

Set- π : Set Membership π -calculus

Alessandro Bruni, Sebastian Mödersheim, Flemming Nielson, Hanne Riis Nielson
 DTU Compute, Technical University of Denmark
 Emails: {albr,samo,fnie,hrni}@dtu.dk

Revised version as of July 16, 2018

Thanks to the comparison by Cheval et al. [14] we have discovered a flaw in our example key-registration protocol, presented in Section VIII-A, whereby an attacker can obtain a registered secret key. Unfortunately in our original query we required that the key is registered *and* still in use by the owner. This property holds, but does not characterize the security of the protocol, which we corrected with an updated query at the end of the section. Furthermore, we introduce a correct translation of the AIF model of Zeb’s key server from the original AIF paper [21] in Section VIII-B, which is in fact secure and can be verified by the SetPi tool.

All the changes to the original published version of the paper are highlighted using a change bar, like this paragraph.

We would like to thank Vincent Cheval and Veronique Cortier for the fruitful interaction regarding fixing these issues, and the updated models along with updates to the tool are available at <http://www2.imm.dtu.dk/~samo/aifom.html>.

Abstract—Communication protocols often rely on stateful mechanisms to ensure certain security properties. For example, counters and timestamps can be used to ensure authentication, or the security of communication can depend on whether a particular key is registered to a server or it has been revoked. ProVerif, like other state of the art tools for protocol analysis, achieves good performance by converting a formal protocol specification into a set of Horn clauses, that represent a monotonically growing set of facts that a Dolev-Yao attacker can derive from the system. Since this set of facts is not state-dependent, the category of protocols of our interest cannot be precisely analysed by such tools, as they would report false attacks due to the over-approximation.

In this paper we present Set- π , an extension of the Applied π -calculus that includes primitives for handling databases of objects, and propose a translation from Set- π into Horn clauses that employs the set-membership abstraction to capture the non-monotonicity of the state. Furthermore, we give a characterisation of authentication properties in terms of the set properties in the language, and prove the correctness of our approach. Finally we showcase our method with three examples, a simple authentication protocol based on counters, a key registration protocol, and a model of the Yubikey security device.

I. INTRODUCTION

The automated verification of security protocols has been the subject of intensive study for about two decades now. This has resulted in methods and tools that are feasible for finding attacks or proving the absence of attacks for a large class

of protocols. One of the most successful approaches is static analysis, as for instance used in the ProVerif tool [6], [8], [25], [16], [10]. The key idea of this approach is to avoid the exploration of the state space of a transition system, but rather compute an over-approximation of the set of messages that the intruder can ever learn. The abstraction is efficient because it avoids the common state-explosion of model checking and it does not require a limitation to finite state-spaces. While this works fine for many protocols, we get trivial “attacks” if a protocol relies on a notion of state that is not local to a single session. A simple example is the protocol:

$$A \rightarrow B : \{Msg, Counter\}_{Key}$$

where *Key* is a symmetric key known only to *A* and *B*, *Msg* is some payload message and *Counter* is the current value of a counter used for avoiding replay attacks: *B* accepts a message only if *Counter* is strictly greater than in the last accepted message from *A*. This protocol thus ensures *injective agreement* [19] on *Msg*, since *B* can be sure that *A* has sent *Msg* and it is not a replay, i.e., even if *A* chooses to transmit several times the same Payload *Msg*, *B* will not accept it more often than *A* sent it. There are of course several ways to model such a counter in the applied π calculus, the input language of ProVerif, however none is going to work in the abstraction due to its *monotonicity*: roughly speaking, whatever *B* accepts once, he will accept any number of times and we thus get trivial attacks. In fact, verifying injective agreement properties in ProVerif requires a dedicated mechanism [9].

The above message is taken from the CANAuth protocol [17] that is intended for the automotive industry and needs to work under strong limitations on bandwidth and time. Due to these constraints, standard mechanisms like challenge-response (*B* first sends a nonce, then *A* includes it in the message instead of the counter) are no option. But even without such bounds there are practical real-world examples that today’s abstraction approaches cannot support:

- Key update/revocation: after updating an old key with a fresh one, one does not accept messages encrypted with the old key anymore (at least after some grace period).
- Key tokens/hardware security modules: they maintain a set of keys of different status and attribute, and can be communicated with through an API. When changing the status of a key, an operation may no longer be possible with that key.
- Data bases: an online shop that maintains a database of orders along with their current status; a customer may cancel an order, but only as long as it has not entered the status “shipped”.

More generally, systems that have a notion of state (that is not local to a session) and that have a *non-monotonic* behavior—i.e. an action is possible until a certain change of state and that is disabled afterwards are incompatible with the abstraction of tools like ProVerif.

Contributions. In this paper we formally define the novel Set- π calculus that extends the popular applied π calculus [1] by a notion of sets of messages. It allows us to declaratively specify how processes can store, lookup and manipulate information like sets of keys, orders, or simply counters as in the above example. (Note that this does not increase the expressive power of applied π , since one could also simulate sets using private channels.) The semantics gives rise to an infinite-state transition system since we can model unbounded processes that generate any number of fresh messages. We can define state-based queries for Set- π , that ask for attacker-derivable messages, their set membership status, and boolean combinations thereof. A specification is secure iff no query is satisfied in any reachable state. Note that we do not specify a particular attacker, but more generally prove that the protocol is secure in the presence of an *arbitrary* attacker A that can be specified as a Set- π calculus (without access to restricted names and sets).

The second contribution is a stateful abstraction for Set- π . The idea is that the abstraction of a message incorporates the information to which sets it belongs, and we model how this set membership can change. In doing so, we integrate the essential part of the state information into an otherwise stateless abstraction. This fine balance allows us to combine the benefits of stateless abstraction—namely avoiding state explosion and bounds to finite state spaces—and at the same time support a large class of protocols that rely on some state aspects.

Formally, this abstraction is a translation from a Set- π protocol specification and a set of queries into a set of first-order Horn clauses. Our third contribution is to prove a soundness result for this abstraction: every reachable state is abstractly represented by the Horn clauses. In particular, if the Horn clauses have no model, then the given Set- π specification is secure for the given queries and against an arbitrary Set- π -attacker. For checking whether the Horn clauses have a model, we can use various automated tools like ProVerif.

Finally, we demonstrate the practical feasibility of our approach by two major case studies on the MaCAN and CANAuth which are candidates for the next generation automotive systems; we use in this paper only excerpts as illustrating examples, the full analysis is found in [12].

a) Plan of the Paper: The rest of the paper proceeds as follows: Section II introduces the language and presents, as a running example, a simplified version of CANAuth—a protocol for securing in-vehicle communication—a protocol that can be more precisely analysed with Set- π . Section III describes the type system. Section IV gives an instrumented semantics for the language. Section V presents two definitions for weak and strong authentication and provides a mechanised way to encode such properties in Set- π . Section VI presents our translation of Set- π into Horn clauses and Section VII proves the correctness of this approach. In Section VIII we

$$\begin{aligned}
M, N &::= x \mid a^l[M_1, \dots, M_n] \mid f(M_1, \dots, M_n) \\
P, Q &::= 0 \mid !^k P \mid P_1 \mid P_2 \\
&\mid \text{out}(M, N); P \mid \text{in}(M, x : T); P \mid \text{new}^l x : a; P \\
&\mid \text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q \\
&\mid \text{if } b \text{ then } P \text{ else } Q \mid \text{update}(b^+); P \\
&\mid \text{lock}(L); P \mid \text{unlock}(L); P \\
b &::= b_1 \wedge b_2 \mid b_1 \vee b_2 \mid \neg b \mid M \in s \\
b^+ &::= b_1^+; b_2^+ \mid M \in s \mid M \notin s \\
Sys &::= \text{new } s : \text{set } T; Sys \\
&\mid \text{reduc } \forall \vec{x} : \vec{T} . g(M_1, \dots, M_n) \rightarrow M; Sys \mid P
\end{aligned}$$

Fig. 1. The process calculus

present the key registration and the Yubikey examples along with experimental results, and finally in Section IX we discuss related work.

II. SET-PI

The calculus is presented in Figure 1. The syntactic categories are terms M, N , processes P, Q , set-membership expressions b , set-membership transitions b^+ , and system declarations Sys . Types in the language, denoted by T and T^{Sym} , are introduced in the next section.

We mark with \mathbb{P} the set of processes produced by the syntactic category P and with \mathbb{M} the set of terms produced by M . To avoid ambiguity, we mark with $\mathbb{S} = \{s_1, \dots, s_n\}$ the sets declared in a specific Set- π model, while we use s, s', s_1, s_2 and so on to denote any of the sets in \mathbb{S} , and L denotes a collection of locked sets.

Terms are either variables, names or constructor applications. Names are annotated with a sequence of terms that record which copy of the process created them, by the use of session identifiers, and input terms to keep track of data dependencies. Constructors are generally accompanied by destructors defined as rewrite rules that describe cryptographic primitives. For example:

$$\text{reduc } \forall x_m : t_m, x_k : \text{key} . \text{dec}(\text{enc}(x_m, x_k), x_k) \rightarrow x_m;$$

models symmetric key encryption: for every message x_m , key x_k , if a process knows an encrypted message $\text{enc}(x_m, x_k)$ and the key x_k then it can obtain the message x_m . For any rewrite rule of the form $\text{reduc } \forall \vec{x} : \vec{T} . g(M_1, \dots, M_n) \rightarrow M; Sys$ we require that $\text{fv}(M) \subseteq \text{fv}(M_1, \dots, M_n) \subseteq \{\vec{x}\}$.

Processes P, Q are: the stuck process 0, replication – which is marked by a label k , parallel composition of two processes, output, typed input, restriction – marked by a label l – and destructor application. We require that processes are closed and that they are properly alpha-renamed. Note that the user does not specify annotated names and labels in the initial process, hence the grey color. Names are introduced by the semantic step for restriction, and unique labels are automatically inserted by the parser.

The distinguishing feature of our calculus is the ability to track values in databases: the membership test (if b then P else Q) allows us to check a membership condition b , while a set transition ($\text{update}(b^+); P$) inserts and removes terms from sets, according to b^+ . Finally we use locks on sets to ensure linearity of set transitions: the construct $\text{lock}(L)$ prevents all other processes to modify sets in L in the continuation, while $\text{unlock}(L)$ releases the locks on L .

A system Sys is the context in which the process operates. A new set s containing elements of type T is declared with new $s : \text{set } T$, and the reduc construct specifies a rewrite rule.

As syntactic sugar we add the following features to $\text{Set-}\pi$:

- n -tuples $\langle M_1, \dots, M_n \rangle$, which can be encoded with a constructor $\text{mktpl}_n(M_1, \dots, M_n)$ and n destructors $\text{reduc } \forall \vec{x} : \vec{T} . \text{proj}_n^i(\text{mktpl}_n(x_1, \dots, x_n)) \rightarrow x_i$;
- pattern matching on tuples for let bindings and inputs, which can be encoded using multiple let bindings with the destructors proj_n^i and equality tests;
- $!^{\{s_1, \dots, s_n\}} P$ means the replication of P that locks sets s_1, \dots, s_n before its execution; the semantics releases the locks when P reduces to 0;
- let variable assignments when the right hand side is not a defined destructor, denoting by $M \notin s$ the expression $\neg(M \in s)$, and omitted else branches where not needed.

a) *CANAuth example*: As a running example we use CANAuth, a protocol that runs on top of resource limited CAN bus networks and—due to the real-time requirements of CAN bus networks—uses one way communication from source to destination, avoiding challenge-response patterns. The low level mechanism that is used to ensure freshness properties is the use of counters together with message authentication codes. Comparing a counter with the highest value previously received allows to ensure that a message cannot be replayed.

Here we model a simplified form of its message authentication procedure, that assumes that the two communicating parties, a sender Alice and a receiver Bob, have established a session key k and are both keeping track of their own local copy of a counter c .

