscript $\#$ to annotate usual multiset operation, e. g. $S_1 \cup^\# S_2$ denotes the multiset union of multisets S_1, S_2 . Given a multiset S we denote by $\text{set}(S)$ the set of elements in S . The sequence consisting of elements e_1, \dots, e_n will be denoted by $[e_1, \dots, e_n]$ and the empty sequence is denoted by $[]$. We denote by $|S|$ the length, i. e., the number of elements of the sequence. We use \cdot for the operation of adding an element either to the start or to the end, e. g., $e_1 \cdot [e_2, e_3] = [e_1, e_2, e_3] = [e_1, e_2] \cdot e_3$. Given a sequence S , we denote by $\text{idx}(S)$ the set of positions in S , i. e., \mathbb{N}_n when S has n elements, and for $i \in \text{idx}(S)$ S_i denotes the i th element of the sequence. Set membership modulo E is denoted by \in_E and defined as $e \in_E S$ if $\exists e' \in S. e' =_E e$. \subset_E and $=_E$ are defined for sets in a similar way. Application of substitutions are lifted to sets, sequences and multisets as expected. By abuse of notation we sometimes interpret sequences as sets or multisets; the applied operators should make the implicit cast clear.

III. A CRYPTOGRAPHIC PI CALCULUS WITH EXPLICIT STATE

A. Syntax and informal semantics

Our calculus is a variant of the applied pi calculus [1]. In addition to the usual operators for concurrency, replication, communication and name creation, it offers several constructs for reading and updating an explicit global state. The grammar for processes is described in Figure 1.

0 denotes the terminal process. $P \mid Q$ is the parallel execution of processes P and Q and $!P$ the replication of P , allowing an unbounded number of sessions in protocol executions. The construct $\nu n; P$ binds the name n in P and models the generation of a fresh, random value. Processes $\text{out}(M, N); P$ and $\text{in}(M, N); P$ represent the output, respectively input, of message N on channel M . Readers familiar with the applied pi calculus [1] may note that we opted for the possibility of pattern matching in the input construct, rather than merely binding the input to a variable x . The process $\text{if } M=N \text{ then } P \text{ else } Q$ will execute P if $M =_E N$ and Q otherwise. The event construct is merely used for annotating processes and will be useful for stating security properties. For readability we sometimes omit to write **else** Q when Q is 0, as well as trailing 0 processes.

The remaining constructs are used for manipulating state and are new compared to the applied pi calculus. We offer two different mechanisms for state. The first construct is *functional* and allows to associate a value to a key. The construct $\text{insert } M, N$ binds the value N to a key M . Successive inserts allow to change this binding. The delete M operation simply “undefines” the mapping for the key M . The lookup M as x in P else Q allows to retrieve the value associated to M , binding it to the variable x in P . If the mapping is undefined for M the process behaves as Q . The lock and unlock constructs allow to gain exclusive access to a resource M . This is essential for writing protocols where parallel processes may read and update a common memory. We additionally offer another kind of global state in form of a multiset of ground facts, as opposed to the previously introduced functional store. This multiset can be altered using the construct $[L] \text{ } \neg[A] \rightarrow [R]; P$, which tries to match each fact in the sequence L to facts in the current multiset and, if successful, adds the corresponding instance of facts R to the store. The facts A are used as annotations in a similar way to events. The purpose of this construct is to provide access to the underlying notion of state in tamarin, but we stress that it is distinct from the previously introduced functional state, and its use is only advised to expert users. We allow this “low-level” form of state manipulation in addition to the *functional* state, as it offers a great flexibility and has shown useful in one of our case studies. This style of state manipulation is similar to the state extension in the strand space model [14] and the underlying specification language of the tamarin tool [21], [22]. Note that, even though those stores are distinct (which is a restriction imposed by our translation), data can be moved from one to another, for example as follows: $\text{lookup 'store1' as } x \text{ in } [] \text{ } \neg[\] \rightarrow [\text{store2}(x)]$.

In the following example, which will serve as our running example, we model a security API that, even though much simplified, illustrates the most salient issues that occur in the analysis of security APIs such as PKCS#11 [11], [8].

Example. We consider a security device that allows the creation of keys in its secure memory. The user can access the device via an API. If he creates a key, he obtains a handle, which he can use to let the device perform operations on his behalf. For each handle the device also stores an attribute which defines what operations are permitted for this handle. The goal is that the user can never gain knowledge of the key, as the user’s machine might be compromised. We model the device by the following process (we use $\text{out}(m)$ as a shortcut for $\text{out}(c, m)$ for a public channel c):

$!P_{new} \mid !P_{set} \mid !P_{dec} \mid !P_{wrap}$, where

$P_{new} := \nu h; \nu k; \text{event NewKey}(h, k);$
 $\text{insert } \langle \text{'key'}, h \rangle, k;$
 $\text{insert } \langle \text{'att'}, h \rangle, \text{'dec'}; \text{out}(h)$

In the first line, the device creates a new handle h and a key k and, by the means of the event $\text{NewKey}(h, k)$, logs the creation of this key. It then stores the key that belongs to the handle by associating the pair $\langle \text{'key'}, h \rangle$ to the value of the key k . In the next line, $\langle \text{'att'}, h \rangle$ is associated to a public constant ‘dec’. Intuitively, we use the public constants ‘key’ and ‘att’ to distinguish two databases. The process

$P_{set} := \text{in}(h); \text{insert } \langle \text{'att'}, h \rangle, \text{'wrap'}$

allows the attacker to change the attribute of a key from the initial value ‘dec’ to another value ‘wrap’. If a handle has the ‘dec’ attribute set, it can be used for decryption:

$P_{dec} := \text{in}(\langle h, c \rangle); \text{lookup } \langle \text{'att'}, h \rangle \text{ as } a \text{ in}$
 $\text{if } a = \text{'dec'} \text{ then}$
 $\text{lookup } \langle \text{'key'}, h \rangle \text{ as } k \text{ in}$
 $\text{if } \text{encSucc}(c, k) = \text{true} \text{ then}$
 $\text{event DecUsing}(k, \text{sdec}(c, k));$
 $\text{out}(\text{sdec}(c, k))$

The first lookup stores the value associated to $\langle \text{'att'}, h \rangle$ in a . The value is compared against ‘dec’. If the comparison and another lookup for the associated key value k succeeds, we check whether decryption succeeds and, if so, output the plaintext.

If a key has the ‘wrap’ attribute set, it might be used to encrypt the value of a second key:

$P_{wrap} := \text{in}(\langle h_1, h_2 \rangle); \text{lookup } \langle \text{'att'}, h_1 \rangle \text{ as } a_1 \text{ in}$
 $\text{if } a_1 = \text{'wrap'} \text{ then}$
 $\text{lookup } \langle \text{'key'}, h_1 \rangle \text{ as } k_1 \text{ in}$
 $\text{lookup } \langle \text{'key'}, h_2 \rangle \text{ as } k_2 \text{ in}$
 $\text{event Wrap}(k_1, k_2);$
 $\text{out}(\text{senc}(k_2, k_1))$

The bound names of a process are those that are bound by νn . We suppose that all names of sort *fresh* appearing in the process are under the scope of such a binder. Free names must be of sort *pub*. A variable x can be bound in

three ways: (i) by the construct `lookup M as x`, or (ii) $x \in \text{vars}(N)$ in the construct `in(M, N)` and x is not under the scope of a previous binder, (iii) $x \in \text{vars}(L)$ in the construct `[L] -[A]→ [R]` and x is not under the scope of a previous binder. While the construct `lookup M as x` always acts as a binder, the input and `[L] -[A]→ [R]` constructs do not rebind an already bound variable but perform pattern matching. For instance in the process

$$P = \text{in}(c, f(x)); \text{in}(c, g(x))$$

x is bound by the first input and pattern matched in the second. It might seem odd that `lookup` acts as a binder, while input does not. We justify this decision as follows: as P_{dec} and P_{wrap} in the previous example show, lookups appear often after input was received. If `lookup` were to use pattern matching, the following process

$$P = \text{in}(c, x); \text{lookup 'store' as } x \text{ in } P'$$

might unexpectedly perform a check if ‘store’ contains the message given by the adversary, instead of binding the content of ‘store’ to x , due to an undetected clash in the naming of variables.

A process is ground if it does not contain any free variables. We denote by $P\sigma$ the application of the homomorphic extension of the substitution σ to P . As usual we suppose that the substitution only applies to free variables. We sometimes interpret the syntax tree of a process as a term and write $P|_p$ to refer to the subprocess of P at position p (where $|$, `if` and `lookup` are interpreted as binary symbols, all other constructs as unary).

