



Efficiently Deciding Equivalence for Standard Primitives and Phases

Véronique Cortier¹, Antoine Dallon^{1,2,3(✉)}, and Stéphanie Delaune³

¹ LORIA, CNRS, Nancy, France

dallon@lsv.fr

² LSV, CNRS and ENS Paris-Saclay, Cachan, France

³ Univ Rennes, CNRS, IRISA, Rennes, France

Abstract. Privacy properties like anonymity or untraceability are now well identified, desirable goals of many security protocols. Such properties are typically stated as equivalence properties. However, automatically checking equivalence of protocols often yields efficiency issues.

We propose an efficient algorithm, based on graph planning and SAT-solving. It can decide equivalence for a bounded number of sessions, for protocols with standard cryptographic primitives and phases (often necessary to specify privacy properties), provided protocols are well-typed, that is encrypted messages cannot be confused. The resulting implementation, SAT-Equiv, demonstrates a significant speed-up w.r.t. other existing tools that decide equivalence, covering typically more than 100 sessions. Combined with a previous result, SAT-Equiv can now be used to prove security, for some protocols, for an unbounded number of sessions.

1 Introduction

Security protocols are notoriously difficult to design. A common good practice is to formally analyse protocols using symbolic techniques, in order to spot flaws possibly before their deployment (e.g. TLS 1.3 [4, 20], an avionic protocol [5]). These symbolic techniques are mature for reachability properties like confidentiality or authentication. More recently, this approach has been extended to privacy properties, such as vote secrecy, anonymity, untraceability, or unlinkability. These properties are expressed through equivalences. For example, in the case of biometric passports, an attacker should not be able to distinguish whether she is in contact with Alice's passport or Bob's passport.

Recently, a new tool, SAT-Equiv [16], has been proposed to decide such equivalence properties for security protocols, for a bounded number of sessions.

The research leading to these results has received funding from the European Research Council under the European Union's horizon 2020 research and innovation program (ERC grant agreement n° 714955-POPSTAR and n° 645865-SPOOC), as well as from the French National Research Agency (ANR) under the project TECAP, and the DGA.

It is based on a standard model-checking approach, namely graph planning [7, 23] and SAT-solving. Intuitively, protocols executions are over-approximated as a graph planning problem, which allows to consider several possible interleavings in parallel, allowing the analysis of dozen of sessions of a protocol in a few seconds. However, this result is limited to a very small set of primitives, namely symmetric encryption and concatenation.

Our Contributions. Building upon this novel approach, we enrich SAT-Equiv in order to cover protocols using asymmetric primitives and/or phases. As for the original SAT-Equiv, we assume a non confusion property: encrypted messages should not be confused, a condition automatically checked by our tool and which can be enforced e.g. through appropriate labelling.

First, we extend SAT-Equiv to cover all standard primitives: symmetric and asymmetric encryption, signatures, and hashes. Since graph planning is a bounded model-checking technique, SAT-Equiv relies on a small model property, that bounds the size of messages. More precisely, [12] guarantees that if there is an attack, then there is a well-typed attack, where messages follow a fix format. This result has been recently extended to standard primitives [14]. The straightforward extension of SAT-Equiv to standard primitives however yields severe efficiency issues. Indeed, unlike the symmetric encryption case, checking whether two sequences of messages are equivalent (i.e. in static equivalence) may require complex tests where the attacker *construct* messages (that is, hash or asymmetrically encrypt messages). We therefore provide a precise characterisation of the set of tests that need to be considered when checking for static equivalence. This characterisation is of independent interest and could be used in other contexts. We also extend SAT-Equiv to consider protocols with phases, which are useful to model game-based properties.

Our extension of SAT-Equiv now provably terminates. In [16], termination can be guaranteed by checking that any state of the planning graph is indeed reachable, which requires to query a SAT-solver at each step. While this provides termination in theory, this yields a non practical algorithm and has not been implemented. Instead, we exhibit a bound on the maximal length of the smallest attack (bounding the attacker steps as well). It is therefore sufficient to stop the construction of the graph planning once this bound has been reached, enforcing termination for free (no computation overhead).