In order for Alice to send a message m to Bob, she signs m and her own counter c increased by one with the shared key k . Here we denote with $\text{hmac}(\text{msg}(c), k)$ such a signature. Bob receives the message and checks whether the counter c is already in the set received ; if not, it accepts the message.

$$A \triangleq \text{new}^1 c : \text{cnt};$$

let $m = \text{msg}(c)$ in
event $\text{send}(m)$;
out($ch, \langle m, \text{hmac}(m, k) \rangle$); 0

$$B \triangleq \text{in}(ch, \langle x_m, x_s \rangle : \langle \text{msg}(\text{cnt}), \text{hmac}(\text{msg}(\text{cnt}), key) \rangle);$$

let $x_c = \text{getcnt}(x_m)$ in
let $_ = \text{checksign}(x_m, x_s, k)$ in
if $x_c \notin \text{received}$ then
 update($x_c \in \text{received}$);
 event $\text{accept}(x_m)$; 0

$T ::= t \mid a \mid f(T_1, \dots, T_n)$	algebraic data types
$T^{Sym} ::= T$	data types
$(T_1, \dots, T_n) \rightarrow T$	destructor type
set T	set type

Fig. 2. Type system

$$S \triangleq \text{new } \text{received} : \text{set } \text{cnt};$$

reduc $\forall x : \text{cnt} . \text{getcnt}(\text{msg}(x)) \rightarrow x$;
reduc $\forall x : t, k : \text{key} . \text{checksign}(x, \text{hmac}(x, k), k) \rightarrow x$;
new² $k : \text{key}$;
(!³ $A \mid$!⁴ $\{\text{received}\} B$)

In order to express authentication we insert two events: *send* and *accept*. These are just syntactic sugar for set operations and, as we show in Section V, they can be translated into set operations.

III. TYPE SYSTEM

The type system presented in Figure 2 is constructed to track the membership of values in sets. We denote by Sym the set of symbols for terms, destructors and sets that occur in a process; the category of types for Sym is T^{Sym} .

Data types T are either type variables, name types or constructors over types. Name types (a) are atomic types like key or cnt , type variables (t) are used to make destructors polymorphic, and constructor types define the shape of a constructor. For example the term $\text{pk}(\text{skey}^l[_])$ has type $\text{pkey} \triangleq \text{pk}(\text{skey})$, and a possible type for the constructor $\text{enc}(x_m, x_k)$ could be $\text{enc}(\text{pair}(\text{id}, \text{pkey}), \text{key})$, if we give the type $\text{pair}(\text{id}, \text{pkey})$ to x_m and key to x_k .

Destructor types are of the form $(T_1, \dots, T_n) \rightarrow T$ where we require $\text{fv}(T) \subseteq \text{fv}(T_1, \dots, T_n)$. A destructor can therefore be applied to different types of data in the process, as long as the typing judgment instantiates a ground type when it is applied. For example the destructor: $\text{reduc } \forall x_m : t, x_k : \text{key} . \text{dec}(\text{enc}(x_m, x_k), x_k) \rightarrow x_m$; has type $(\text{enc}(t, \text{key}), \text{key}) \rightarrow t$, while the instantiation $\text{dec}(\text{enc}(x_m, x_k), x_k)$ has type $(\text{enc}(\text{pair}(\text{id}, \text{pkey}), \text{key})) \rightarrow \text{pair}(\text{id}, \text{pkey})$ considering the previous type assignment.

Set types specify the type of terms contained in sets. For example a set of type $\text{set } \text{pk}(\text{skey})$ contains public keys.

The typing rules (Figure 3) enforce the correct typing of processes. Γ is the type environment, a map from identifiers of terms, destructors and sets to their type.

The typing rules for terms check whether the environment contains the right types for variables, and build types accordingly for the constructors. The rules for systems create the type environment required for typing destructors and set operations in processes. The rule for destructors applies the substitution $\sigma = \{T_i/x_i\}$ to the terms M_1, \dots, M_n, M in order to obtain the type of the destructor, while the rule for sets simply adds the type to the environment.

Processes are typed recursively on their syntactic form. The rules for the stuck process, parallel, replication and output

Terms:

$$\frac{}{\Gamma \vdash x : T} \quad \Gamma(x) = T \quad \frac{}{\Gamma \vdash a[M_1, \dots, M_n] : a} \quad \frac{\Gamma \vdash M_1 : T_1 \quad \dots \quad \Gamma \vdash M_n : T_n}{\Gamma \vdash f(M_1, \dots, M_n) : f(T_1, \dots, T_n)}$$

Conditions:

$$\frac{L, \Gamma \vdash M : T}{L, \Gamma \vdash M \in s} \quad \Gamma(s) = \text{set } T, s \in L \quad \frac{L, \Gamma \vdash b_1 \quad L, \Gamma \vdash b_2}{L, \Gamma \vdash b_1 \wedge b_2} \quad \frac{L, \Gamma \vdash b_1 \quad L, \Gamma \vdash b_2}{L, \Gamma \vdash b_1 \vee b_2} \quad \frac{L, \Gamma \vdash b_1^+ \quad L, \Gamma \vdash b_2^+}{L, \Gamma \vdash b_1^+; b_2^+} \quad \frac{L, \Gamma \vdash b}{L, \Gamma \vdash \neg b}$$

Systems:

$$\frac{\Gamma[g \mapsto (\sigma M_1, \dots, \sigma M_n) \rightarrow \sigma M] \vdash Sys}{\Gamma \vdash \text{reduc } \forall \vec{x} : \vec{T} . g(M_1, \dots, M_n) \rightarrow M; Sys} \quad \sigma = \{T_i/x_i\} \quad \frac{\Gamma[s \mapsto \text{set } T] \vdash Sys}{\Gamma \vdash \text{new } s : \text{set } T; Sys} \quad \frac{\emptyset, \Gamma \vdash P}{\Gamma \vdash P}$$

Processes:

$$\frac{}{\emptyset, \Gamma \vdash 0} \quad \frac{\emptyset, \Gamma \vdash P_1 \quad \emptyset, \Gamma \vdash P_2}{\emptyset, \Gamma \vdash P_1 | P_2} \quad \frac{\emptyset, \Gamma \vdash P}{\emptyset, \Gamma \vdash !P} \quad \frac{L, \Gamma \vdash P}{L, \Gamma \vdash \text{out}(M, N); P} \quad \frac{L, \Gamma[x \mapsto T] \vdash P}{L, \Gamma \vdash \text{in}(M, x : T); P}$$

$$\frac{L, \Gamma[x \mapsto a] \vdash P}{L, \Gamma \vdash \text{new}^l x : a; P} \quad \frac{L, \Gamma \vdash b \quad L, \Gamma \vdash P \quad L, \Gamma \vdash Q}{L, \Gamma \vdash \text{if } b \text{ then } P \text{ else } Q} \quad \frac{L, \Gamma \vdash b \quad L, \Gamma \vdash P}{L, \Gamma \vdash \text{update}(b^+); P} \quad \frac{L \cup L', \Gamma \vdash P}{L, \Gamma \vdash \text{lock}(L'); P} \quad L' \cap L = \emptyset$$

$$\frac{\Gamma \vdash M_i : \sigma T_i \quad L, \Gamma[x \mapsto \sigma T] \vdash P \quad L, \Gamma \vdash Q}{L, \Gamma \vdash \text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q} \quad \Gamma(g) = (T_1, \dots, T_n) \rightarrow T \quad \frac{L \setminus L', \Gamma \vdash P}{L, \Gamma \vdash \text{unlock}(L'); P} \quad L' \subseteq L$$

Fig. 3. Typing rules for terms, rewrite rules, processes and boolean expressions.

simply try to type the continuation under the same Γ . The rules for input and restriction add to Γ the type of the new bound variables. The rules for if and update check that the membership test or the set transitions are well-formed (i.e. M is of type T when s is of type set T) and that the interested sets are locked. Finally the rule for let infers the type of the result of the destructor application given the types of M_1, \dots, M_n , by finding a type substitution that allows typing all arguments of the destructor and by applying such substitution to the result type.

We allow destructor definitions to contain type variables, while we require processes and sets to have only terms of ground types. In the Section IV we introduce a formal semantics for the language, together with the necessary subject reduction results for the type system.

IV. SEMANTICS

We define in Figure 4 an instrumented operational semantics for the language. We have transitions of the form $\rho, S, \mathcal{P} \rightarrow \rho', S', \mathcal{P}'$ where:

- $\rho : \text{Var} \rightarrow \mathbb{M}$, an environment mapping process variables (x, y, \dots) to name instances $(a^l[\dots])$,
- $S \subseteq \mathbb{S} \times \mathbb{M}$ records the set-membership states,
- $\mathcal{P} \subseteq \mathbb{P} \times \wp(\mathbb{S}) \times \wp(\mathbb{M})$ is a multiset of concurrent processes, which are represented as triplets (P, L, V) where P is a process, L is the set of locks held by P , and V is a list of terms that influenced the process (either session identifiers or inputs).

A configuration ρ, S, \mathcal{P} represents the parallel execution of all processes in \mathcal{P} :

$$|(P_i, L_i, V_i) \in \mathcal{P} \rho(P_i)$$

In the semantic rules we assume Γ to contain the type definitions for sets, constructors and destructors, and the initial process to be well-typed according to Γ .

The concrete semantics presented here is a synchronous semantics, which we choose for simplicity and in accordance with the previous related work on ProVerif [9].

The rule NIL removes the process 0 when it holds no locks. The rule COM matches an input and an output processes if the output has the type required by the input; note that the set V_1 of influencing terms for the input process is increased with the term N' constructed from type T using the function pt_x^V . The purpose of pt_x^V is to substitute any type T with a term N' that is homomorphic to T : that is, for every occurrence of a name a in the type T it produces a variable $x_{a,V}$ that is syntactically different from all other variable occurrences, and every occurrence of a constructor type produces a constructor term of the same form.

The rule PAR splits the process into two parallel processes. Replication REPL is annotated with $k \in \mathbb{N}$ and produces a fresh copy of P , adding x_k to V and the substitution $\{k/x_k\}$ to the environment ρ ; the replication process is annotated with the index $k + 1$ after the transition.

The rule for restriction NEW maps x to $a^l[V]$ in the continuation of the process, where l is a unique label for the process $\text{new}^l x : a; P$. Here we extend the terms of Figure 1 to annotate names with a list of variables V under square brackets.

The rules for let reduce the process to P_1 where x is substituted with the result of the rewrite rule in case of success, and to P_2 otherwise.

To that end, we define the relation \rightarrow_ρ as follows. Let s be a term that has only variables of atomic types and such that

$\rho, S, \mathcal{P} \uplus \{(0, \emptyset, V)\} \rightarrow \rho, S, \mathcal{P}$	NIL
$\rho, S, \mathcal{P} \uplus \{(\text{in}(M, x : T); P_1, L_1, V_1), (\text{out}(M, N); P_2, L_2, V_2)\} \rightarrow$ $\text{mgu}(N', N) \circ \rho, S, \mathcal{P} \uplus \{(P_1\{N'/x\}, L_1, N' :: V_1), (P_2, L_2, V_2)\}$ where $\Gamma \vdash N : T$ and $N' = pt_x^{ri(V_1)}(T)$	COM
$\rho, S, \mathcal{P} \uplus \{(P_1 \mid P_2, \emptyset, V)\} \rightarrow \rho, S, \mathcal{P} \uplus \{(P_1, \emptyset, V), (P_2, \emptyset, V)\}$	PAR
$\rho, S, \mathcal{P} \uplus \{(!^k P, \emptyset, V)\} \rightarrow \rho\{^k/x_k\}, S, \mathcal{P} \uplus \{(P, \emptyset, x_k :: V), (!^{k+1} P, \emptyset, V)\}$	REPL
$\rho, S, \mathcal{P} \uplus \{(\text{new}^l x : a; P, L, V)\} \rightarrow \rho, S, \mathcal{P} \uplus \{(P\{a^l[V]/x\}, L, V)\}$	NEW
$\rho, S, \mathcal{P} \uplus \{(\text{let } x = g(M_1, \dots, M_n) \text{ in } P_1 \text{ else } P_2, L, V)\} \rightarrow \rho, S, \mathcal{P} \uplus \{(P_1\{M/x\}, L, V)\}$ if $g(M_1, \dots, M_n) \rightarrow_\rho M$	LET1
$\rho, S, \mathcal{P} \uplus \{(\text{let } x = g(M_1, \dots, M_n) \text{ in } P_1 \text{ else } P_2, L, V)\} \rightarrow \rho, S, \mathcal{P} \uplus \{(P_2, L, V)\}$ if $g(M_1, \dots, M_n) \not\rightarrow_\rho$	LET2
$\rho, S, \mathcal{P} \uplus \{(\text{if } b \text{ then } P_1 \text{ else } P_2, L, V)\} \rightarrow \rho, S, \mathcal{P} \uplus \{(P_1, L, V)\}$ if $\rho, S \models b$	IF1
$\rho, S, \mathcal{P} \uplus \{(\text{if } b \text{ then } P_1 \text{ else } P_2, L, V)\} \rightarrow \rho, S, \mathcal{P} \uplus \{(P_2, L, V)\}$ if $\rho, S \not\models b$	IF2
$\rho, S, \mathcal{P} \uplus \{(\text{update}(b^+); P, L, V)\} \rightarrow \rho, \text{update}(S, \rho(b^+)), \mathcal{P} \uplus \{(P, L, V)\}$	SET
$\rho, S, \mathcal{P} \uplus \{(\text{lock}(L'); P, L, V)\} \rightarrow \rho, S, \mathcal{P} \uplus \{(P, L \cup L', V)\}$ if $\forall (P'', L'', V'') \in \mathcal{P} . L' \cap L'' = \emptyset$	LCK
$\rho, S, \mathcal{P} \uplus \{(\text{unlock}(L'); P, L, V)\} \rightarrow \rho, S, \mathcal{P} \uplus \{(P, L \setminus L', V)\}$ if $L' \subseteq L$	ULCK