B. Semantics

Frames and deduction: Before giving the formal semantics of our calculus we introduce the notions of frame and deduction. A *frame* consists of a set of fresh names \tilde{n} and a substitution σ and is written $\nu\tilde{n}.\sigma$. Intuitively a frame represents the sequence of messages that have been observed by an adversary during a protocol execution and secrets \tilde{n} generated by the protocol, a priori unknown to the adversary. Deduction models the capacity of the adversary to compute new messages from the observed ones.

Definition 1 (Deduction). *We define the deduction relation $\nu\tilde{n}.\sigma \vdash t$ as the smallest relation between frames and terms defined by the deduction rules in Figure 2.*

Example. If one key is used to wrap a second key, then, if the intruder learns the first key, he can deduce the second. For $\tilde{n} = k_1, k_2$ and $\sigma = \{ \text{senc}(k_2, k_1) / x_1, k_1 / x_2 \}$, $\nu\tilde{n}.\sigma \vdash k_2$, as witnessed by the proof tree given in Figure 3.

Operational semantics: We can now define the operational semantics of our calculus. The semantics is defined by a labelled transition relation between process configurations. A *process configuration* is a 6-tuple $(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P}, \sigma, \mathcal{L})$ where

- $\mathcal{E} \subseteq FN$ is the set of fresh names generated by the processes;
- $\mathcal{S} : \mathcal{M}_\Sigma \rightarrow \mathcal{M}_\Sigma$ is a partial function modeling the functional store;
- $\mathcal{S}^{\text{MS}} \subseteq \mathcal{G}^\#$ is a multiset of ground facts and models the multiset of stored facts;
- \mathcal{P} is a multiset of ground processes representing the processes executed in parallel;
- σ is a ground substitution modeling the messages output to the environment;
- $\mathcal{L} \subseteq \mathcal{M}_\Sigma$ is the set of currently acquired locks.

The transition relation is defined by the rules described in Figure 4. Transitions are labelled by sets of ground facts. For readability we omit empty sets and brackets around singletons, i.e., we write \rightarrow for $\xrightarrow{\emptyset}$ and \xrightarrow{f} for $\xrightarrow{\{f\}}$. We write \rightarrow^* for the reflexive, transitive closure of \rightarrow (the transitions that are labelled by the empty sets) and write \xrightarrow{f}^* for $\rightarrow^* \xrightarrow{f} \rightarrow^*$. We can now define the set of traces, i.e., possible executions, that a process admits.

Definition 2 (Traces of P). *Given a ground process P we define the set of traces of P as*

$$\begin{aligned} \text{traces}^{pi}(P) = & \left\{ [F_1, \dots, F_n] \mid (\emptyset, \emptyset, \emptyset, \{P\}, \emptyset, \emptyset) \right. \\ & \xrightarrow{F_1} (\mathcal{E}_1, \mathcal{S}_1, \mathcal{S}_1^{\text{MS}}, \mathcal{P}_1, \sigma_1, \mathcal{L}_1) \\ & \xrightarrow{F_2} \dots \xrightarrow{F_n} (\mathcal{E}_n, \mathcal{S}_n, \mathcal{S}_n^{\text{MS}}, \mathcal{P}_n, \sigma_n, \mathcal{L}_n) \left. \right\} \end{aligned}$$

Example. In Figure 5 we display the transitions that illustrate how the first key is created on the security device in our running example and witness that $[\text{NewKey}(h', k')] \in \text{traces}^{pi}(P)$.

IV. LABELLED MULTISSET REWRITING

We now define the syntax and semantics of labelled multiset rewriting rules, which constitute the input language of the tamarin tool [21].

Definition 3 (Multiset rewrite rule). *A labelled multiset rewrite rule ri is a triple (l, a, r) , $l, a, r \in \mathcal{F}^*$, written $l -[a]→ r$. We call $l = \text{prems}(ri)$ the premises, $a = \text{actions}(ri)$ the actions, and $r = \text{conclusions}(ri)$ the conclusions of the rule.*

Definition 4 (Labelled multiset rewriting system). *A labelled multiset rewriting system is a set of labelled multiset rewrite rules R , such that each rule $l -[a]→ r \in R$ satisfies the following conditions:*

- l, a, r do not contain fresh names
- r does not contain Fr-facts

A labelled multiset rewriting system is called well-formed, if additionally

- *for each $l' -[a']→ r' \in E$ $\text{ginsts}(l -[a]→ r)$ we have that $\cap_{r''=Er'} \text{names}(r'') \cap FN \subseteq \cap_{l''=El'} \text{names}(l'') \cap FN$.*

$$\begin{array}{c}
\frac{a \in FN \cup PN \quad a \notin \tilde{n}}{\nu\tilde{n}.\sigma \vdash a} \text{ DNAME} \qquad \frac{\nu\tilde{n}.\sigma \vdash t \quad t =_E t'}{\nu\tilde{n}.\sigma \vdash t'} \text{ DEQ} \\
\frac{x \in \mathbf{D}(\sigma)}{\nu\tilde{n}.\sigma \vdash x\sigma} \text{ DFRAME} \qquad \frac{\nu\tilde{n}.\sigma \vdash t_1 \cdots \nu\tilde{n}.\sigma \vdash t_n \quad f \in \Sigma^k}{\nu\tilde{n}.\sigma \vdash f(t_1, \dots, t_n)} \text{ DAPPL}
\end{array}$$

Figure 2. Deduction rules.

$$\frac{\frac{x_1 \in \mathbf{D}(\sigma)}{\nu\tilde{n}.\sigma \vdash \text{senc}(k_2, k_1)} \quad \frac{x_2 \in \mathbf{D}(\sigma)}{\nu\tilde{n}.\sigma \vdash k_1}}{\nu\tilde{n}.\sigma \vdash \text{sdec}(\text{senc}(k_2, k_1), k_1)} \quad \text{sdec}(\text{senc}(k_2, k_1), k_1) =_E k_2}{\nu\tilde{n}.\sigma \vdash k_2}$$

Figure 3. Proof tree witnessing that $\nu\tilde{n}.\sigma \vdash k_2$

Standard operations:

$$\begin{array}{l}
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{0\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P}, \sigma, \mathcal{L}) \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P|Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P, Q\}, \sigma, \mathcal{L}) \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{!P\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{!P, P\}, \sigma, \mathcal{L}) \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\nu a; P\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E} \cup \{a'\}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\{a'/a\}\}, \sigma, \mathcal{L}) \\
\hspace{15em} \text{if } a' \text{ is fresh} \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P}, \sigma, \mathcal{L}) \xrightarrow{K(M)} (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P}, \sigma, \mathcal{L}) \quad \text{if } \nu\mathcal{E}.\sigma \vdash M \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{out}(M, N); P\}, \sigma, \mathcal{L}) \xrightarrow{K(M)} (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\}, \sigma \cup \{N/x\}, \mathcal{L}) \\
\hspace{15em} \text{if } x \text{ is fresh and } \nu\mathcal{E}.\sigma \vdash M \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{in}(M, N); P\}, \sigma, \mathcal{L}) \xrightarrow{K(\langle M, N \rangle)} (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\tau\}, \sigma, \mathcal{L}) \\
\hspace{15em} \text{if } \exists \tau. \tau \text{ is grounding for } N, \nu\mathcal{E}.\sigma \vdash M, \nu\mathcal{E}.\sigma \vdash N\tau \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{out}(M, N); P; \text{in}(M', N'); Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P, Q\tau\}, \sigma, \mathcal{L}) \\
\hspace{15em} \text{if } M =_E M' \text{ and } \exists \tau. N =_E N'\tau \text{ and } \tau \text{ grounding for } N' \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{if } M = N \text{ then } P \text{ else } Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\}, \sigma, \mathcal{L}) \quad \text{if } M =_E N \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{if } M = N \text{ then } P \text{ else } Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{Q\}, \sigma, \mathcal{L}) \quad \text{if } M \neq_E N \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{event}(F); P\}, \sigma, \mathcal{L}) \xrightarrow{F} (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\}, \sigma, \mathcal{L})
\end{array}$$