Finally, we have considerably revisited and improved the original implementation of SAT-Equiv. This significant speedup now allows for security proofs for an unbounded number of sessions. Indeed, [13] shows decidability of equivalence, for an unbounded number of sessions, for protocols with an acyclic dependency graph. The notion of dependency graph is introduced in [13] and intuitively captures how the input/output actions of the protocol may use messages from other steps of the protocol. As a corollary, [13] induces a bound on the number of sessions that needs to be considered for an attack, which depends on the size and structure of the graph. This bound can be rather large (50 to 100 sessions, even on small examples) but SAT-Equiv is now able to reach such bounds.

These novelties are implemented in an extension of SAT-Equiv and compared with the other tools of the literature, namely Spec [25], Akiss [8] and the very recent DeepSec [11] tool. Our experiments show that SAT-Equiv is much faster on all the examples, allowing to reach typically more than 100 sessions. As an application, we consider two protocols, Denning-Sacco and Needham-Schroeder symmetric keys, shown to have acyclic dependency graphs in [13]. Considering the necessary number of sessions as induced by [13], we establish trace equivalence for these two protocols, for an unbounded number of sessions.

Due to lack of space, the reader is referred to the companion technical report [17] for the missing proofs and additional details.

Related Work. There are two main families of tools to analyse equivalence properties on security protocols. Some tools prove equivalence for an arbitrary number of sessions, that is, no matter how often a protocol is used. The main tools in this category are ProVerif [6], Tamarin [24], Maude-NPA [22], Type-Eq [18]. Maude-NPA often suffers from termination issues when used for equivalence properties. Type-Eq [18, 19] is a sound (but incomplete) type-checker for equivalence properties that has good performance. It requires that protocols have a similar structure. ProVerif and Tamarin work well in practice. They actually prove a stronger notion of equivalence, diff-equivalence, that also requires that the two considered protocols have a very similar structure. Moreover, equivalence properties are undecidable in general for an unbounded number of sessions. Therefore, ProVerif may not terminate and Tamarin may need some user guidance.

A second approach consists in *deciding* equivalence, for a bounded number of sessions. Spec [25] is one of the first tool that decides equivalence of security protocols but it does not scale well when the number of sessions grows (it can typically handle up to three sessions for small protocols). DeepSec [11] is a very recent tool that builds upon Akiss [8] and Apte [9]. All these tools analyse symbolic executions and typically have to consider all possible interleavings between the roles of the protocol, which often raises efficiency issues.

2 Model

Protocols are modeled through a process algebra, in the spirit of the applied-pi calculus [1]. We consider here a model similar to the ones used *e.g.* in [14, 16].

2.1 Term Algebra

As usual, messages are modeled by terms. Private data are represented through an infinite set \mathcal{N} of *names* used to model *e.g.* keys or nonces. We consider an infinite set \mathcal{C}_0 of constants to represent public data such as agent names or attacker's nonces or keys. We consider also two sets of *variables* \mathcal{X} and \mathcal{W} . Variables in \mathcal{X} model arbitrary data expected by the protocol, while variables in \mathcal{W} are used to store messages learnt by the attacker. A *data* is either a constant, a variable, or a name. Cryptographic primitives are represented by function symbols. We consider the *signature* Σ parameterised by $n \geq 2$:

- $\Sigma_c = \{\text{senc}, \text{aenc}, \text{hash}, \text{pub}, \text{sign}, \text{vk}, \text{ok}\} \cup \{\langle \rangle_k \mid 2 \leq k \leq n\}$;
- $\Sigma_d = \{\text{sdec}, \text{adec}, \text{getmsg}\} \cup \{\text{proj}_j^k \mid 2 \leq k \leq n \text{ and } 1 \leq j \leq k\}$; and
- $\Sigma = \Sigma_c \cup \Sigma_d \cup \{\text{check}\}$.