$\rho, S \models b_1 \wedge b_2$ iff $\rho, S \models b_1$ and $\rho, S \models b_2$

$\rho, S \models b_1 \vee b_2$ iff $\rho, S \models b_1$ or $\rho, S \models b_2$

$\rho, S \models \neg b$ iff $\rho, S \not\models b$

$\rho, S \models M \in s_i$ iff $\rho(M) \in S(s_i)$

$pt_x^V(a) = x_{a,V}$

$pt_x^V(f(T_1, \dots, T_n)) = f(pt_x^{1::V}(T_1), \dots, pt_x^{n::V}(T_n))$

$ri(V)$ denotes the set of variables x_k in V produced by replication.

$\text{update}(S, M \in s) = S \cup \{(s, M)\}$

$\text{update}(S, M \notin s) = S \setminus \{(s, M)\}$

$\text{update}(S, b_1^+; b_2^+) = \text{update}(\text{update}(S, b_1^+), b_2^+)$

Fig. 4. Semantics for the process algebra

$\rho(s)$ is ground. Then $s \rightarrow_\rho t$ holds iff for some reduction rule $\text{reduc } \forall \vec{x} : \vec{T} . l \rightarrow r$, there is a σ such that:

- σ is the most general unifier of l and s ; w.l.o.g. we can assume that $\text{fv}(\text{Img}(\sigma)) \subseteq \text{Dom}(\rho)$;
- $t = \sigma(r)$. (Note that $\rho(t)$ is ground.)

Otherwise (if no such σ exists), we write $s \not\rightarrow_\rho$.

The rules for $\text{if } b \text{ then } P_1 \text{ else } P_2$ execute P_1 in case the set-membership state S satisfies the boolean formula b , and P_2 otherwise. The rule for update updates the current state according to the expression b^+ . Finally, lock and unlock respectively acquire and release the locks on the sets in L' for the current process.

Having presented the semantic for $\text{Set-}\pi$, we need to prove that our typing judgments are preserved over the transition relation. Hence we introduce Lemma 1 to then prove subject reduction (Theorem 1).

We define the mapping $\widehat{a^l[V]} = a$ that recovers the type from an instrumented name, and its extension to environments $\widehat{\rho}(x) = \rho(x)$.

Lemma 1 (Type substitution). *Let P be a process, Γ and Γ' two type environments, M a term and T a type. If $x \notin \text{Dom}(\Gamma')$ and $L, \Gamma[x \mapsto T]\Gamma' \vdash P$ and $\Gamma\Gamma' \vdash M : T$ then $L, \Gamma\Gamma' \vdash P\{M/x\}$.*

Proof sketch. The proof is carried out by induction on the shape of P , and its sub-terms and boolean conditions.

In particular, when x is encountered in P , we know that:

$$\frac{(\Gamma[x \mapsto T]\Gamma')(x) = T}{\Gamma[x \mapsto T]\Gamma' \vdash x : T}$$

is applied for the proof of $L, \Gamma[x \mapsto T]\Gamma' \vdash P$. Since $x\{M/x\} = M$, our statement directly follows from the hypothesis. \square

Lemma 2 (Environment extension). *Let ρ and ρ' be two environments such that $\text{Dom}(\rho) \subseteq \text{Dom}(\rho')$ and for all x in $\text{Dom}(\rho)$ we have $\rho(x) = \rho'(x)$, let P be a process, Γ a type environment. If $L, \Gamma[\widehat{\rho}] \vdash P$, then $L, \Gamma[\widehat{\rho'}] \vdash P$.*

Proof. By induction on the shape of P . \square

Theorem 1 (Subject reduction). *Let Γ be a type environment, ρ, S, \mathcal{P} a configuration. If for all $(P, L, V) \in \mathcal{P}$ we have $L, \Gamma[\widehat{\rho}] \vdash P$, and if $\rho, S, \mathcal{P} \rightarrow \rho', S', \mathcal{P}'$, then for all $(P', L', V') \in \mathcal{P}'$ we have $L', \Gamma[\widehat{\rho'}] \vdash P'$.*

Proof sketch. The statement can be proven by a case-by-case analysis of the semantic step $\rho, S, \mathcal{P} \rightarrow \rho', S', \mathcal{P}'$. A detailed version of this proof is available in the appendix. \square

Subject reduction enforces that well-typed processes remain well-typed over transitions, and in particular that values of the right type are inserted and removed from sets, and that if and update only occur when the concerned sets are locked.

a) *Attacker processes*: Because sets represent private information of the protocol, our attacker model does not have access to sets. We define it as follows:

Definition 1 (Attacker process). *An attacker process is a well-typed process that shares a channel x_{ch} with the honest protocol, and cannot perform set operations (if, update, lock and unlock constructs are excluded).*

V. AUTHENTICATION IN SET-PI

A general definition of authentication goals for security protocols that has become standard in formal verification are Lowe's notions of non-injective and injective agreement [19]. These notions are hard to combine with an abstract interpretation approach as they inherently include a form of negation that is incompatible with over-approximation. For this reason, ProVerif has a special notion of *events* that are handled in a special way by its resolution procedure. We now show that in our Set- π calculus we can directly express both non-injective and injective agreement using sets, and we can thus define *events* practically as syntactic sugar.

The definition of authentication is based on events $e_1(M)$ and $e_2(M)$ where message M typically contains the (claimed) sender and (intended) receiver name, as well as the data that the participants want to agree on; the event $e_1(M)$ is issued by the sender typically at the beginning of the session and $e_2(M)$ by the receiver at the end of the session, but the precise content and placement can be chosen by the modeler. One can then define non-injective agreement as follows: whenever an event $e_2(M)$ happens for a message M , then previously the event $e_1(M)$ must have happened. Thus, it is an attack if somebody accepts a message that has not been sent that way. The injective agreement additionally requires that if $e_2(M)$ has occurred n times, then $e_1(M)$ must have previously occurred at least n times (i.e., there is an injective mapping from e_2 events to e_1 events). Roughly speaking, it is an attack if a message is accepted more often than it was actually sent. It is hard to automatically verify injective agreement in this formulation. To simplify matters, it is common to require that the authenticated message M includes something fresh, i.e., a unique identifier that the sender chooses [24], [9]. Thanks to this construction, the same e_1 event cannot occur more than once. Then, the injective agreement goal boils down to checking that no e_2 event occurs twice (and that non-injective agreement holds).

Let us thus extend Set- π with event declarations:

$$Sys ::= \dots \mid \text{new } e : \text{event}(T); Sys$$

and event processes:

$$P, Q ::= \dots \mid \text{event } e(M); P$$

and introduce the semantic rule (EVT):

$$\rho, S, \mathcal{P} \uplus \{(\text{event } e(M); P, L, V)\} \xrightarrow{e_1(M)} \rho, S, \mathcal{P} \uplus \{(P, L, V)\}$$

With this extension of the language, we can reason about non-injective and injective agreement properties according to Lowe's definitions. We then encode processes in the extended

calculus with events into processes in standard Set- π and show how our encoding simulates the events.

Definition 2 (Non-injective agreement). *There is a non-injective agreement between event $e_1(M)$ and event $e_2(M)$ if and only if, for every possible trace $\rho, S_0, \mathcal{P}_0 \rightarrow \rho, S_1, \mathcal{P}_1 \rightarrow \dots \rightarrow \rho, S_n, \mathcal{P}_n$ produced by the protocol, if $\rho_i, S_i, \mathcal{P}_i \xrightarrow{e_2(M)} \rho_{i+1}, S_{i+1}, \mathcal{P}_{i+1}$ occurs in the trace, then also $\rho_j, S_j, \mathcal{P}_j \xrightarrow{e_1(M)} \rho_{j+1}, S_{j+1}, \mathcal{P}_{j+1}$ occurs, for $j < i$.*

We construct a transformation from the extended language with events into the language without events, then prove the equivalence between Definition 2 in the original process (extended with events) and a set-property of the transformed process.

The transformation is as follows:

$$\begin{aligned} \text{new } e : \text{event}(T); Sys &\rightarrow \text{new } e : \text{set } T; Sys \\ \text{event } e(M); P &\rightarrow \text{lock}(e); \text{update}(M \in e); \\ &\quad \text{unlock}(e); P \end{aligned}$$

Every event declaration becomes a set declaration in the translated process (assuming that the names for sets and events are disjoint). Whenever an event $e(M); P$ occurs, where M is of type T , we substitute it with the process that locks e , inserts M in the set e , unlocks e and continues with P ; we also add a set declaration for e in its scope. Furthermore, to gain precision in the analysis we merge a set transition followed by an event into a single operation. That is, if we have a process $\text{update}(b^+); \text{event } e(M); P$, we transform it into the process $\text{lock}(e); \text{update}(b^+; M \in e); \text{unlock}(e); P$.

Note that this transformation is sound, although two semantic steps are merged into one: for the purpose of finding violations to an agreement property where $e(M)$ should happen before $e'(M)$, if there is a trace where event $e'(M)$ happens between $\text{update}(b^+)$ and event $e(M)$, then there is also a trace where event $e'(M)$ happens before the set operation. Given a process P we denote its event-free encoding as $\text{agree}(P)$.

Theorem 2. *Let P be an extended process with events. If there is no reachable state S from $P' = \text{agree}(P)$ that satisfies the expression $M \in e_2 \wedge M \notin e_1$, then there is a non-injective agreement between $e_1(M)$ and $e_2(M)$ in P .*

Proof sketch. To prove the correctness of our transformation we construct a simulation relation between P and P' , where the semantic step of an event is simulated by our construction. The full proof is found in the appendix. \square

Definition 3 (Injective agreement). *There is an injective agreement between event $e_1(M)$ and event $e_2(M)$ if and only if, for every possible trace $\rho, S_0, \mathcal{P}_0 \rightarrow \rho, S_1, \mathcal{P}_1 \rightarrow \dots \rightarrow \rho, S_n, \mathcal{P}_n$ produced by the protocol, if $\rho_i, S_i, \mathcal{P}_i \xrightarrow{e_2(M)} \rho_{i+1}, S_{i+1}, \mathcal{P}_{i+1}$ occurs in the trace, then also $\rho_j, S_j, \mathcal{P}_j \xrightarrow{e_1(M)} \rho_{j+1}, S_{j+1}, \mathcal{P}_{j+1}$ occurs, for some $j < i$; furthermore, there does not exist $k > i$ such that $\rho_k, S_k, \mathcal{P}_k \xrightarrow{e_2(M)} \rho_{k+1}, S_{k+1}, \mathcal{P}_{k+1}$.*

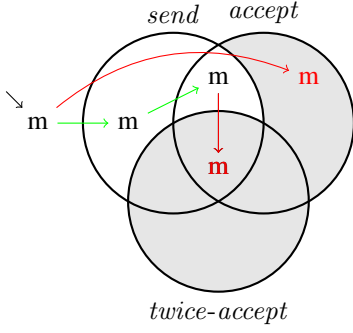


Fig. 5. State transitions of messages in the CANAuth example

For proving injective agreement properties the transformation becomes:

$$\begin{aligned}
 \text{new } e : \text{event}(T); Sys &\rightarrow \text{new } e : \text{set } T; \\
 &\quad \text{new } \textit{twice-e} : \text{set } T; Sys \\
 \text{event } e(M); P &\rightarrow \text{lock}(e, \textit{twice-e}); \\
 &\quad \text{if } M \notin e \text{ then} \\
 &\quad \quad \text{update}(M \in e); \\
 &\quad \quad \text{unlock}(e, \textit{twice-e}); P \\
 &\quad \text{else} \\
 &\quad \quad \text{update}(M \in \textit{twice-e}); \\
 &\quad \quad \text{unlock}(e, \textit{twice-e}); P
 \end{aligned}$$

Every event declaration becomes a pair of set declarations for e and $\textit{twice-e}$ in the translated process (assuming that the names for sets and events are disjoint). Whenever an event $e(M); P$ occurs, where M is of type T , we substitute it with the process that locks e and $\textit{twice-e}$, and performs $\text{update}(M \in e); P$ when M has not yet been inserted in e ; when it is already present in e it performs $\text{update}(M \in \textit{twice-e}); P$, and in both cases unlocks e and $\textit{twice-e}$; finally we add set declarations for e and $\textit{twice-e}$ in the scope. Similarly to the non-injective case, we merge a set operation with the event that follows. Given a process P we denote its event-free encoding as $\textit{inj-agree}(P)$.