Operations on global state:

$$\begin{array}{l}
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{insert } M, N; P\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}[M \mapsto N], \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\}, \sigma, \mathcal{L}) \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{delete } M; P\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}', \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\}, \sigma, \mathcal{L}) \\
\hspace{15em} \text{where } \mathcal{S}'(x) = \begin{cases} \mathcal{S}(x) & \text{if } x \neq_E M \\ \perp & \text{otherwise} \end{cases} \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{lookup } M \text{ as } x \text{ in } P \text{ else } Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\{V/x\}\}, \sigma, \mathcal{L}) \\
\hspace{15em} \text{if } \mathcal{S}(N) =_E V \text{ is defined and } N =_E M \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{lookup } M \text{ as } x \text{ in } P \text{ else } Q\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{Q\}, \sigma, \mathcal{L}) \\
\hspace{15em} \text{if } \mathcal{S}(N) \text{ is undefined for all } N =_E M \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{lock } M; P\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\}, \sigma, \mathcal{L} \cup \{M\}) \\
\hspace{15em} \text{if } M \notin_E \mathcal{L} \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{\text{unlock } M; P\}, \sigma, \mathcal{L}) \longrightarrow (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{P\}, \sigma, \mathcal{L} \setminus \{M' \mid M' =_E M\}) \\
(\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}}, \mathcal{P} \cup\# \{[l \text{ } \neg[a] \rightarrow r]; P\}, \sigma, \mathcal{L}) \xrightarrow{a'} (\mathcal{E}, \mathcal{S}, \mathcal{S}^{\text{MS}} \setminus \text{lfacts}(l') \cup\# r', \mathcal{P} \cup\# \{P\tau\}^\#, \sigma, \mathcal{L}) \\
\hspace{15em} \text{if } \exists \tau, l', a', r'. \tau \text{ grounding for } l \text{ } \neg[a] \rightarrow r, l' \text{ } \neg[a'] \rightarrow r' =_E (l \text{ } \neg[a] \rightarrow r)\tau, \\
\hspace{15em} \text{lfacts}(l') \subseteq\# \mathcal{S}^{\text{MS}}, \text{pfacts}(l') \subseteq \mathcal{S}^{\text{MS}}
\end{array}$$

Figure 4. Operational semantics

$$\begin{aligned}
& (\emptyset, \emptyset, \emptyset, \underbrace{\{!P_{new}, !P_{set} !P_{dec} !P_{wrap}\}^\#}_{=: \mathcal{P}'}, \emptyset, \emptyset) \rightarrow (\emptyset, \emptyset, \emptyset, \{P_{new}\}^\# \cup^\# \mathcal{P}', \emptyset, \emptyset) \\
& \rightarrow (\emptyset, \emptyset, \emptyset, \{\nu h; \nu k; \text{event NewKey}(h, k); \dots\}^\# \cup^\# \mathcal{P}', \emptyset, \emptyset) \\
& \rightarrow^* (\{h', k'\}, \emptyset, \emptyset, \{\text{event NewKey}(h', k'); \dots\}^\# \cup^\# \mathcal{P}', \emptyset, \emptyset) \\
& \xrightarrow{\text{NewKey}(h', k')} (\{h', k'\}, \emptyset, \emptyset, \{\text{insert } \langle \text{'key'}, h' \rangle, k'; \dots\}^\# \cup^\# \mathcal{P}', \emptyset, \emptyset) \\
& \rightarrow^* (\{h', k'\}, \mathcal{S}, \emptyset, \{\text{out}(h'); 0\}^\# \cup^\# \mathcal{P}', \emptyset, \emptyset) \rightarrow^* (\{h', k'\}, \mathcal{S}, \emptyset, \mathcal{P}', \{h'/x_1\}, \emptyset) \\
& \text{where } \mathcal{S}(\langle \text{'key'}, h' \rangle) = k' \text{ and } \mathcal{S}(\langle \text{'att'}, h' \rangle) = \text{'dec'}.
\end{aligned}$$

Figure 5. Example of transitions modelling the creation of a key on a PKCS#11-like device

We define one distinguished rule FRESH which is the only rule allowed to have Fr-facts on the right-hand side

$$\text{FRESH} : [] \dashv\vdash [\text{Fr}(x : \text{fresh})]$$

The semantics of the rules is defined by a labelled transition relation.

Definition 5 (Labelled transition relation). *Given a multiset rewriting system R we define the labeled transition relation $\rightarrow_R \subseteq \mathcal{G}^\# \times \mathcal{P}(\mathcal{G}) \times \mathcal{G}^\#$ as*

$$S \xrightarrow{a}_R ((S \setminus^\# \text{lfacts}(l)) \cup^\# r)$$

if and only if $l \dashv\vdash a \rightarrow r \in_E \text{ginsts}(R \cup \text{FRESH})$, $\text{lfacts}(l) \subseteq^\# S$ and $\text{pfacts}(l) \subseteq S$.

Definition 6 (Executions). *Given a multiset rewriting system R we define its set of executions as*

$$\begin{aligned}
\text{exec}^{mstr}(R) &= \left\{ \emptyset \xrightarrow{A_1}_R \dots \xrightarrow{A_n}_R S_n \mid \right. \\
&\quad \forall a, i, j: 0 \leq i \neq j < n. \\
&\quad \left. (S_{i+1} \setminus^\# S_i) = \{\text{Fr}(a)\} \Rightarrow (S_{j+1} \setminus^\# S_j) \neq \{\text{Fr}(a)\} \right\}
\end{aligned}$$

The set of executions consists of transition sequences that respect freshness, i. e., for a given name a the fact $\text{Fr}(a)$ is only added once, or in other words the rule FRESH is at most fired once for each name. We define the set of traces in a similar way as for processes.

Definition 7 (Traces). *The set of traces is defined as*

$$\begin{aligned}
\text{traces}^{mstr}(R) &= \left\{ [A_1, \dots, A_n] \mid \forall 0 \leq i \leq n. A_i \neq \emptyset \right. \\
&\quad \left. \text{and } \emptyset \xrightarrow{A_1}_R \dots \xrightarrow{A_n}_R S_n \in \text{exec}^{mstr}(R) \right\}
\end{aligned}$$

where \xrightarrow{A}_R is defined as $\xrightarrow{\emptyset}_R^* \xrightarrow{A}_R \xrightarrow{\emptyset}_R^*$.

Note that both for processes and multiset rewrite rules the set of traces is a sequence of sets of facts.

V. SECURITY PROPERTIES

In the tamarin tool [21] security properties are described in an expressive two-sorted first-order logic. The sort temp is used for time points, $\mathcal{V}_{\text{temp}}$ are the temporal variables.

Definition 8 (Trace formulas). *A trace atom is either false \perp , a term equality $t_1 \approx t_2$, a timepoint ordering $i < j$, a timepoint equality $i \doteq j$, or an action $F@i$ for a fact $F \in \mathcal{F}$ and a timepoint i . A trace formula is a first-order formula over trace atoms.*

As we will see in our case studies this logic is expressive enough to analyze a variety of security properties, including complex injective correspondence properties.

To define the semantics, let each sort s have a domain $\mathbf{D}(s)$. $\mathbf{D}(\text{temp}) = \mathcal{Q}$, $\mathbf{D}(\text{msg}) = \mathcal{M}$, $\mathbf{D}(\text{fresh}) = \text{FN}$, and $\mathbf{D}(\text{pub}) = \text{PN}$. A function $\theta : \mathcal{V} \rightarrow \mathcal{M} \cup \mathcal{Q}$ is a valuation if it respects sorts, that is, $\theta(\mathcal{V}_s) \subset \mathbf{D}(s)$ for all sorts s . If t is a term, $t\theta$ is the application of the homomorphic extension of θ to t .

Definition 9 (Satisfaction relation). *The satisfaction relation $(tr, \theta) \models \varphi$ between trace tr , valuation θ and trace formula φ is defined as follows:*

$$\begin{aligned}
(tr, \theta) \models \perp & \quad \text{never} \\
(tr, \theta) \models F@i & \quad \text{iff } \theta(i) \in \text{idx}(tr) \text{ and } F\theta \in_E tr_{\theta(i)} \\
(tr, \theta) \models i < j & \quad \text{iff } \theta(i) < \theta(j) \\
(tr, \theta) \models i \doteq j & \quad \text{iff } \theta(i) = \theta(j) \\
(tr, \theta) \models t_1 \approx t_2 & \quad \text{iff } t_1\theta =_E t_2\theta \\
(tr, \theta) \models \neg\varphi & \quad \text{iff } \text{not } (tr, \theta) \models \varphi \\
(tr, \theta) \models \varphi_1 \wedge \varphi_2 & \quad \text{iff } (tr, \theta) \models \varphi_1 \text{ and } (tr, \theta) \models \varphi_2 \\
(tr, \theta) \models \exists x : s. \varphi & \quad \text{iff } \text{there is } u \in \mathbf{D}(s) \text{ such that} \\
& \quad (tr, \theta[x \mapsto u]) \models \varphi
\end{aligned}$$

Definition 10 (Validity, satisfiability). *Let $Tr \subseteq (\mathcal{P}(\mathcal{G}))^*$ be a set of traces. A trace formula φ is said to be valid for Tr , written $Tr \models^\forall \varphi$, if for any trace $tr \in Tr$ and any valuation θ we have that $(tr, \theta) \models \varphi$.*

A trace formula φ is said to be satisfiable for Tr , written $Tr \models^\exists \varphi$, if there exist a trace $tr \in Tr$ and a valuation θ such that $(tr, \theta) \models \varphi$.

Note that $Tr \models^\forall \varphi$ iff $Tr \not\models^\exists \neg\varphi$. Given a multiset rewriting system R we say that φ is valid, written $R \models^\forall \varphi$, if $\text{traces}^{mstr}(R) \models^\forall \varphi$. We say that φ is satisfied in R , written $R \models^\exists \varphi$, if $\text{traces}^{mstr}(R) \models^\exists \varphi$. Similarly, given a ground process P we say that φ is valid, written $P \models^\forall \varphi$,

if $\text{traces}^{pi}(P) \models^\forall \varphi$, and that φ is satisfied in P , written $P \models^\exists \varphi$, if $\text{traces}^{pi}(P) \models^\exists \varphi$.

Example. The following trace formula expresses secrecy of keys generated on the security API, which we introduced in Section III.

$$\neg(\exists h, k: \text{msg}, i, j: \text{temp}. \text{NewKey}(h, k)@i \wedge \text{K}(k)@j)$$

VI. A TRANSLATION FROM PROCESSES INTO MULTISSET REWRITE RULES

In this section we define a translation from a process P into a set of multiset rewrite rules $\llbracket P \rrbracket$ and a translation on trace formulas such that $P \models^\forall \varphi$ if and only if $\llbracket P \rrbracket \models^\forall \llbracket \varphi \rrbracket$. Note that the result also holds for satisfiability, as an immediate consequence. For a rather expressive subset of trace formulas (see [21] for the exact definition of the fragment), checking whether $\llbracket P \rrbracket \models^\forall \llbracket \varphi \rrbracket$ can then be discharged to the tamarin prover that we use as a backend.

A. Definition of the translation of processes

To model the adversary’s message deduction capabilities, we introduce the set of rules MD defined in Figure 6.

In order for our translation to be correct, we need to make some assumptions on the set of processes we allow. These assumptions are however, as we will see, rather mild and most of them without loss of generality. First we define a set of reserved variables that will be used in our translation and whose use we therefore forbid in the processes.

Definition 11 (Reserved variables and facts). *The set of reserved variables is defined as the set containing the elements n_a for any $a \in FN$, lock_l for any l .*

The set of reserved facts \mathcal{F}_{res} is defined as the set containing facts $f(t_1, \dots, t_n)$ where $t_1, \dots, t_n \in \mathcal{T}$ and $f \in \{\text{Init}, \text{Insert}, \text{Delete}, \text{IsIn}, \text{IsNotSet}, \text{state}, \text{Lock}, \text{Unlock}, \text{Out}, \text{Fr}, \text{In}, \text{Msg}, \text{ProtoNonce}, \text{Eq}, \text{NotEq}, \text{Event}, \text{InEvent}\}$.

Similar to [2], for our translation to be sound, we require that for each process, there exists an injective mapping assigning to every `unlock` t in a process a `lock` t that precedes it in the process’ syntax tree. Moreover, given a process `lock` t ; P the corresponding `unlock` in P may not be under a parallel or replication. These conditions allow us to annotate each corresponding `lock` t , `unlock` t pair with a unique label l . The annotated version of a process P is denoted \bar{P} . The formal definition of \bar{P} is given in Section A. In case the annotation fails, i.e., P violates one of the above conditions, the process \bar{P} contains \perp .

Definition 12 (well-formed). *A ground process P is well-formed if*

- no reserved variable nor reserved fact appear in P ,
- any name and variable in P is bound at most once and
- \bar{P} does not contain \perp .
- For each action $l \text{--[}a\text{]}\rightarrow r$ that appears in the process, the following holds: for each $l' \text{--[}a'\text{]}\rightarrow r' \in E$

$$\text{ginsts}(l \text{--[}a\text{]}\rightarrow r) \text{ we have that } \cap_{r''=E r'} \text{names}(r'') \cap FN \subseteq \cap_{l''=E l'} \text{names}(l'') \cap FN.$$

A trace formula φ is well-formed if no reserved variable nor reserved fact appear in φ .

The two first restrictions of well-formed processes are not a loss of generality as processes and formulas can be consistently renamed to avoid reserved variables and α -converted to avoid binding names or variables several times. Also note that the second condition is not necessarily preserved during an execution, e.g. when unfolding a replication, $!P$ and P may bind the same names. We only require this condition to hold on the initial process for our translation to be correct.

The annotation of locks restricts the set of protocols we can translate, but allows us to obtain better verification results, since we can predict which `unlock` is “supposed” to close a given `lock`. This additional information is helpful for tamarin’s backward reasoning. We think that our locking mechanism captures all practical use cases for locking. However, using our calculus’ “low-level” multiset manipulation construct, the user is free to implement locks himself, e.g., as

$$[\text{NotLocked}()] \rightarrow []; \text{code}; [] \rightarrow [\text{NotLocked}()]$$

(In this case the user does not benefit from the optimisation we put into the translation of locks.) Obviously, locks can be modelled both in tamarin’s multiset rewriting calculus (this is actually what the translation does) and Mödersheim’s set rewriting calculus [19]. However, protocol steps typically consist of a single input, followed by several database lookups, and finally an output. In practice, they tend to be modelled as a single rule, and are therefore atomic. Real implementations are however different, as several entities might be involved, database lookups could be slow, etc. In this case, such simplified models could, e.g., miss race conditions. To the best of our knowledge, StatVerif is the only comparable tool that models locks explicitly and it has stronger restrictions.

Definition 13. *Given a well-formed ground process P we define the labelled multiset rewriting system $\llbracket P \rrbracket$ as*

$$\text{MD} \cup \{\text{INIT}\} \cup \{\bar{P}, [], []\}$$

- where the rule INIT is defined as

$$\text{INIT} : [] \text{--[}Init()\text{]}\rightarrow [\text{state}_\square()]$$

- $\llbracket P, p, \tilde{x} \rrbracket$ is defined inductively for process P , position $p \in \mathbb{N}^*$ and sequence of variables \tilde{x} in Figure 7.
- For a position p of P we define state_p to be persistent if $P|_p = !Q$ for some process Q ; otherwise state_p is linear.

In the definition of $\llbracket P, p, \tilde{x} \rrbracket$ we intuitively use the family of facts state_p to indicate that the process is currently at position p in its syntax tree. A fact state_p will indeed be

$\text{Out}(x)$	$-\lfloor \rfloor \rightarrow$	$!K(x)$	(MDOU)
$!K(x)$	$-\lfloor K(x) \rfloor \rightarrow$	$\text{In}(x)$	(MDIN)
	$-\lfloor \rfloor \rightarrow$	$!K(x : \text{pub})$	(MDPUB)
$\text{Fr}(x : \text{fresh})$	$-\lfloor \rfloor \rightarrow$	$!K(x : \text{fresh})$	(MDFRESH)
$!K(x_1), \dots, !K(x_k)$	$-\lfloor \rfloor \rightarrow$	$!K(f(x_1, \dots, x_k))$	(MDAPPL)
		for $f \in \Sigma^k$	

Figure 6. The set of rules MD.

true in an execution of these rules whenever some instance of P_p (i.e. the process defined by the subtree at position p of the syntax tree of P) is in the multiset \mathcal{P} of the process configuration. Note in particular that the translation of $!P$ results in a persistent fact as $!P$ always remains in \mathcal{P} . Some of the labels are used to restrict the set of executions. For instance the label $\text{Eq}(M, N)$ will be used to indicate that we only consider executions in which $M =_E N$. As we will see these restrictions will be encoded in the trace formula.

Example. Figure 8 illustrates the above translation by presenting the set of msr rules $\llbracket !P_{new} \rrbracket$ (omitting the rules in MD already shown in Figure 6).

A graph representation of an example trace, generated by the tamarin tool, is depicted in Figure 9. Every box in this picture stands for the application of a multiset rewriting rule, where the premises are at the top, the conclusions at the bottom, and the actions (if any) in the middle. Every premise needs to have a matching conclusion, visualized by the arrows, to ensure the graph depicts a valid msr execution. (This is a simplification of the dependency graph representation tamarin uses to perform backward-induction [21], [22].) Note that the machine notation for $\text{state}_p()$ predicates omits brackets $\lfloor \rfloor$ in the position p and denotes the empty sequence by '0'. We also note that in the current example $!\text{state}_{[1]}()$ is persistent and can therefore be used multiple times as a premise. As $\text{Fr}()$ facts are generated by the FRESH rule which has an empty premise and action, we omit instances of FRESH and leave those premises, but only those, disconnected.

Remark 1. One may note that, while for all other operators, the translation produces well-formed multiset rewriting rules (as long as the process is well-formed itself), this is not the case for the translation of the *lookup* operator, i.e., it violates the well-formedness condition from Definition 4. The constraint solving algorithm employed by tamarin requires all rules, with the exception of FRESH, to be well-formed. We show however that, under these specific conditions, the solution procedure is still correct. See Appendix C in the full version for the proof.

B. Definition of the translation of trace formulas

We can now define the translation for formulas.