Theorem 3. *Let P be an extended process with events. If no reachable state S from $P' = \textit{inj-agree}(P)$ satisfies the expression $(M \in e_2 \wedge M \notin e_1) \vee (M \in \textit{twice-e}_2)$, then the injective agreement between $e_1(M)$ and $e_2(M)$ holds in P .*

Proof sketch. To prove the correctness of our transformation we construct a simulation relation between P and P' , where the semantic step of an event is simulated by our construction. The full proof is found in the appendix. \square

Relating back to our example, Figure 5 shows in green the desired transitions and in red the undesired ones. Our model satisfies non-injective agreement if no message is being accepted without being previously sent by the honest principal. It satisfies injective agreement if no message is accepted twice.

VI. TRANSLATION

The translation takes a process in $\text{Set-}\pi$ and produces a set of Horn clauses that are then solved by a saturation based resolution engine, like ProVerif or SPASS.

$$\begin{aligned}
 \langle H_1 \wedge \dots \wedge H_n \Rightarrow C \rangle_\alpha &= \langle H_1 \rangle_\alpha \wedge \dots \wedge \langle H_n \rangle_\alpha \Rightarrow \langle C \rangle_\alpha \\
 \langle \mathbf{p}(M_1, \dots, M_n) \rangle_\alpha &= \mathbf{p}(\langle M_1 \rangle_\alpha, \dots, \langle M_n \rangle_\alpha) \\
 \langle f(M_1, \dots, M_n) \rangle_\alpha &= f(\langle M_1 \rangle_\alpha, \dots, \langle M_n \rangle_\alpha) \\
 \langle a^l[V] \rangle_\alpha &= \begin{cases} \text{val}(a^l[V], \alpha(s_1, a^l[V]), \dots, \alpha(s_n, a^l[V])) \\ \quad \text{if } l \in \text{labels}(P_0) \\ \text{val}(a^\top[], \alpha(s_1, a^l[V]), \dots, \alpha(s_n, a^l[V])) \quad \text{otr} \end{cases} \\
 \langle x \rangle_\alpha &= \text{val}(x, \alpha(s_1, x), \dots, \alpha(s_n, x))
 \end{aligned}$$

Fig. 6. Applying the set-abstraction

At the end of the section we show how the translation is carried out for our CANAuth example. We now present the general concepts of the translation at an intuitive level, which we then refine with details later in the section. The translation produces clauses with predicates of the form $\text{msg}(M, N)$ to denote that the system has produced an output of N on channel M , predicates of the form $\text{att}(M)$ to denote that the attacker process knows M , predicates of the form $\text{name}(a)$ to denote that a new name is produced by the protocol, and clauses that conclude $\text{transfer}(\cdot, \cdot)$ to denote set-transitions.

The body of a Horn clause represents the inputs that are required to reach a specific point in the process, while the head of the clause represents the output that is generated. For example:

$$\text{in}(ch, x : a); \text{in}(ch, y : b); \text{out}(ch, f(x, y))$$

produces the clause:

$$\text{msg}(ch, x) \wedge \text{msg}(ch, y) \Rightarrow \text{msg}(ch, f(x, y))$$

Names and variables in the predicates are annotated with a special constructor val that defines their current membership class. For example if we have three sets in our system s_1, s_2, s_3 , the term $\text{val}(a, 1, 0, x_{s_3, a})$ represents a name a in the process algebra in a state where a is in s_1 , it is not in s_2 and its membership to the set s_3 is not constrained, as denoted by the variable $x_{s_3, a}$. By doing so two clauses can be unified only if the terms are in consistent states. For example $\text{val}(a, 1, 0, x_{s_3, a})$ unifies with $\text{val}(a, x_{s_1, a}, 0, 1)$ but not with $\text{val}(a, x_{s_1, a}, 1, 1)$, because the first term represents a name a that is not in s_2 , while the third term represents a in a state where it belongs to s_2 .

Now we look at how the translation is constructed. We use a special function α , which we call the set-abstraction, to record whether a particular term belongs to some sets or not, and introduce the rules of Figure 6 to transform clauses, predicates and terms into annotated ones. The set-abstraction is a function of type:

$$\alpha : (\mathbb{S} \times \mathbb{M}) \rightarrow (\{0, 1\} \cup \{x_{s, M} \mid s \in \mathbb{S}, M \in \mathbb{M}\})$$

where we require:

$$\alpha(s, M) \notin \{0, 1\} \implies \alpha(s, M) = x_{s, M}$$

It takes a process set s and a term M , and returns either the constant 1, to enforce that M is in s , the constant 0,

$$\begin{aligned}
\text{restrict}(\alpha, M \in s) &= \text{if } \alpha(s, M) \neq 0 \text{ then } \{\alpha'\} \text{ else } \emptyset \\
\text{where } \alpha'(s', M') &= \begin{cases} 1 & \text{if } M' = M \wedge s' = s \\ \alpha(s', M') & \text{otherwise} \end{cases} \\
\text{restrict}(\alpha, M \notin s) &= \text{if } \alpha(s, M) \neq 1 \text{ then } \{\alpha'\} \text{ else } \emptyset \\
\text{where } \alpha'(s', M') &= \begin{cases} 0 & \text{if } M' = M \wedge s' = s \\ \alpha(s', M') & \text{otherwise} \end{cases} \\
\text{restrict}(\alpha, b_1 \wedge b_2) &= \bigcup \{\text{restrict}(\alpha', b_2) \mid \alpha' \in \text{restrict}(\alpha, b_1)\} \\
\text{restrict}(\alpha, b_1 \vee b_2) &= \text{restrict}(\alpha, b_1) \cup \text{restrict}(\alpha, b_2) \\
\text{restrict}(\alpha, \neg(b_1 \wedge b_2)) &= \text{restrict}(\alpha, (\neg b_1) \vee (\neg b_2)) \\
\text{restrict}(\alpha, \neg(b_1 \vee b_2)) &= \text{restrict}(\alpha, (\neg b_1) \wedge (\neg b_2)) \\
\text{restrict}(\alpha, \neg\neg b) &= \text{restrict}(\alpha, b) \\
\text{zero}(\alpha, a) &= \alpha' \\
\text{where } \alpha'(s, M) &= \begin{cases} 0 & \text{if } a \text{ occurs in } M \\ \alpha(s, M) & \text{otherwise} \end{cases} \\
\text{relax}(\alpha, L) &= \alpha' \\
\text{where } \alpha'(s, M) &= \begin{cases} \alpha(s, M) & \text{if } s \in L \\ x_{s, M} & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 7. Functions for updating α

to enforce that M is not in s , or the variable $x_{s, M}$, to allow one of the two choices to be picked consistently across the hypotheses.

The function $\langle\!\langle p \rangle\!\rangle_\alpha$ of Figure 6 recursively applies the set-abstraction to the clauses. When it encounters an annotated name $a^l[V]$ in the protocol, it produces a constructor $\text{val}(\dots)$ where the first parameter is the name itself—with no annotations in case of attacker names—and the remaining parameters represent the membership of $a^l[V]$ to the sets s_1, \dots, s_n ; similar clauses are generated for occurring variables. For the purpose of making the analysis feasible, as a well-formedness condition we require set types (of the form set T) to contain only name types and monadic constructors over name types. For example, set Seed , set $\text{pk}(\text{Seed})$ and set $\text{sk}(\text{Seed})$ are acceptable set types, while $\text{key}(\text{Seed}, \text{Nonce})$ is not a monadic constructor, hence set $\text{key}(\text{Seed}, \text{Nonce})$ is not an acceptable set type.

The function $\llbracket P \rrbracket_{HVL\alpha}$ of Figure 8 takes a process P , a set of hypothesis predicates H , that intuitively represent the set of messages required to reach P , a list of influencing terms for the process V , a set of locks L held by the process, and the set-abstraction α , and produces a set of clauses representing the protocol behaviour.

Lastly the set of functions restrict , zero and relax modify the set-abstraction for various constructs of Set- π . The function restrict takes a set-abstraction α and a boolean formula b and produces the set of all consistent abstractions that satisfy b , while zero inserts the constant 0 for fresh names, and relax introduces variables in the image of α for unlocked sets.

Having introduced the auxiliary functions for manipulating the set-abstraction, we now come back to explaining the

translation process. The function $\text{relax}(\alpha, L)$ is applied at each step of the translation, as it inserts variables in the image of α for all sets that are not locked, as they may be changed by other processes.

The translation for 0 produces an empty set of clauses. Replication $!^l P$ translates P with the introduction of a new session variable x_l in the list of influencing variables V . Parallel composition $P_1 \mid P_2$ is translated as the union of the clauses generated by both processes.

Input $\text{in}(M, x : T)$ adds the predicate $\text{msg}(M, N')$ as an hypothesis in H , where N' is a copy of T where every occurrence of a name type is replaced with a unique variable using $\text{pt}_x^{\text{ri}(V)}(T)$; the substitution $\{N'/x\}$ is then applied on the continuation. Output $\text{out}(M, N)$ produces a clause with head $\{\text{msg}(M, N)\}_\alpha$ and with hypotheses $\langle\!\langle H \rangle\!\rangle_\alpha$. The rule for $\text{new}^l x : a$ introduces a restricted name: the value class of $a^l[V]$ is set to 0 for every set, the predicate $\text{name}(a^l[V])$ is introduced both in the hypotheses for analysing the continuation and as a fact that follows the current set of hypotheses H . This ensures that all the set-abstraction variables occurring in the head of a clause are closed under the hypotheses.

The rule for $\text{let } x = g(M_1, \dots, M_n) \text{ in } P_1 \text{ else } P_2$ looks for a substitution σ that successfully unifies a definition of the rewrite rule for the destructor g with the actual parameters M_1, \dots, M_n , and then finds a substitution θ that unifies the terms in the set-abstraction α accordingly to the unification on the process algebra terms; if both substitutions are found then $\sigma(P_1)$ is analysed where x is substituted with the result of the reduction. The rule also includes the clauses generated for P_2 on the updated state α' , with no restriction, which is a standard over-approximation found in similar works, e.g. [7].

The rule for $\text{if } b \text{ then } P_1 \text{ else } P_2$ translates P_1 with all the set-abstractions that satisfy the formula b , and P_2 with all the set-abstractions that satisfy the formula $\neg b$. The rule for $\text{lock}(s)$; translates the continuation by first introducing s in the locked sets L , and applying relax to take into account state changes from other processes before the lock takes place. Similarly, the rule for $\text{unlock}(s)$; translates the continuation by removing s from the set L , and applying relax . The rule for $\text{update}(b^+)$; , for every name and variable occurring in b^+ that we denote by M , creates a clause of the form $\langle\!\langle H \rangle\!\rangle_{\alpha'} \Rightarrow \text{transfer}(\langle\!\langle M \rangle\!\rangle_{\alpha'}, \langle\!\langle M \rangle\!\rangle_{\alpha' \cup \{(s, M)\}})$, in order to mark that whenever M appears in a predicate on state α' , we will also have the same predicate on state $\alpha' \cup \{(s, M)\}$, then proceeds to translate the continuation.

a) *Clauses representing the attacker:* We add the following set of clauses to represent a Dolev-Yao attacker. The attacker can eavesdrop messages form known channels:

$$\text{msg}(x_{ch}, x_{msg}) \wedge \text{att}(x_{ch}) \Rightarrow \text{att}(x_{msg})$$

The attacker can insert known messages into channels:

$$\text{att}(x_{ch}) \wedge \text{att}(x_{msg}) \Rightarrow \text{msg}(x_{ch}, x_{msg})$$

For every n -ary constructor f occurring in the protocol we produce a clause:

$$\text{att}(x_1) \wedge \dots \wedge \text{att}(x_n) \Rightarrow \text{att}(f(x_1, \dots, x_n))$$

$$\begin{aligned}
\llbracket 0 \rrbracket HV L \alpha &= \emptyset \\
\llbracket !^l P \rrbracket HV \emptyset \alpha &= \llbracket P \rrbracket H(x_l :: V) \emptyset (\text{relax}(\alpha, \emptyset)) \\
\llbracket P_1 \mid P_2 \rrbracket HV \emptyset \alpha &= \llbracket P_1 \rrbracket HV \emptyset (\text{relax}(\alpha, \emptyset)) \cup \llbracket P_2 \rrbracket HV \emptyset (\text{relax}(\alpha, \emptyset)) \\
\llbracket \text{in}(M, x : T); P \rrbracket HV L \alpha &= \llbracket P \{N'/x\} \rrbracket (H \wedge \text{msg}(M, N')) (N' :: V) L \alpha' \\
&\quad \text{where } \alpha' = \text{relax}(\alpha, L), N' = \text{pt}_x^{\text{ri}(V)}(T) \\
\llbracket \text{out}(M, N); P \rrbracket HV L \alpha &= \llbracket P \rrbracket HV L (\text{relax}(\alpha, L)) \cup \{ \llbracket H \Rightarrow \text{msg}(M, N) \rrbracket_\alpha \} \\
\llbracket \text{new}^l x : a; !^l P \rrbracket HV L \alpha &= \llbracket P \{a^l[V]/x\} \rrbracket (H \wedge \text{name}(a^l[V])) V L \alpha' \cup \{ \llbracket H \Rightarrow \text{name}(a^l[V]) \rrbracket_{\alpha'} \} \\
&\quad \text{where } \alpha' = \text{zero}(\alpha, a^l[V]) \\
\llbracket \text{let } x = g(M_1, \dots, M_n) \text{ in } P_1 \text{ else } P_2 \rrbracket HV L \alpha &= \\
&\quad \{ \llbracket \sigma(P_1) \rrbracket \sigma(H) \sigma(V) L \alpha'' \mid \text{reduc } \forall \vec{x}' : \vec{T}' . g(M'_1, \dots, M'_n) \rightarrow M'; \text{ is in the scope of let,} \\
&\quad \sigma \text{ is an m.g.u. that satisfies } M_1 \doteq M'_1 \wedge \dots \wedge M_n \doteq M'_n \wedge x \doteq M', \\
&\quad \theta \text{ is an m.g.u. that satisfies } \forall s, N_1, N_2, \sigma(N_1) = \sigma(N_2) \Rightarrow \alpha'(s, N_1) \doteq \alpha'(s, N_2), \\
&\quad \text{and } \alpha'' \text{ satisfies } \forall N . \alpha''(s, \sigma(N)) = \theta(\alpha'(s, N)) \} \\
&\quad \cup \llbracket P_2 \rrbracket HV L \alpha' \\
&\quad \text{where } \alpha' = \text{relax}(\alpha, L) \\
\llbracket \text{if } b \text{ then } P_1 \text{ else } P_2 \rrbracket HV L \alpha &= \{ \llbracket P_1 \rrbracket HV L \alpha' \mid \alpha' \in \text{restrict}(\text{relax}(\alpha, L), b) \} \cup \\
&\quad \{ \llbracket P_2 \rrbracket HV L \alpha' \mid \alpha' \in \text{restrict}(\text{relax}(\alpha, L), \neg b) \} \\
\llbracket \text{lock}(L'); P \rrbracket HV L \alpha &= \llbracket P \rrbracket HV (L \cup L') (\text{relax}(\alpha, L)) \\
\llbracket \text{unlock}(L'); P \rrbracket HV L \alpha &= \llbracket P \rrbracket HV (L \setminus L') (\text{relax}(\alpha, L)) \\
\llbracket \text{update}(b^+); P \rrbracket HV L \alpha &= \{ \llbracket H \rrbracket_{\alpha'} \Rightarrow \text{transfer}(\llbracket M \rrbracket_{\alpha'}, \llbracket M \rrbracket_{\alpha''}) \mid M \in \text{fv}(b^+) \cup \text{fn}(b^+) \} \cup \llbracket P \rrbracket HV L \alpha''' \\
&\quad \text{where } \alpha' = \text{relax}(\alpha, L) \text{ and } \alpha'' = \text{update}(\alpha, b^+) \text{ and } \alpha''' = \text{relax}(\alpha'', L)
\end{aligned}$$

Fig. 8. Translation rules for processes into Horn clauses

For all destructors of the form $g(M_1, \dots, M_n) \rightarrow M$ we produce a clause:

$$\text{att}(M_1) \wedge \dots \wedge \text{att}(M_n) \Rightarrow \text{att}(M)$$

Finally, the attacker knows a name for each name type a in the initial state S_0 :

$$\Rightarrow \llbracket \text{att}(a^\top \square) \rrbracket_{S_0} \Rightarrow \llbracket \text{name}(a^\top \square) \rrbracket_{S_0}$$

as well as the public channel shared with the honest protocol:

$$\Rightarrow \llbracket \text{att}(ch^{lo} \square) \rrbracket_{S_0}$$

b) Clauses representing the set-transitions: We now introduce clauses that express the meaning of the $\text{transfer}(\cdot, \cdot)$ predicate: roughly speaking, when the fixedpoint contains $\text{transfer}(M, M')$, then for every fact $C[M]$ also $C[M']$ holds (for any context $C[\cdot]$). Since there are infinitely many contexts $C[\cdot]$, we cannot directly write this as Horn clauses, but it actually suffices to restrict ourselves to contexts that occur on the right-hand side of a Horn clause. More precisely, let Cl be the set of clauses produced by the translation $\llbracket P_0 \rrbracket \emptyset \emptyset \emptyset \alpha_0$. Let M_a, M'_a be two terms of type a , and consider every clause $H \Rightarrow C[M_a] \in Cl$ for some context $C[\cdot]$ that is not a transfer predicate. Finally let α and α' be two set-abstractions such that $H' \Rightarrow \text{transfer}(\llbracket M'_a \rrbracket_\alpha, \llbracket M'_a \rrbracket_{\alpha'}) \in Cl$, and let $\sigma = \text{mgu}(M_a, M'_a)$. For each such case we add the following Horn clause:

$$\llbracket C[M'_a] \rrbracket_{\alpha \sigma} \wedge \text{transfer}(\llbracket M'_a \rrbracket_\alpha, \llbracket M'_a \rrbracket_{\alpha'}) \Rightarrow \llbracket C[M'_a] \rrbracket_{\alpha' \sigma}$$

This set of clauses transfers messages and attacker knowledge between states. Hence if a predicate is derivable in the saturation in state α and there is a transition from α to α' , then the predicate will be derivable in state α' .

Intuitively this set of rules suffices for the translation because only the honest protocol produces state transitions; everything that the attacker can derive in a state, it can also derive in the successor state. Therefore it is only necessary to transfer the conclusions for the protocol. Lemma 5 establishes the correctness of this approach.

c) Translation of CANAuth into Horn clauses: To show how the translation is applied to produce Horn clauses from the original description, we have taken an excerpt from our running example, namely the receiving process, translated the events into set transitions, and label each point of the program:

$$\begin{aligned}
P_1 &\triangleq !^{l_1} \{r, a, \text{at}\} \text{ } l_2 \text{ in}(ch, \langle x_m, x_s \rangle : \\
&\quad \langle \text{msg}(\text{cnt}), \text{hmac}(\text{msg}(\text{cnt}), \text{key}) \rangle);^{l_3} \\
&\quad \text{let } _ = \text{eq}(x_s, \text{hmac}(x_m, k)) \text{ in }^{l_4} \\
&\quad \text{let } x_c = \text{getc}nt(x_m) \text{ in }^{l_5} \\
&\quad \text{if } x_c \notin r \text{ then }^{l_6} \\
&\quad \quad \text{if } x_m \notin a \text{ then }^{l_7} \text{update}(x_c \in r; x_m \in a);^{l_8} \\
&\quad \quad \text{else }^{l_9} \text{update}(x_c \in r; x_m \in \text{at});^{l_{10}}
\end{aligned}$$

Figure 9 shows the clauses that are generated, together with the recursive calls of $\llbracket \cdot \rrbracket$ that are required to produce

them. We use here the notation like $(H_1 = \emptyset)$ to indicate the development of the parameters H , V , L and α over the recursive calls of $\llbracket \cdot \rrbracket$, and P^l to denote the subprocess of P_1 at label l .

VII. CORRECTNESS

In this section we want to establish the correctness of our translation with respect to the semantics of Section IV. We use the inference system of Figure 10 to express our correctness results. Intuitively, this set of rules relates the instrumented semantics to the Horn clauses generated by the translation, namely that the fixed-point \mathcal{F}_{P_0} covers all possible behaviours of a process, when started in the given configuration (ρ, L, V, S) and in any environment that cannot change sets in L .

Let S and S' again be states of the sets (i.e., $S(s)$ yields the elements that are members of set s in state S); we can view a state as a special case of an abstraction α that has no variables (i.e., indetermined set memberships) and we can thus can write $\langle \cdot \rangle_S$ accordingly. We will now show: if the semantic relation induces a reachable state S at which output N on channel M is produced, then the Horn clauses generated for the protocol entails the ground fact $\langle \text{msg}(M, N) \rangle_S$. This ensures that whatever behaviour is present in the semantics is also captured by the translation.

We denote by \mathcal{C}_{P_0} the set of clauses produced by the translation, including the fixed clauses, and by \mathcal{F}_{P_0} the set of ground facts derivable from \mathcal{C}_{P_0} . First we introduce the order relation \preceq_L on the facts \mathcal{F}_{P_0} derivable from the initial protocol.

Definition 4 (Order relation \preceq_L). *The order relation $S_1 \preceq_L S_2$ between states S_1 and S_2 holds iff:*

- (i) $\forall s_j \in L . S_1(s_j) = S_2(s_j)$;
- (ii) $\forall p(M_1, \dots, M_k) . \langle p(M_1, \dots, M_k) \rangle_{S_1} \in \mathcal{F}_{P_0} \Rightarrow \langle p(M_1, \dots, M_k) \rangle_{S_2} \in \mathcal{F}_{P_0}$.

Intuitively the \preceq_L relation captures the causal relation of the semantic rules, as condition (ii) requires all predicates of the form `msg`, `name` and `att` to be transferred from state S_1 to state S_2 . Furthermore condition (i) imposes that the locked sets L are not modified between the two states. The most general of such relations is \preceq_\emptyset , as it allows any set to be modified.

Next we formalise the definition for set-abstraction α that was introduced in Section VI.

Definition 5 (Set-abstraction). *The mapping α abstracts S under the environment ρ iff for every set s , term M , either $\alpha(s, M) = 1$ and $\rho(M) \in S(s)$, or $\alpha(s, M) = 0$ and $\rho(M) \notin S(s)$, or $\alpha(s, M) = x_{s,M}$.*

A set abstraction α abstracts a state S if every pair (s, M) that maps to a variable in α is mapped to a variable that is unique in the image (this is ensured syntactically by the use of $x_{s,M}$), and whenever $\alpha(s, M)$ maps to the constants 1 and 0 then $\rho(M) \in s$ and $\rho(M) \notin s$, respectively, in the state S .