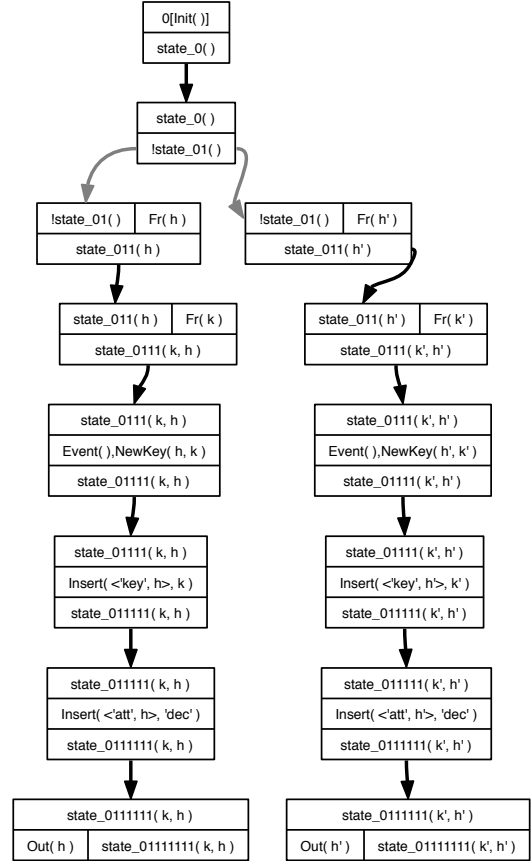


Figure 9. Example trace for the translation of $!P_{new}$.

Definition 14. Given a well-formed trace formula φ we define

$$\llbracket \varphi \rrbracket_{\forall} := \alpha \Rightarrow \varphi \quad \text{and} \quad \llbracket \varphi \rrbracket_{\exists} := \alpha \wedge \varphi$$

where α is defined in Figure 10.

The hypotheses of $\llbracket \varphi \rrbracket$ use the actions of the generated rules to filter out executions that we wish to discard:

- α_{init} ensures that the init rule is only fired once.
- α_{eq} and α_{noteq} ensure that we only consider traces where all (dis)equalities hold.
- α_{in} and α_{notin} ensure that all successful lookups were preceded by an insert that was neither revoked nor overwritten while all unsuccessful lookups where either

$$\begin{aligned}
\llbracket 0, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}) \rightarrow []\} \\
\llbracket P \mid Q, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}) \rightarrow [\text{state}_{p.1}(\tilde{x}), \text{state}_{p.2}(\tilde{x})]\} \\
&\quad \cup \llbracket P, p \cdot 1, \tilde{x} \rrbracket \cup \llbracket Q, p \cdot 2, \tilde{x} \rrbracket \\
\llbracket !P, p, \tilde{x} \rrbracket &= \{[\text{!state}_p(\tilde{x}) \rightarrow [\text{state}_{p.1}(\tilde{x})]] \cup \llbracket P, p \cdot 1, \tilde{x} \rrbracket\} \\
\llbracket \nu a; P, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}), \text{Fr}(n_a : \text{fresh}) \rightarrow [\text{ProtoNonce}(n_a : \text{fresh})] \rightarrow \\
&\quad [\text{state}_{p.1}(\tilde{x}, n_a : \text{fresh})]\} \cup \llbracket P, p \cdot 1, (\tilde{x}, n_a : \text{fresh}) \rrbracket \\
\llbracket \text{Out}(M, N); P, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}), \text{In}(M) \rightarrow [\text{InEvent}(M)] \rightarrow [\text{Out}(N), \text{state}_{p.1}(\tilde{x})], \\
&\quad \text{state}_p(\tilde{x}) \rightarrow [\text{Msg}(M, N), \text{state}_p^{\text{semi}}(\tilde{x})], \\
&\quad \text{state}_p^{\text{semi}}(\tilde{x}), \text{Ack}(M, N) \rightarrow [\text{state}_{p.1}(\tilde{x})]\} \cup \llbracket P, p \cdot 1, \tilde{x} \rrbracket \\
\llbracket \text{In}(M, N); P, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}), \text{In}(\langle M, N \rangle) \rightarrow [\text{InEvent}(\langle M, N \rangle)] \rightarrow \\
&\quad [\text{state}_{p.1}(\tilde{x} \cup \text{vars}(N))], [\text{state}_p(\tilde{x}), \text{Msg}(M, N)] \rightarrow \\
&\quad [\text{state}_{p.1}(\tilde{x} \cup \text{vars}(N)), \text{Ack}(M, N)]\} \\
&\quad \cup \llbracket P, p \cdot 1, \tilde{x} \cup \text{vars}(N) \rrbracket \\
\llbracket \text{if } M = N \text{ then } P \text{ else } Q, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}) \rightarrow [\text{Eq}(M, N)] \rightarrow [\text{state}_{p.1}(\tilde{x})], \\
&\quad \text{state}_p(\tilde{x}) \rightarrow [\text{NotEq}(M, N)] \rightarrow [\text{state}_{p.2}(\tilde{x})]\} \\
&\quad \cup \llbracket P, p \cdot 1, \tilde{x} \rrbracket \cup \llbracket Q, p \cdot 2, \tilde{x} \rrbracket \\
\llbracket \text{event } F; P, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}) \rightarrow [\text{Event}(), F] \rightarrow [\text{state}_{p.1}(\tilde{x})]\} \cup \llbracket P, p \cdot 1, \tilde{x} \rrbracket \\
\llbracket \text{insert } s, t; P, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}) \rightarrow [\text{Insert}(s, t)] \rightarrow [\text{state}_{p.1}(\tilde{x})]\} \cup \llbracket P, p \cdot 1, \tilde{x} \rrbracket \\
\llbracket \text{delete } s; P, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}) \rightarrow [\text{Delete}(s)] \rightarrow [\text{state}_{p.1}(\tilde{x})]\} \cup \llbracket P, p \cdot 1, \tilde{x} \rrbracket \\
\llbracket \text{lookup } M \text{ as } v \text{ in } P \text{ else } Q, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}) \rightarrow [\text{IsIn}(M, v)] \rightarrow [\text{state}_{p.1}(\tilde{M}, v)], \\
&\quad \text{state}_p(\tilde{x}) \rightarrow [\text{IsNotSet}(M)] \rightarrow [\text{state}_{p.2}(\tilde{x})]\} \\
&\quad \cup \llbracket P, p \cdot 1, (\tilde{x}, v) \rrbracket \cup \llbracket Q, p \cdot 2, \tilde{x} \rrbracket \\
\llbracket \text{lock}^l s; P, p, \tilde{x} \rrbracket &= \{[\text{Fr}(\text{lock}_l), \text{state}_p(\tilde{x}) \rightarrow [\text{Lock}(\text{lock}_l, s)] \rightarrow [\text{state}_{p.1}(\tilde{x}, \text{lock}_l)]]\} \\
&\quad \cup \llbracket P, p \cdot 1, \tilde{x} \rrbracket \\
\llbracket \text{unlock}^l s; P, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}) \rightarrow [\text{Unlock}(\text{lock}_l, s)] \rightarrow [\text{state}_{p.1}(\tilde{x})]\} \cup \llbracket P, p \cdot 1, \tilde{x} \rrbracket \\
\llbracket [l \rightarrow a] \rightarrow r; P, p, \tilde{x} \rrbracket &= \{\text{state}_p(\tilde{x}), l \rightarrow [\text{Event}(), a] \rightarrow [r, \text{state}_{p.1}(\tilde{x} \cup \text{vars}(l))]\} \\
&\quad \cup \llbracket P, p \cdot 1, \tilde{x} \cup \text{vars}(l) \rrbracket
\end{aligned}$$

Figure 7. Translation of processes: definition of $\llbracket P, p, \tilde{x} \rrbracket$

never inserted, or at one point deleted and never re-inserted.

- α_{lock} checks that between each two matching locks there must be an unlock. Furthermore, between the first of these locks and the corresponding unlock, there is neither a lock nor an unlock.
- α_{inev} ensures that whenever an instance of MDIN is required to generate an In-fact, it is generated as late as possible, i.e., there is no visible event between the action $K(t)$ produced by MDIN, and a rule that requires In(t).

C. Correctness of the translation

The correctness of our translation is stated by the following theorem.

Theorem 1. *Given a well-formed ground process P and a well-formed trace formula φ we have that*

$$\text{traces}^{pi}(P) \models^* \varphi \text{ iff } \text{traces}^{msr}(\llbracket P \rrbracket) \models^* \llbracket \varphi \rrbracket_*$$

where \star is either \forall or \exists .

We here give an overview of the main propositions and lemmas needed to prove Theorem 1. To show the result we need two additional definitions. We first define an operation that allows to restrict a set of traces to those that satisfy the trace formula α as defined in Definition 14.

We also note that $Tr \models^\forall \llbracket \varphi \rrbracket_\forall$ iff $Tr \not\models^\exists \llbracket \neg \varphi \rrbracket_\exists$.

\square	$\neg[\text{Init}()] \rightarrow$	$[\text{state}_{\square}()]$
$[\text{state}_{\square}()]$	$\neg[\] \rightarrow$	$[\text{!state}_{[1]}()]$
$[\text{!state}_{[1]}(), \text{Fr}(h)]$	$\neg[\] \rightarrow$	$[\text{state}_{[11]}(h)]$
$[\text{state}_{[11]}(h), \text{Fr}(k)]$	$\neg[\] \rightarrow$	$[\text{state}_{[111]}(k, h)]$
$[\text{state}_{[111]}(k, h)]$	$\neg[\text{Event}(), \text{NewKey}(h, k)] \rightarrow$	$[\text{state}_{[1111]}(k, h)]$
$[\text{state}_{[1111]}(k, h)]$	$\neg[\text{Insert}(\langle \text{'key'}, h \rangle, k)] \rightarrow$	$[\text{state}_{[11111]}(k, h)]$
$[\text{state}_{[11111]}(k, h)]$	$\neg[\text{Insert}(\langle \text{'att'}, h \rangle, \text{'dec'})] \rightarrow$	$[\text{state}_{[111111]}(k, h)]$
$[\text{state}_{[111111]}(k, h)]$	$\neg[\] \rightarrow$	$[\text{Out}(h), \text{state}_{[1111111]}(k, h)]$

Figure 8. The set of multiset rewrite rules $\llbracket !P_{new} \rrbracket$ (omitting the rules in MD)

$\alpha := \alpha_{init} \wedge \alpha_{eq} \wedge \alpha_{noteq} \wedge \alpha_{in} \wedge \alpha_{notin} \wedge \alpha_{lock} \wedge \alpha_{inev}$ and

$$\begin{aligned}
\alpha_{init} &:= \forall i, j. && \text{Init}()@i \wedge \text{Init}()@j \implies i = j \\
\alpha_{eq} &:= \forall x, y, i. && \text{Eq}(x, y)@i \implies x \approx y \\
\alpha_{noteq} &:= \forall x, y, i. && \text{NotEq}(x, y)@i \implies \neg(x \approx y) \\
\alpha_{in} &:= \forall x, y, t_3. && \text{IsIn}(x, y)@t_3 \implies \exists t_2. \text{Insert}(x, y)@t_2 \wedge t_2 < t_3 \\
&&& \wedge \forall t_1. \text{Delete}(x)@t_1 \implies (t_1 < t_2 \vee t_3 < t_1) \\
&&& \wedge \forall t_1, y. \text{Insert}(x, y)@t_1 \implies (t_1 < t_2 \vee t_1 \doteq t_2 \vee t_3 < t_1) \\
\alpha_{notin} &:= \forall x, y, t_3. && \text{IsNotSet}(x)@t_3 \implies (\forall t_1, y. \text{Insert}(x, y)@t_1 \implies t_3 < t_1) \vee \\
&&& (\exists t_1. \text{Delete}(x)@t_1 \wedge t_1 < t_3 \\
&&& \wedge \forall t_2, y. (\text{Insert}(x, y)@t_2 \wedge t_2 < t_3) \implies t_2 < t_1) \\
\alpha_{lock} &:= \forall x, l, l', i, j. && \text{Lock}(l, x)@i \wedge \text{Lock}(l', x)@j \wedge i < j \\
&&& \implies \exists k. \text{Unlock}(l, x)@k \wedge i < k \wedge k < j \\
&&& \wedge (\forall l', m. \text{Lock}(l', x)@m \implies \neg(i < m \wedge m < k)) \\
&&& \wedge (\forall l', m. \text{Unlock}(l', x)@m \implies \neg(i < m \wedge m < k)) \\
\alpha_{inev} &:= \forall t, i. && \text{InEvent}(t)@i \implies \exists j. \text{K}(t)@j \\
&&& \wedge (\forall k. \text{Event}()@k \implies (k < j \vee i < k)) \\
&&& \wedge (\forall k, t'. \text{K}(t')@k \implies (k < j \vee i < k \vee k \approx j))
\end{aligned}$$

Figure 10. Definition of α .

Definition 15. Let α be the trace formula as defined in Definition 14 and Tr a set of traces. We define

$$\text{filter}(Tr) := \{tr \in Tr \mid \forall \theta. (tr, \theta) \models \alpha\}$$

The following proposition states that if a set of traces satisfies the translated formula then the filtered traces satisfy the original formula.

Proposition 1. Let Tr be a set of traces and φ a trace formula. We have that

$$Tr \models^* \llbracket \varphi \rrbracket_{\star} \text{ iff } \text{filter}(Tr) \models^* \varphi$$

where \star is either \forall or \exists .

The proof (detailed in the full version) follows directly from the definitions. Next we define the *hiding* operation which filters out all reserved facts from a trace.

Definition 16 (hide). Given a trace tr and a set of facts F

we inductively define $\text{hide}(\square) = \square$ and

$$\text{hide}(F \cdot tr) := \begin{cases} \text{hide}(tr) & \text{if } F \subseteq \mathcal{F}_{res} \\ (F \setminus \mathcal{F}_{res}) \cdot \text{hide}(tr) & \text{otherwise} \end{cases}$$

Given a set of traces Tr we define $\text{hide}(Tr) = \{\text{hide}(t) \mid t \in Tr\}$.

As expected well-formed formulas that do not contain reserved facts evaluate the same whether reserved facts are hidden or not.

Proposition 2. Let Tr be a set of traces and φ a well-formed trace formula. We have that

$$Tr \models^* \varphi \text{ iff } \text{hide}(Tr) \models^* \varphi$$

where \star is either \forall or \exists .

We can now state our main lemma which is relating the set of traces of a process P and the set of traces of its translation into multiset rewrite rules.

Lemma 1. *Let P be a well-formed ground process. We have 4 that*

$$\text{traces}^{pi}(P) = \text{hide}(\text{filter}(\text{traces}^{msr}(\llbracket P \rrbracket))).$$

Our main theorem can now be proven by applying 1, Proposition 2 and Proposition 1.

Proof of Theorem 1:

$$\begin{aligned} \text{traces}^{pi}(P) &\models^* \varphi \\ &\Leftrightarrow \text{hide}(\text{filter}(\text{traces}^{msr}(\llbracket P \rrbracket))) \models^* \varphi && \text{(by Lemma 1)} \\ &\Leftrightarrow \text{filter}(\text{traces}^{msr}(\llbracket P \rrbracket)) \models^* \varphi && \text{(by Proposition 2)} \\ &\Leftrightarrow \text{traces}^{msr}(\llbracket P \rrbracket) \models^* \llbracket \varphi \rrbracket_* && \text{(by Proposition 1)} \end{aligned}$$

■

VII. CASE STUDIES

In this section we briefly overview some of the case studies we performed. These case studies include a simple security API similar to PKCS#11 [20], the Yubikey security token, the optimistic contract signing protocol by Garay, Jakobsson and MacKenzie (GJM) [13] and a few other examples discussed in Arapinis et al. [2] and Mödersheim [19]. The results are summarized in Figure 11. For each of the case studies we provide the number of typing lemmas that were needed by the tamarin prover and whether manual guidance of the tool was required. We do not detail all the formal models of the protocols and properties that we studied, and sometimes present slightly simplified versions. All files of our prototype implementation and our case studies are available at <https://sites.google.com/site/protocolswithglobalstate/>.

A. Security API à la PKCS#11

This example illustrates how our modelling might be useful for the analysis of Security APIs in the style of the PKCS#11 standard [20]. We expect studying a complete model of PKCS#11, such as in [11], to be a straightforward extension of this example. In addition to the processes presented in the running example in Section III the actual case study models the following two operations: (i) *encryption*: given a handle and a plain-text, the user can request an encryption under the key the handle points to. (ii) *unwrap* given a ciphertext $\text{enc}(k_2, k_1)$, and a handle h_1 , the user can request the ciphertext to be *unwrapped*, i.e. decrypted, under the key pointed to by h_1 . If decryption is successful the result is stored on the device, and a handle pointing to k_2 is returned. Moreover, contrary to the running example, at creation time keys are assigned the attribute ‘init’, from which they can move to either ‘wrap’, or ‘unwrap’, see the following snippet:

```

1 in(⟨‘set_dec’, h⟩); lock ⟨‘att’, h⟩;
2   lookup ⟨‘att’, h⟩ as a in
3   if a=‘init’ then

```

```

insert ⟨‘att’, h⟩, ‘dec’; unlock ⟨‘att’, h⟩

```

Note that, in contrast to the running example, it is necessary to encapsulate the state changes between lock and unlock. Otherwise an adversary can stop the execution after line 3, set the attribute to ‘wrap’ in a concurrent process and produce a wrapping. After resuming operation at line 4, he can set the key’s attribute to ‘dec’, even though the attribute is set to ‘wrap’. Hence, the attacker is allowed to decrypt the wrapping he has produced and can obtain the key. Such subtleties can produce attacks that our modeling allows to detect. If locking is handled correctly, we show secrecy of keys produced on the device, proving the property introduced in Example 5. If locks are removed the attack described before is found.