The following lemmata establish the relation between the operations used in the translation and the order relation \preceq_L . The interested reader can find the full proofs in the extended version of this article .

Lemma 3 (*relax* preserves the set-abstraction over \preceq_L). *Let S, S' be two states such that $S \preceq_L S'$, and assume α abstracts S under ρ . Then $\alpha' = \text{relax}(\alpha, L)$ abstracts S' under ρ .*

Since *relax* inserts unique variables in α' for all sets that are not locked, and for all sets s that are locked $\alpha(s) = \alpha'(s)$ and $S(s) = S'(s)$ holds by condition (i) of the order relation, then α satisfies the properties of Definition 5.

Lemma 4 (*restrict* preserves the set-abstraction). *Let α be a set abstraction, ρ an environment, S a state and $A = \text{restrict}(\alpha, b)$. If $\rho, S \models b$ and α abstracts S , then there exists an $\alpha' \in A$ such that α' abstracts S .*

Restrict produces a set of set-abstractions each representing a possible way of satisfying the formula b . Lemma 4 establishes that if α abstracts S then at least one of these restrictions on α satisfies the abstraction of S .

Lemma 5 (*transfer* preserves $S \preceq_L S'$). *Let S be a set-membership state, and $S' = S \cup \{(s_1, M_1), \dots, (s_j, M_j)\} \setminus \{(s_{j+1}, M_{j+1}), \dots, (s_n, M_n)\}$. If for all $M \in \{M_1, \dots, M_n\}$ we have $\text{transfer}(\langle M \rangle_S, \langle M \rangle_{S'}) \in \mathcal{F}_{P_0}$ then for any set of locks L such that $\{s_1, \dots, s_n\} \cap L = \emptyset$ we have $S \preceq_L S'$.*

Lemma 5 establishes that *transfer* predicates actually capture the state transitions, hence following the definition of the order \preceq_L the set of predicates derivable in the updated state is larger than that derivable in the original state.

Next we type the attacker process A and the honest protocol P_0 , under the initial environment $\rho_0 = [x_{ch} \mapsto ch^{l_0}]$.

Lemma 6 (Typability of A). *Let A be an attacker process, then $\rho_0, \emptyset, \emptyset, S_0 \Vdash A$.*

Proof of sketch. Let B be a subprocess of A , ρ an environment, S a state, V a list of terms. We prove that if:

- (i) $\rho(B)$ is a closed process, $\rho(V)$ is ground,
- (ii) $S_0 \preceq_\emptyset S$, and
- (iii) for every maximal subterm M of B closed under ρ , we have $\langle \rho(\text{att}(M)) \rangle_S \in \mathcal{F}_{P_0}$,

then:

$$\rho, V, \emptyset, S \Vdash B$$

Proof by induction over the depth of B .

In particular, we have that (i) $fv(A) = x_{ch}$, hence $\rho_0(A)$ is closed; (ii) $S_0 \preceq_\emptyset S_0$ by reflexivity; and (iii) the only maximal subterm of A that is bound by ρ_0 is x_{ch} , and by construction of the translation $\langle \rho_0(\text{att}(x_{ch})) \rangle_{S_0} \in \mathcal{F}_{P_0}$. Hence the attacker process types. \square

Lemma 7 (Typability of P_0). $\rho_0, \emptyset, \emptyset, S_0 \Vdash P_0$.

Proof sketch. Let Q be a process. We prove that, given a list of terms V , a set of locks L , a state S , a set-abstraction α , an environment ρ ; if:

- (i) $\rho(Q)$ is a closed process, $\rho(V)$ and $\rho(H)$ are ground,
- (ii) α abstracts S under ρ ,
- (iii) $\mathcal{C}_{P_0} \supseteq \llbracket Q \rrbracket HVL\alpha$,
- (iv) for every predicate p in H , we have that $\langle \rho(p) \rangle_S \in \mathcal{F}_{P_0}$

Then $\rho, V, L, S \Vdash Q$.

$$\begin{aligned}
\llbracket B \rrbracket \emptyset \emptyset \emptyset \alpha_0 &= \llbracket P^{l_1} \rrbracket (H_1 = \emptyset)(V_1 = [x_{l_1}])(L_1 = \emptyset)(\alpha_1 = \text{relax}(\alpha_0, \emptyset)) \\
&= \llbracket P^{l_2} \rrbracket (H_2 = \emptyset)(V_2 = V_1)(L_2 = \{r, a, at\})(\alpha_2 = \text{relax}(\alpha_1, \emptyset)) \\
&= \llbracket P^{l_3} \rrbracket (H_3 = \{\text{msg}(ch[]), T_3 = \langle \text{msg}(x_{cnt,1}), \text{hmac}(\text{msg}(x_{cnt,2}), x_{key,3}) \rangle \}) \\
&\quad (V_3 = T_3 :: V_2)(L_3 = L_2)(\alpha_3 = \text{relax}(\alpha_2, L_2)) \\
&= \llbracket P^{l_4} \rrbracket (H_4 = \{\text{msg}(ch[]), T_4 = \langle \text{msg}(x_{cnt,1}), \text{hmac}(\text{msg}(x_{cnt,1}), k[]) \rangle \}) \\
&\quad (V_4 = T_4 :: V_2)(L_4 = L_3)(\alpha_4 = \theta(\text{relax}(\alpha_3, L_3))) \\
&= \llbracket P^{l_5} \rrbracket (H_5 = H_4)(V_5 = V_4)(L_5 = L_4)(\alpha_5 = \text{relax}(\alpha_4, L_4)) \\
&= \llbracket P^{l_6} \rrbracket (H_6 = H_5)(V_6 = V_5)(L_6 = L_5)(\alpha_6 = \text{restrict}(\text{relax}(\alpha_5, L_5), x_{cnt,1} \notin r)) \\
&= \llbracket P^{l_7} \rrbracket (H_7 = H_6)(V_7 = V_6)(L_7 = L_6)(\alpha_7 = \text{restrict}(\text{relax}(\alpha_6, L_6), \text{msg}(x_{cnt,1}) \notin a)) \cup \\
&\quad \llbracket P^{l_{10}} \rrbracket (H_{10} = H_6)(V_{10} = V_6)(L_{10} = L_6)(\alpha_{10} = \text{restrict}(\text{relax}(\alpha_6, L_6), \text{msg}(x_{cnt,1}) \in a)) \\
\llbracket P^{l_7} \rrbracket H_7 V_7 L_7 \alpha_7 &= \{\text{msg}(\text{msg}(\text{val}(ch[])), \langle \text{msg}(\text{val}(x_1, x_{1,s}, 0, 0, x_{1,at})), \text{hmac}(\text{msg}(\text{val}(x_1, x_{1,s}, 0, 0, x_{1,at}), \text{val}(k[]))) \rangle) \} \\
&\quad \Rightarrow \text{transfer}(\text{val}(x_1, x_{1,s}, 0, 0, x_{1,at}), \text{val}(x_1, x_{1,s}, 1, 1, x_{1,at})) \} \\
\llbracket P^{l_{10}} \rrbracket H_{10} V_{10} L_{10} \alpha_{10} &= \{\text{msg}(\text{msg}(\text{val}(ch[])), \langle \text{msg}(\text{val}(x_1, x_{1,s}, 0, 1, x_{1,at})), \text{hmac}(\text{msg}(\text{val}(x_1, x_{1,s}, 0, 1, x_{1,at}), \text{val}(k[]))) \rangle) \} \\
&\quad \Rightarrow \text{transfer}(\text{val}(x_1, x_{1,s}, 0, 1, x_{1,at}), \text{val}(x_1, x_{1,s}, 1, 1, 1)) \}
\end{aligned}$$

Fig. 9. Horn clauses generated by the translation

The proof is carried out by induction on the structure of the process Q .

In particular, (i) ρ_0 closes P_0 by construction, $\rho_0(\emptyset)$ is trivially ground, (ii) α_0 abstracts S_0 under ρ_0 by construction, (iii) $C_{P_0} \supseteq \llbracket P_0 \rrbracket \emptyset \emptyset \emptyset \alpha_0$ by definition of the translation, (iv) holds vacuously. Therefore the conditions (i–iv) are satisfied and hence $\rho_0, \emptyset, \emptyset, S_0 \Vdash P_0$. \square

Theorem 4 (Subject reduction). *If $\rho, S, \mathcal{P} \rightarrow \rho', S', \mathcal{P}'$ and for all $(P, L, V) \in \mathcal{P}$ we have $\rho, V, L, S \Vdash P$ then for all $(P', L', V') \in \mathcal{P}'$ we have $\rho', V', L', S' \Vdash P'$.*

The proof is a case-by-case analysis on the semantic rules for the language.

Theorem 5 (Correctness of the analysis). *Let $\text{Sys}[\cdot]$ be the system context, let P_0 be the protocol, let T be the set of types used by P_0 , let A be any attacker process using only types in T , and $\rho_0 = [x_{ch} \mapsto ch^{l_0}]$.*

If the typing $\emptyset \vdash \text{Sys}[\text{new}^{l_0} x_{ch} : ch; P_0 \mid A]$ holds, and if $\rho_0, S_0, \mathcal{P}_0 = \{(P_0 \mid A, \emptyset, \emptyset)\} \rightarrow^ \rho_n, S_n, \mathcal{P}_n = \mathcal{P}_n' \uplus \{(\text{out}(M, N); P', L, V)\}$; then $\langle \rho(\text{msg}(M, N)) \rangle_{S_n} \in \mathcal{F}_{P_0}$.*

Proof. By Lemma 7 we know that $\rho_0, \emptyset, \emptyset, S_0 \Vdash P_0$; by Lemma 6 we know that $\rho_0, \emptyset, \emptyset, S_0 \Vdash A$, hence all processes in \mathcal{P}_0 type in the initial state S_0 .

Let $\rho_0, S_0, \mathcal{P}_0 \rightarrow \rho_1, S_1, \mathcal{P}_1 \rightarrow \dots \rightarrow \rho_n, S_n, \mathcal{P}_n$. By inductively applying Theorem 4 on the length of the trace n we can conclude that all processes in \mathcal{P}_n type in the S_n .

In particular, the process $\text{out}(M, N); P'$ types in state S_n and hence $\langle \rho_n(\text{msg}(M, N)) \rangle_{S_n} \in \mathcal{F}_{P_0}$. \square

Theorem 5 establishes the final relation between the inference system of Figure 10 and the instrumented semantics. We use this result to link the facts generated by the translation to a query of interest.

Corollary 1 (Checking queries). *If $\rho_0, S_0, \mathcal{P}_0 = \{(P_0 \mid A, \emptyset, \emptyset)\} \rightarrow^* \rho_j, S_j, \mathcal{P}_j = \mathcal{P}_j' \uplus \{(\text{out}(M, N); P', L, V)\} \rightarrow^* \rho_n, S_n, \mathcal{P}_n$ and $\rho_n, S_n \models b$ then there exists an $\alpha_b \in \text{restrict}(\alpha_0, b)$ where $\theta = \text{mgu}(S_n, \alpha_b)$ and $\langle \rho_n(\text{msg}(M, N)) \rangle_{\theta \circ \alpha_b}$.*

Proof. Follows from Theorem 5 and because $S_j \preceq_{\emptyset} S_n$. \square

Therefore we can express any query of the form:

$$\text{msg}(M, N) \text{ where } b$$

where b is a boolean expression ranging over names and monadic constructors in M and N . Queries of this form are general enough to model secrecy from the attacker's perspective (assuming that the channel M is public), as well as the authentication properties discussed in Section V.

VIII. EXPERIMENTAL EVALUATION

We implemented our analysis into a prototype tool written in Haskell, that translates processes specified in Set- π and uses ProVerif as a back-end resolution engine for Horn clauses. The tool is available for download at [11].

Figure 11 shows the results for our examples: the running example on CANAuth, a flawed version of MaCAN [12], a key registration protocol, and an implementation of the Yubikey protocol modeled after [18].

We recorded the running times for our test suite on a Intel Core i3-7100U at 2.40GHz with 8 GB of RAM running Linux.

They are comparable to similar ProVerif models in applied- π , which shows that there is little overhead induced by our specific translation. In the next two subsections we present the key registration protocol and the Yubikey example.