B. Yubikey

The Yubikey [24] is a small hardware device designed to authenticate a user against network-based services. Manufactured by Yubico, a Swedish company, the Yubikey itself is a low cost (\$25), thumb-sized USB device. In its typical configuration, it generates one-time passwords based on encryptions of a secret value, a running counter and some random values using a unique AES-128 key contained in the device. The Yubikey authentication server accepts a one-time password only if it decrypts under the correct AES key to a valid secret value containing a counter larger than the last counter accepted. The counter is thus used as a means to prevent replay attacks. To date, over a million Yubikeys have been shipped to more than 30,000 customers including governments, universities and enterprises, e.g. Google, Microsoft, Agfa and Symantec [25].

Besides the counter values used in the one-time password, the Yubikey stores three additional pieces of information: the public id pid that is used to identify the Yubikey, a secret id $secretid$ that is transmitted as part of the one-time password and only known to the server and the Yubikey, as well as the AES key k , which is also shared with the server. The following process $P_{Yubikey}$ models a single Yubikey, as well as its initial configuration, where an entry in the server’s database for the public id pid is created. This entry contains a tuple consisting of the Yubikey’s secret id, AES key, and an initial counter value (‘one’).

```

PYubikey =
ν k; ν pid; ν secretid;
insert ⟨‘Server’, pid⟩, ⟨secretid, k, ‘one’⟩;
insert ⟨‘Yubikey’, pid⟩, ‘one’;
out(pid);
!PPlugin | !PButtonPress

```

Here, the processes $!P_{Plugin}$ and $!P_{ButtonPress}$ model the Yubikey being unplugged and plugged in again (possibly on a different computer), and the emission of the one-time password. We will only discuss $P_{ButtonPress}$ here. When the user presses the button on the Yubikey, the device outputs

Example	Typing Lemmas	Automated Run
Security API à la PKCS#11	1	yes
Yubikey Protocol [16], [24]	3	no
GJM Contract-Signing protocol [2], [13]	0	yes
Mödersheim’s Example (locks/inserts) [19]	0	little interaction*
Mödersheim’s Example (embedded MSRs) [19]	0	yes
Security Device Example [2]	1	yes
Needham-Schroeder-Lowe [18]	1	yes

* (7 manual rule selections)

Figure 11. Case studies.

a one-time-password consisting of a counter tc , the secret id $secretid$ and additional randomness npr encrypted using the AES key k .

```

PButtonPress =
lock pid;
lookup ⟨‘Yubikey’, pid⟩ as tc in
  insert ⟨‘Yubikey’, pid⟩, tc + ‘one’;
  ν nonce; ν npr;
  event YubiPress(pid, secretid, k, tc);
  out(⟨pid, nonce, senc(⟨secretid, tc, npr⟩, k)⟩);
unlock pid

```

The one-time password $senc(\langle secretid, tc, npr \rangle, k)$ can be used to authenticate against a server that shares the same secret key, which we model in the process P_{Server} . The process receives the encrypted one-time password along with the public id pid of a Yubikey and a $nonce$ that is part of the protocol, but is irrelevant for the authentication of the Yubikey on the server.

The server looks up the secret id and the AES key associated to the public id, i. e., to the Yubikey sending the request, as well as the last recorded counter value otc . If the key and secret id used in the request match the values retrieved from the database, then the event $Smaller(otc, tc)$ is logged along with the event $Login(pid, k, tc)$, which marks a successful login of the Yubikey pid with key k for the counter value tc . Afterwards, the old tuple $\langle secretid, k, otc \rangle$ is replaced by $\langle secretid, k, tc \rangle$, to update the latest counter value received.

```

PServer =
! in(⟨pid, nonce, senc(⟨secretid, tc, npr⟩, k)⟩);
lock pid;
lookup ⟨‘Server’, pid⟩ as tuple in
  if fst(tuple) = secretid then
    if fst(snd(tuple)) = k then
      event Smaller(snd(snd(tuple)), tc)
      event Login(pid, k, tc);
      insert ⟨‘Server’, pid⟩, ⟨secretid, k, tc⟩;
unlock pid

```

Note that, in our modelling, the server keeps one lock per public id, which means that it is possible to have several

active instances of the server thread in parallel as long as all requests concern different Yubikeys.

An important part of the modelling of the protocol is to determine whether one counter value is smaller than another. To this end, our modelling employs a feature added to the development version of tamarin as of October 2012, a union operator $\cup^\#$ for multisets of message terms. The operator is denoted with a plus sign (“+”). We model the counter as a multiset only consisting of the symbol “one”. The cardinality of the multiset is the value of the counter. A counter value is considered smaller than another one, if the first multiset is included in the second. A test $a < b$ is included by adding the event $Smaller(a, b)$ and an axiom that requires that a is a subset of b :

$$\alpha_{Smaller} := \forall i : temp, a, b : msg. Smaller(a, b)@i \Rightarrow \exists z : msg. a + z = b$$

We incorporate this axiom into the security properties just like in Definition 14. Intuitively, we are only interested in traces where a is indeed smaller than b .

The process we analyse models a single authentication server (that may run arbitrary many threads) and an arbitrary number of Yubikeys, i. e., $P_{Server} \mid !P_{Yubikey}$. Among other properties, we show by the means of an injective correspondence property that an attacker that controls the network cannot perform replay attacks, and that each successful login was preceded by a user “pressing the button”, formally:

$$\begin{aligned} & \forall pid, k, x, t_2. \\ & Login(pid, k, x)@t_2 \Rightarrow \\ & \exists sid, t_1. YubiPress(pid, sid, k, x)@t_1 \wedge t_1 < t_2 \\ & \wedge \forall t_3. Login(pid, k, x)@t_3 \Rightarrow t_3 = t_2 \end{aligned}$$

Besides injective correspondence, we show the absence of replay attacks and the property that a successful login invalidates previously emitted one-time passwords. All three properties follow more or less directly from a stronger invariant, which itself can be proven in 295 steps. To find these steps, tamarin needs some additional human guidance, which can be provided using the interactive mode. This mode still allows the user to complement his manual efforts

with automated backward search. The example files contain the modelling in our calculus, the complete proof, and the manual part of the proof which can be completed by tamarin without interaction.

Our analysis makes three simplifications: First, in P_{Server} , we use pattern matching instead of decryption as demonstrated in the process P_{dec} we introduced in Section III. Second, we omit the CRC checksum and the time-stamp that are part of the one-time password in the actual protocol, since they do not add to the security of the protocol in the symbolic setting. Third, the Yubikey has actually two counters instead of one, a session counter, and a token counter. We treat the session and token counter on the Yubikey as a single value, which we justify by the fact that the Yubikey either increases the session counter and resets the token counter, or increases only the token counter, thereby implementing a complete lexicographical order on the pair (*session counter*, *token counter*).

A similar analysis has already been performed by Künnemann and Steel, using tamarin’s multiset rewriting calculus [16]. However, model in our new calculus is more fine-grained and we believe more readable. Security-relevant operations like locking and tests on state are written out in detail, resulting in a model that is closer to the real-life operation of such a device. The modeling of the Yubikey takes approximately 38 lines in our calculus, which translates to 49 multiset rewrite rules. The model of [16] contains only four rules, but they are quite complicated, resulting in 23 lines of code. More importantly, the gap between their model and the actual Yubikey protocol is larger – in our calculus, it becomes clear that the server can treat multiple authentication requests in parallel, as long as they do not claim to stem from the same Yubikey. An implementation on the basis of the model from Künnemann and Steel would need to implement a global lock accessible to the authentication server and all Yubikeys. While a server-side global lock might be conceivable (albeit impractical for performance reasons), a real global lock could not be implemented for the Yubikey as deployed.

C. Further Case Studies

We also investigated the case study presented by Mödersheim [19], a key-server example. We encoded two models of this example, one using the insert construct, the other manipulating state using the embedded multiset rewrite rules. For this example the second model turned out to be more natural and more convenient allowing for a direct automated proof without any additional typing lemma.

We furthermore modeled the contract signing protocol by Garay et al. [13] and a simple security device which both served as examples in [2]. In the contract signing protocol a trusted party needs to maintain a database with the current status of all contracts (aborted, resolved, or no decision has been taken). In our calculus the status information is

naturally modelled using our insert and lookup constructs. The use of locks is indispensable to avoid the status to be changed between a lookup and an insert. Arapinis et al. [2] showed the crucial property that the same contract can never be both aborted and resolved. However, due to the fact that StatVerif only allows for a finite number of memory cells, they have shown this property for a single contract and provide a manual proof to lift the result to an unbounded number of contracts. We directly prove this property for an unbounded number of contracts. Finally we also illustrate the tool’s ability to analyze classical security protocols, by analyzing the Needham Schroeder Lowe protocol [18].