A. Key Registration (with Attack)

Here we present a key-registration protocol where an honest principal A registers its current pair of asymmetric keys

$$\begin{array}{c}
\frac{}{\rho, V, L, S \Vdash 0} \text{T-NIL} \quad \frac{\forall S' \text{ s.t. } S \preceq_{\emptyset} S' \ (\rho, V, \emptyset, S' \Vdash Q_1 \wedge \rho, V, \emptyset, S' \Vdash Q_2)}{\rho, V, \emptyset, S \Vdash Q_1 \mid Q_2} \text{T-PAR} \\
\frac{\forall S' \text{ s.t. } S \preceq_{\emptyset} S' \ (\rho\{l/x_l\}, (x_l :: V), \emptyset, S' \Vdash Q)}{\rho, V, \emptyset, S \Vdash l^! Q} \quad l \in \mathbb{N} \text{T-REPL} \\
\frac{\forall S' \text{ s.t. } S \preceq_L S' \ \forall N \text{ s.t. } \Gamma \vdash N : T \quad \langle \rho(\text{msg}(M, N)) \rangle_{S'} \in \mathcal{F}_{P_0} \Rightarrow (\text{mgu}(N', N) \circ \rho), (N' :: V), L, S' \Vdash Q\{N'/x\}}{\rho, V, L, S \Vdash \text{in}(M, x : T); Q} \quad N' = p_x^{ri(V)}(T) \text{T-IN} \\
\frac{\langle \rho(\text{msg}(M, N)) \rangle_S \in \mathcal{F}_{P_0} \wedge \forall S' \text{ s.t. } S \preceq_L S' \ (\rho, V, L, S' \Vdash Q)}{\rho, V, L, S \Vdash \text{out}(M, N); Q} \text{T-OUT} \\
\frac{\langle \rho(\text{name}(a^l[V])) \rangle_S \in \mathcal{F}_{P_0} \wedge \forall S' \text{ s.t. } S \preceq_L S' \ \rho, V, L, S' \Vdash Q\{a^l[V]/x\}}{\rho, V, L, S \Vdash \text{new}^l x : a; Q} \text{T-NEW} \\
\frac{\forall S' \text{ s.t. } S \preceq_L S' \ (\forall M \text{ s.t. } g(M_1, \dots, M_n) \rightarrow_{\rho} M \ \rho, V, L, S' \Vdash Q_1\{M/x\}) \wedge \rho, V, L, S' \Vdash Q_2}{\rho, V, L, S \Vdash \text{let } x = g(M_1, \dots, M_n) \text{ in } Q_1 \text{ else } Q_2} \text{T-LET} \\
\frac{\forall S' \text{ s.t. } S \preceq_L S' \ (\rho, S' \models b \Rightarrow \rho, V, L, S' \Vdash Q_1) \wedge (\rho, S' \models \neg b \Rightarrow \rho, V, L, S' \Vdash Q_2)}{\rho, V, L, S \Vdash \text{if } b \text{ then } Q_1 \text{ else } Q_2} \text{T-IF} \\
\frac{\forall S' \text{ s.t. } S \preceq_L S' \ \rho, V, (L \cup L'), S' \Vdash Q}{\rho, V, L, S \Vdash \text{lock}(L'); Q} \text{T-LOCK} \quad \frac{\forall S' \text{ s.t. } S \preceq_L S' \ \rho, V, (L \setminus L'), S' \Vdash Q}{\rho, V, L, S \Vdash \text{unlock}(L'); Q} \text{T-UNLOCK} \\
\frac{\forall S' \text{ s.t. } S \preceq_L S' \ (\forall M \in \text{fv}(b^+) \cup \text{fn}(b^+) \ \text{transfer}(\langle \rho(M) \rangle_{S'}, \langle \rho(M) \rangle_{S''}) \in \mathcal{F}_{P_0}) \wedge (\forall S''' \text{ s.t. } S'' \preceq_L S''', \ \rho, V, L, S''' \Vdash Q)}{\rho, V, L, S \Vdash \text{update}(b^+); Q} \quad S'' = \text{update}(S', \rho(b^+)) \text{T-SET}
\end{array}$$

Fig. 10. Inference system for correctness

Example	Average time	Vulnerable
CANAuth	0.011s	no
MaCAN	0.019s	yes
Key registration	0.043s	yes
Zeb's key server	0.160s	no
Yubikey single	0.014s	no

Fig. 11. Experimental results

(pk_A, sk_A) to the server S . An initial pair of keys is distributed securely to A and S , where A knows both public and secret keys while S only knows the public key.

Later in the protocol, before the current key expires, A registers a new key to the server by sending the following message:

$$A \rightarrow S : \text{senc}(sk, (\text{new}, a, pk'))$$

which encodes the new public key pk' with the old secret key sk . S will be able to decrypt A 's message with the old public key pk , move pk from the set of valid keys to the database of revoked keys and send back an acknowledgment to A .

$$S \rightarrow A : \text{penc}(pk', (\text{confirm}))$$

In turn A will be able to decrypt this message with sk' and remove the old sk from its key-ring.

$$\begin{array}{l}
A \triangleq \text{in}(kdb_a, sk_a : SKey); \\
\text{if } sk_a \in ring_a \text{ then} \\
\text{new } s'_a : Seed; \\
\text{update}(sk(s'_a) \in ring_a); \\
\text{out}(ch, \text{senc}(sk_a, (\text{new_key}, a, pk(s'_a)))); \\
\text{out}(kdb_a, sk(s'_a)); \\
\text{in}(ch, x_c : \text{senc}(PKey, x_t)); \\
\text{let } x_r = \text{pdec}(sk(s'_a), x_c) \text{ in} \\
\text{if } x_r = \text{confirm} \text{ then} \\
\text{update}(sk_a \notin ring_a); \\
\text{out}(ch, sk_a); 0 \\
\\
S \triangleq \text{in}(ch, x_s : \text{senc}(SKey, (x_t, x_t', PKey))); \\
\text{in}(kdb_s, pk_a : PKey); \\
\text{let } (= \text{new}, = a, pk'_a) = \text{pdec}(pk_a, x_s) \text{ in} \\
\text{if } pk_a \in valid_a \wedge pk'_a \notin valid_a \wedge pk'_a \notin revoked_a \text{ then} \\
\text{update}(pk_a \in revoked_a; pk_a \notin valid_a; pk'_a \in valid_a); \\
\text{out}(ch, \text{penc}(pk'_a, \text{confirm})); \\
\text{out}(kdb_s, pk'_a); 0
\end{array}$$

$$\begin{aligned}
Sys \triangleq & \text{reduc } \forall x : Seed, m : t . \\
& \quad pdec(pk(x), senc(sk(x), m)) \rightarrow m; \\
& \text{reduc } \forall x : Seed, m : t . \\
& \quad \quad sdec(sk(x), penc(pk(x), m)) \rightarrow m; \\
& \text{reduc } \forall x : Seed . \text{keypair}(sk(x), pk(x)) \rightarrow true; \\
& \text{new } ring_a : \text{set } sk(Seed); \\
& \text{new } valid_a : \text{set } pk(Seed); \\
& \text{new } revoked_a : \text{set } pk(Seed); \\
& \text{new } kdb_a : SKKey; \text{new } kdb_s : SKKey; \text{new } s_a : Seed; \\
& \text{update}(sk(s_a) \in ring_a; pk(s_a) \in valid_a); \\
& \text{out}(kdb_a, sk(s_a)); \\
& \text{out}(kdb_s, pk(s_a)); \\
& (!\{ring_a\} A \mid !\{valid_a, revoked_a\} S)
\end{aligned}$$

Once a new key is established and the client receives confirmation from the server that the secret key sk has been revoked, sk can be revealed to the attacker. An attacker succeeds in breaking the protocol when she discovers a secret key that is still registered to the server.

We express the property with the following query:

$$\text{att}(sk) \text{ where } pk(sk) \text{ in } valid_a$$

with this query we find the attack that was originally discovered by Cheval et al. [14].

B. Zeb's Key Registration in SetPi

Here we present an adaptation of Zeb's Key Registration protocol from [21] with two clients, a and b , and one server s , each running the protocol. The *Init* process models the out-of-band key registration, the *Client* process models the client requesting a new key, and the *Server* process models the server-side handling of the request. All processes are parameterized on the agent names and sets on which they operate:

$$\begin{aligned}
Init(a, ring, valid) \triangleq & \\
& \text{lock}(ring, valid); \\
& \text{new } sk_1 : skey; \\
& \text{let } pk_1 : pkey = pk(sk_1) \text{ in} \\
& \text{update}(sk_1 \text{ in } ring \wedge pk_1 \text{ in } valid); \\
& \text{out}(d, sk_1,); \\
& \text{out}(ch, pk_1,); \\
& \text{unlock}(ring, valid); .
\end{aligned}$$

$$\begin{aligned}
Client(a, ring, revoked) \triangleq & \\
& \text{lock}(ring); \\
& \text{new } sk_2 : skey; \\
& \text{in}(d, sk_1 : skey); \\
& \text{if } sk_1 \text{ in } ring \text{ then} \\
& \quad \text{let } pk_2 : pkey = pk(sk_2) \text{ in} \\
& \quad \text{update}((sk_2 \text{ in } ring) \wedge (\neg sk_1 \text{ in } ring));
\end{aligned}$$

$$\begin{aligned}
& \text{out}(ch, sign(sk_1, (a, pk_2)),); \\
& \text{unlock}(ring); \\
& \text{lock}(revoked); \\
& \text{let } pk_1 : pkey = pk(sk_1) \text{ in} \\
& \text{if } pk_1 \text{ in } revoked \text{ then} \\
& \quad \text{out}(ch, sk_1,); \\
& \quad \text{unlock}(revoked); .
\end{aligned}$$

$$\begin{aligned}
Server(a, valid, revoked) \triangleq & \\
& \text{lock}(valid, revoked); \\
& \text{in}(ch, pk_1 : pkey); \\
& \text{if } pk_1 \text{ in } valid \text{ then} \\
& \quad \text{in}(ch, x_{msg} : sign(skey, (agent, pkey))); \\
& \quad \text{let } (=a, x_{pk_2} : pkey) = \text{checksign}(pk_1, x_{msg}) \text{ in} \\
& \quad \text{if } \neg x_{pk_2} \text{ in } revoked \text{ then} \\
& \quad \quad \text{update}((\neg pk_1 \text{ in } valid) \wedge (pk_1 \text{ in } revoked) \\
& \quad \quad \quad \wedge (x_{pk_2} \text{ in } valid)); \\
& \quad \text{unlock}(valid, revoked); .
\end{aligned}$$

$$\begin{aligned}
Sys \triangleq & \\
& \text{reduc } \forall x : SKKey, y : t . \\
& \quad \text{checksign}(pk(x), sign(x, y)) = y; \\
& \text{reduc } \forall x : SKKey, y : t . \\
& \quad \text{getmsg}(sign(x, y)) = y; \\
& \text{new } ring_a : \text{set } SKKey; \text{new } ring_b : \text{set } SKKey; \\
& \text{new } valid_{as} : \text{set } pk(SKKey); \\
& \text{new } valid_{bs} : \text{set } pk(SKKey); \\
& \text{new } revoked_{as} : \text{set } pk(SKKey); \\
& \text{new } revoked_{bs} : \text{set } pk(SKKey); \\
& \text{new } a : agent; \text{new } b : agent; \\
& \text{new } d : channel; \\
& (Init(a, ring_a, valid_{as})) \mid \\
& (Init(b, ring_b, valid_{bs})) \mid \\
& (!Client(a, ring_a, revoked_{as})) \mid \\
& (!Client(b, ring_b, revoked_{bs})) \mid \\
& (!Server(a, valid_{as}, revoked_{as})) \mid \\
& (!Server(b, valid_{bs}, revoked_{bs}))
\end{aligned}$$

As in the previous model, we express the security property with the following query:

$$\text{att}(sk) \text{ where } pk(sk) \text{ in } valid$$

for $valid \in \{valid_{as}, valid_{bs}\}$.

C. Yubikey

Yubikey is a small USB token used to authenticate to supported online services. It works by maintaining a pair of a secret identity (shared with the server) and a public identity (shared publicly), and by sending to the Yubikey server its own public identity, together with the one time password encrypted with the current value of a counter using a shared key k .