VIII. CONCLUSION

We present a process calculus which extends the applied pi calculus with constructs for accessing a global, shared memory together with an encoding of this calculus in labelled msr rules which enables automated verification using the tamarin prover as a backend. Our prototype verification tool, automating this translation, has been successfully used to analyze several case studies. As future work we plan to increase the degree of automation of the tool by automatically generating helping lemmas. To achieve this goal we can exploit the fact that we generate the msr rules, and hence control their form. We also plan to use the tool for more complex case studies including a complete model of PKCS#11 and a study of the TPM 2.0 standard, currently in public review. Finally, we wish to investigate how our constructs for manipulating state can be used to encode loops, needed to model stream protocols such as TESLA.

REFERENCES

- [1] M. Abadi and C. Fournet. Mobile values, new names, and secure communication. In *Proc. 28th ACM Symp. on Principles of Programming Languages (POPL’01)*, pages 104–115. ACM Press, 2001.
- [2] M. Arapinis, E. Ritter, and M. Ryan. Statverif: Verification of stateful processes. In *Proc. 24th IEEE Computer Security Foundations Symposium (CSF’11)*, pages 33–47. IEEE Press, 2011.
- [3] A. Armando, D. A. Basin, Y. Boichut, Y. Chevalier, L. Compagna, J. Cuéllar, P. H. Drielsma, P.-C. Héam, O. Kouchnarenko, J. Mantovani, S. Mödersheim, D. von Oheimb, M. Rusinowitch, J. Santiago, M. Turuani, L. Viganò, and L. Vigneron. The AVISPA tool for the automated validation of internet security protocols and applications. In *Proc. 17th International Conference on Computer Aided Verification (CAV’05)*, LNCS, pages 281–285. Springer, 2005.
- [4] A. Armando, R. Carbone, L. Compagna, J. Cuéllar, and L. T. Abad. Formal analysis of saml 2.0 web browser single sign-on: Breaking the saml-based single sign-on for google apps. In *Proc. 6th ACM Workshop on Formal Methods in Security Engineering (FMSE’08)*, pages 1–10, 2008.

- [5] S. Bistarelli, I. Cervesato, G. Lenzini, and F. Martinelli. Relating multiset rewriting and process algebras for security protocol analysis. *Journal of Computer Security*, 13(1):3–47, 2005.
- [6] B. Blanchet. An Efficient Cryptographic Protocol Verifier Based on Prolog Rules. In *Proc. 14th Computer Security Foundations Workshop (CSFW'01)*, pages 82–96. IEEE Press, 2001.
- [7] M. Bond and R. Anderson. API level attacks on embedded systems. *IEEE Computer Magazine*, pages 67–75, October 2001.
- [8] M. Bortolozzo, M. Centenaro, R. Focardi, and G. Steel. Attacking and fixing PKCS#11 security tokens. In *Proc. 17th ACM Conference on Computer and Communications Security (CCS'10)*, pages 260–269. ACM Press, 2010.
- [9] *CCA Basic Services Reference and Guide*, Oct. 2006. Available online.
- [10] S. Delaune, S. Kremer, M. D. Ryan, and G. Steel. Formal analysis of protocols based on TPM state registers. In *Proc. 24th IEEE Computer Security Foundations Symposium (CSF'11)*, pages 66–82. IEEE Press, 2011.
- [11] S. Delaune, S. Kremer, and G. Steel. Formal analysis of PKCS#11 and proprietary extensions. *Journal of Computer Security*, 18(6):1211–1245, Nov. 2010.
- [12] S. Escobar, C. Meadows, and J. Meseguer. Maude-*npa*: Cryptographic protocol analysis modulo equational properties. In *Foundations of Security Analysis and Design V*, volume 5705 of *LNCS*, pages 1–50. Springer, 2009.
- [13] J. A. Garay, M. Jakobsson, and P. D. MacKenzie. Abuse-free optimistic contract signing. In *Advances in Cryptology—Crypto'99*, volume 1666 of *LNCS*, pages 449–466. Springer, 1999.
- [14] J. D. Guttman. State and progress in strand spaces: Proving fair exchange. *J. Autom. Reasoning*, 48(2):159–195, 2012.
- [15] J. Herzog. Applying protocol analysis to security device interfaces. *IEEE Security & Privacy Magazine*, 4(4):84–87, July-Aug 2006.
- [16] R. Künnemann and G. Steel. YubiSecure? Formal security analysis results for the Yubikey and YubiHSM. In *Proc. 8th Workshop on Security and Trust Management (STM'12)*, volume 7783 of *LNCS*, pages 257–272, 2012.
- [17] D. Longley and S. Rigby. An automatic search for security flaws in key management schemes. *Computers and Security*, 11(1):75–89, March 1992.
- [18] G. Lowe. Breaking and fixing the Needham-Schroeder public-key protocol using FDR. In *Proc. 2nd International Workshop on Tools and Algorithms for Construction and Analysis of Systems (TACAS'96)*, volume 1055 of *LNCS*, pages 147–166. Springer, 1996.
- [19] S. Mödersheim. Abstraction by set-membership: verifying security protocols and web services with databases. In *Proc. 17th ACM Conference on Computer and Communications Security (CCS'10)*, pages 351–360. ACM, 2010.
- [20] RSA Security Inc., v2.20. *PKCS #11: Cryptographic Token Interface Standard.*, June 2004.
- [21] B. Schmidt, S. Meier, C. Cremers, and D. Basin. Automated analysis of Diffie-Hellman protocols and advanced security properties. In *Proc. 25th IEEE Computer Security Foundations Symposium (CSF'12)*, pages 78–94. IEEE Press, 2012.
- [22] B. Schmidt, S. Meier, C. Cremers, and D. Basin. The tamarin prover for the symbolic analysis of security protocols. In *Proc. 25th International Conference on Computer Aided Verification (CAV'13)*, volume 8044 of *LNCS*, pages 696–701. Springer, 2013.
- [23] Trusted Computing Group. TPM Specification version 1.2. Parts 1–3, revision 103. http://www.trustedcomputinggroup.org/resources/tpm_main_specification, 2007.
- [24] Yubico AB, Kungsgatan 37, 111 56 Stockholm Sweden. *The YubiKey Manual - Usage, configuration and introduction of basic concepts (Version 2.2)*, available at: <http://www.yubico.com/documentation>, June 2010.
- [25] Yubico AB. Yubico customer list, 2013. Accessed: Wed 17 Jul 2013 11:40:50 CEST.

APPENDIX

Definition 17 (Process annotation). *Given a ground process P we define the annotated ground process \bar{P} as follows:*

$$\begin{aligned}
\bar{0} &:= 0 \\
\overline{P|Q} &:= \bar{P}|\bar{Q} \\
\overline{!P} &:= !\bar{P} \\
\overline{\text{if } t_1 = t_2 \text{ then } P \text{ else } Q} &:= \text{if } t_1 = t_2 \text{ then } \bar{P} \text{ else } \bar{Q} \\
\overline{\text{lookup } M \text{ as } x \text{ in } P \text{ else } Q} &:= \text{lookup } M \text{ as } x \text{ in } \bar{P} \text{ else } \bar{Q} \\
\overline{\alpha; P} &:= \alpha; \bar{P} \\
&\text{where } \alpha \notin \{\text{lock } t, \text{unlock } t : t \in \mathcal{T}\} \\
\overline{\text{lock } t; P} &:= \text{lock}^l t; \overline{au(P, t, l)} \\
&\text{where } l \text{ is a fresh label} \\
\overline{\text{unlock}^l t; P} &:= \text{unlock}^l t; \bar{P} \\
\overline{\text{unlock } t; P} &:= \perp
\end{aligned}$$

where $au(P, t, l)$ annotates the first unlock that has parameter t with the label l , i. e.:

$$\begin{aligned}
au(P|Q, t, l) &:= \perp \\
au(!P, t, l) &:= \perp \\
au(\text{if } t_1 = t_2 \text{ then } P \text{ else } Q, t, l) &:= \text{if } t_1 = t_2 \text{ then } au(P, t, l) \text{ else } au(Q, t, l) \\
au(\text{lookup } M \text{ as } x \text{ in } P \text{ else } Q, t, l) &:= \text{lookup } M \text{ as } x \text{ in } au(P, t, l) \text{ else } au(Q, t, l) \\
au(\alpha; P, t, l) &:= \alpha; au(P, t, l) \\
&\text{where } \alpha \neq \text{unlock } t \\
au(\text{unlock } t; P, t, l) &:= \text{unlock}^l t; P \\
au(0, t, l) &:= 0
\end{aligned}$$