Here we model a simplified version of the Yubikey protocol, where we are interested in the injective agreement between the client Yubikey (YK) and the server (Srv). The process BP represents the process activated by pressing the Yubikey button, which authenticates the user to the server. We define a public channel ch , and a private channel ch_server that is only used to securely exchange the secret identity and shared key to the server.

The Yubikey process YK creates a new fresh key k , its own public and secret identities (x_{pid} and x_{sid}), stores them securely to the server, then reveals its public identity and starts the BP process.

The button press (BP) process initiates the authentication procedure, increasing the counter (this is encoded in our calculus by the creation of a fresh value), producing the nonces x_{nonce} and x_{tpr} , and sending the encrypted message. An event yk_press is inserted to denote that the button has been pressed.

The server on the other end receives the login request from the Yubikey, retrieves its secret identity and key k from its own channel, pattern matching on the Yubikey's public identity to find the right tuple, decrypts the message with the retrieved key k , and finally if the counter has not been used, it issues a yk_login event to conclude the protocol.

```

Sys  $\triangleq$  new  $yk\_press$  : event(cnt);
      new  $yk\_login$  : event(cnt);
      new  $used$  : set cnt;
      reduc  $\forall x : t, k : key . sdec(senc(x, k), k) \rightarrow x$ ;
      new  $ch$  : channel;
      new  $ch\_server$  : channel;
      (! YK | ! Srv)

BP  $\triangleq$  new  $x_c$  : cnt;
      new  $x_{nonce}$  : nonce;
      new  $x_{tpr}$  : nonce;
      event  $yk\_press(x_c)$ ;
      out(ch,  $\langle x_{pid}, x_{nonce}, senc(\langle x_{sid}, x_c, x_{tpr} \rangle, k) \rangle$ ); 0

YK  $\triangleq$  new  $k$  : key;
      new  $x_{pid}$  : pid;
      new  $x_{sid}$  : sid;
      out(ch_server,  $\langle x_{pid}, x_{sid}, k \rangle$ );
      out(ch,  $x_{pid}$ );
      ! BP

Srv  $\triangleq$  in(ch,  $\langle x_{pid}, x_{nonce}, x_{enc} \rangle$  :
      (pid, nonce, senc( $\langle sid, cnt, nonce \rangle$ , key)));
      in(ch_server,  $\langle = x_{pid}, x_{sid}, x_k \rangle : \langle pid, sid, key \rangle$ );
      let  $\langle = x_{sid}, x_{cnt}, x_{tpr} \rangle = sdec(x_{enc}, x_k)$  in
      lock(used);
      if  $x_{cnt} \notin used$  then
      update( $x_{cnt} \in used$ );
      event  $yk\_login(x_{cnt})$ ;

```

```
unlock(used); 0
```

Here we find an injective agreement between the events yk_press and yk_login . Although this example shows only one Yubikey and Server pair, it can be extended by including multiple copies of the client and server processes, and copies of the respective sets and events to prove the injective agreement with a finite number of participants.

IX. CONCLUSIONS AND RELATED WORK

The Set- π calculus and its set-based abstraction method provide an important step to overcome a serious limitation in current automated protocol verification: the limited support verifying protocols that use state. When a change in state can lead to non-monotonicity (things that were possible before the change are not possible after it) then the standard abstraction and resolution approach leads to false positives. Our solution is to enter just the “right amount” of state information into the abstraction of messages: enough to represent the non-monotonic aspects we want to model, but only so much that we do not destroy the benefit of the stateless abstraction approach in the first place. This work has been inspired by several works that go in a similar direction and we discuss here how they relate to us.

The closest works are two articles that similar to this work add state information into abstraction-based approaches. The AIF framework [21] first presented the idea of encoding set memberships into the state abstraction. AIF is based on the low-level AVISPA Intermediate Format [24] and thus does not have the declarativity of a process calculus. For instance, one has to explicitly specify the attacker and cannot derive it from the calculus. Further, AIF uses the “raw” set membership abstraction, while in our abstraction approach we do integrate the context in which messages have been created which gives a finer abstraction. Also Set- π uses locks on sets, while AIF does not have this notion (and the lock exists only as per the scope of each AIF transition rule).

The second similar stateful abstraction approach is StatVerif [2] which also provides an extension of the applied π calculus. While we use state information in the abstraction of messages, StatVerif encodes state information as an additional argument in the generated predicates of the Horn clauses. The state transition that this approach supports are in some sense like “breaking glass”: we can make at some point a global change which cannot be reverted (to avoid cycles in the state transition graph). We believe that Set- π and StatVerif have some complementary strengths as there are examples that cannot be directly expressed in the other. While StatVerif can express that a set of messages makes a state transition at the same time, our abstraction loses this relation between messages. On the other hand, we can flexibly have messages change their set membership independent of each other, and they can return to any previous state. An argument for the expressiveness of Set- π is that we have a systematic way to formalize agreement properties using sets.

There are several model-checking approaches that can deal with stateful protocols, namely the AVISPA/AVANTSSAR platform [3]. Note that here one needs to bound the number of

steps of honest agents which is often fine for finding attacks, but gives limited guarantees for verification. In fact, [15] studies APIs of key tokens using SATMC [4] of AVISPA, and considers abstractions of data similar to our set abstraction. This can in some cases lead to finitely many reachable abstracted states so the analysis despite depth bound is complete. The AVANTSSAR platform also includes the novel specification language ASLan that besides sets also supports the formulation of Horn-clause policies that are freshly evaluated in every state. For certain fragments we can obtain effective model-checking approaches, but again at the price of bounding the number of steps of honest agents [22].

Another verification approach that supports the verification of stateful protocols is the Tamarin prover [20]. Instead of abstraction techniques, it uses backward search and lemmata that allow to cut of search to cope with the infinite state spaces. Even for quite simple stateful protocols, such lemmata have to be supplied by the user to achieve termination. This is demonstrated by the work of [18] that presents another extension of the applied π -calculus where processes can manipulate a global key map and defines a suitable encoding into Tamarin rules. The benefit of Tamarin and related tools is a great amount of flexibility in formalising relationships between data that cannot be captured by a particular abstraction and resolution approach. However it comes at the price of loosing automation, i.e., that the user has to supply insight into the problem by proving auxiliary lemmata.

Other works have proposed a type-based approach to the verification of stateful protocols. Bugliesi et al. [13] propose a type system with resource-aware authorisation policies based on a variant of affine logic with replication. The type system is constructed on a variant of the Applied π -calculus, and can express access revocation to resources once they are consumed. Swamy et al. [23] propose a variant of ML with value-dependent types called F^* for proving security properties in stateful protocol implementations. It has been used to prove the security of a full implementation of the TLS protocol [5], and can encode the type system of [13]. Like in the case of Tamarin, these approaches are very expressive but not fully automated and often require the user to supply additional lemmata. The most interesting current developments are thus to identify fragments of the problem that can be covered for instance using SMT solvers.

We believe that there is potential for further refining our analysis, in particular it seems possible to integrate ideas from the StatVerif approach, possibly even from the Tamarin-based approach into the set-based abstraction to further enlarge the class of protocols we can support.

ACKNOWLEDGMENTS

The work presented in this paper was partially supported by the EU ARTEMIS Project no. 295354 “SESAMO” (sesamo-project.eu) and by the EU FP7 Projects no. 318424, “FutureID: Shaping the Future of Electronic Identity” (futureid.eu). The authors would like to thank Roberto Vigo for valuable discussion and the anonymous reviewers for helpful comments.

REFERENCES

- [1] Martín Abadi and Cédric Fournet. Mobile values, new names, and secure communication. In *Symposium on Principles of Programming Languages*, 2001.
- [2] Myrto Arapinis, Joshua Phillips, Eike Ritter, and Mark Dermot Ryan. StatVerif: Verification of stateful processes. *Journal of Computer Security*, 22(5):743–821, 2014.
- [3] Alessandro Armando, Wihem Arsac, Tigran Avanesov, Michele Bartoletta, Alberto Calvi, Alessandro Cappai, Roberto Carbone, Yannick Chevalier, Luca Compagna, Jorge Cuéllar, Gabriel Erzse, Simone Frau, Marius Minea, Sebastian Mödersheim, David von Oheimb, Giancarlo Pellegrino, Serena Elisa Ponta, Marco Rocchetto, Michaël Rusinowitch, Mohammad Torabi Dashti, Mathieu Turuani, and Luca Viganò. The AVANTSSAR Platform for the Automated Validation of Trust and Security of Service-Oriented Architectures. In *TACAS*, 2012.
- [4] Alessandro Armando and Luca Compagna. Sat-based model-checking for security protocols analysis. *Int. J. Inf. Sec.*, 7(1):3–32, 2008.
- [5] Karthikeyan Bhargavan, Cédric Fournet, Markulf Kohlweiss, Alfredo Pironti, and Pierre-Yves Strub. Implementing TLS with verified cryptographic security. In *Security and Privacy*, 2013.
- [6] Bruno Blanchet. An efficient cryptographic protocol verifier based on prolog rules. In *Computer Security Foundations Workshop*, 2001.
- [7] Bruno Blanchet. From secrecy to authenticity in security protocols. In *Static Analysis Symposium*, 2002.
- [8] Bruno Blanchet. Security protocols: from linear to classical logic by abstract interpretation. *Inf. Process. Lett.*, 95(5):473–479, 2005.
- [9] Bruno Blanchet. Automatic verification of correspondences for security protocols. *Journal of Computer Security*, 17(4):363–434, 2009.
- [10] Chiara Bodei, Mikael Buchholtz, Pierpaolo Degano, Flemming Nielson, and Hanne Riis Nielson. Static validation of security protocols. *Journal of Computer Security*, 13(3):347–390, 2005.
- [11] Alessandro Bruni. SetPi tool page. <http://www2.compute.dtu.dk/~samo/aifom.html>.
- [12] Alessandro Bruni, Michal Sojka, Flemming Nielson, and Hanne Riis Nielson. Formal security analysis of the MaCAN protocol. In *Integrated Formal Methods*, 2014.
- [13] Michele Bugliesi, Stefano Calzavara, Fabienne Eigner, and Matteo Maffei. Resource-aware authorization policies for statically typed cryptographic protocols. In *Computer Security Foundations Symposium*, 2011.
- [14] Vincent Cheval, Véronique Cortier, and Mathieu Turuani. A little more conversation, a little less action, a lot more satisfaction: Global states in ProVerif. Research report, Inria Nancy - Grand Est ; LORIA, UMR 7503, Université de Lorraine, CNRS, Vandoeuvre-lès-Nancy, April 2018.
- [15] Sibylle B. Fröschle and Graham Steel. Analysing PKCS#11 key management APIs with unbounded fresh data. In *ARSPA-WITS*, 2009.
- [16] Jean Goubault-Larrecq. Towards producing formally checkable security proofs, automatically. In *Computer Security Foundations Symposium*, 2008.
- [17] Anthony Van Herrewege, David Singelee, and Ingrid Verbauwhede. CANAuth-A simple, backward compatible broadcast authentication protocol for CAN bus. In *Proceedings of ECRYPT*, 2011.
- [18] Steve Kremer and Robert Künnemann. Automated analysis of security protocols with global state. In *Security and Privacy*, 2014.
- [19] Gavin Lowe. A hierarchy of authentication specification. In *Computer Security Foundations Workshop*, 1997.
- [20] Simon Meier, Benedikt Schmidt, Cas Cremers, and David A. Basin. The TAMARIN prover for the symbolic analysis of security protocols. In *Computer Aided Verification*, 2013.
- [21] Sebastian Mödersheim. Abstraction by set-membership: verifying security protocols and web services with databases. In *Computer and Communications Security*, 2010.
- [22] Sebastian Mödersheim. Deciding security for a fragment of aslan. In *European Symposium on Research in Computer Security*, 2012.
- [23] Nikhil Swamy, Juan Chen, Cédric Fournet, Pierre-Yves Strub, Karthikeyan Bhargavan, and Jean Yang. Secure distributed programming with value-dependent types. *J. Funct. Program.*, 23(4):402–451, 2013.
- [24] The AVISPA Project. AVISPA Project Deliverable 2.3: The Intermediate Format, 2003. <http://www.avispa-project.org/publications.html>.
- [25] Christoph Weidenbach. Towards an automatic analysis of security protocols in first-order logic. In *Conference on Automated Deduction*, 1999.