The symbols `senc`, `aenc`, `sdec`, and `adec` of arity 2 are used to model resp. symmetric and asymmetric encryption. We also consider signature `sign` and hash function `hash`. Concatenation of messages is modeled through tuple operators together with their projection functions. For example, $\langle m_1, m_2, m_3 \rangle_3$ represents the concatenation of the three messages m_1 , m_2 , and m_3 . It is syntactically different from the nested pairs $\langle m_1, \langle m_2, m_3 \rangle_2 \rangle_2$. These two representations correspond to different implementation choices. We distinguish between *constructors* in Σ_c and *destructors* in Σ_d . The symbol `check` of arity 2, which corresponds to the verification of a signature, is neither a destructor nor a constructor. The set of terms built from a signature \mathcal{F} and a set of data D is denoted $\mathcal{T}(\Sigma, D)$. Given a term u , we denote $St(u)$ the set of its *subterms*, $vars(u)$ the set of its *variables*, and $root(u)$ its root symbol. A term is *ground* if it contains no variable. The application of a substitution σ to a term u is written $u\sigma$. We denote $dom(\sigma)$ its *domain* and $img(\sigma)$ its *image*. Two terms u_1 and u_2 are *unifiable* when there exists a substitution σ such that $u_1\sigma = u_2\sigma$.

We consider two *sorts*: `atom` and `bitstring`. The sort `atom` represents atomic data like nonces or keys while `bitstring` models arbitrary messages. Names in \mathcal{N} and constants in \mathcal{C}_0 have sort `atom`. Any $f \in \Sigma_c$ comes with its sorted arity:

$$\begin{array}{ll}
 \langle \rangle_k : \text{bitstring} \times \cdots \times \text{bitstring} \rightarrow \text{bitstring} & \text{ok} : \quad \rightarrow \text{bitstring} \\
 \text{senc} : \text{bitstring} \times \text{atom} \rightarrow \text{bitstring} & \text{pub} : \text{atom} \rightarrow \text{bitstring} \\
 \text{aenc} : \text{bitstring} \times \text{bitstring} \rightarrow \text{bitstring} & \text{vk} : \text{atom} \rightarrow \text{bitstring} \\
 \text{sign} : \text{bitstring} \times \text{atom} \rightarrow \text{bitstring} & \text{hash} : \text{bitstring} \rightarrow \text{bitstring}
 \end{array}$$

Given $D \subseteq \mathcal{C}_0 \uplus \mathcal{X}$, the set $\mathcal{T}_0(\Sigma_c, D)$ is the set of terms t in $\mathcal{T}(\Sigma_c, D)$ such that (i) for any term `pub`(u) (resp. `vk`(u)) in $St(t)$, u is of sort `atom`; (ii) for any `aenc`(u, v) $\in St(t)$, $v = \text{pub}(v')$ for some v' . Terms in $\mathcal{T}_0(\Sigma_c, \mathcal{N} \uplus \mathcal{C}_0)$ are called *messages*. Intuitively, messages are terms with atomic keys.

The properties of the cryptographic primitives are reflected through the following convergent rewriting rules.

$$\begin{array}{ll}
 \text{sdec}(\text{senc}(x, y), y) \rightarrow x & \text{adec}(\text{aenc}(x, \text{pub}(y)), y) \rightarrow x \\
 \text{getmsg}(\text{sign}(x, y)) \rightarrow x & \text{check}(\text{sign}(x, y), \text{vk}(y)) \rightarrow \text{ok} \\
 \text{proj}_j^k(\langle x_1, \dots, x_k \rangle_k) \rightarrow x_j & \text{with } 2 \leq k \leq n \text{ and } 1 \leq j \leq k
 \end{array}$$

A term u can be rewritten into v if there is a position p in u , and a rewriting rule $\mathbf{g}(t_1, \dots, t_n) \rightarrow t$ such that $u|_p = \mathbf{g}(t_1, \dots, t_n)\theta$ for some substitution θ , and $v = u[t\theta]_p$, i.e. u in which the subterm at position p has been replaced by $t\theta$. Moreover, we assume that $t_1\theta, \dots, t_n\theta$ as well as $t\theta$ are messages, in particular they do not contain destructor symbols. As usual, we denote \rightarrow^* the reflexive-transitive closure of \rightarrow , and $u \downarrow$ the *normal form* of a term u .

An attacker builds her own messages by applying public function symbols to terms she already knows and which are available through variables in \mathcal{W} . Formally, a computation done by the attacker is a *recipe*, i.e. a term in $\mathcal{T}(\Sigma, \mathcal{W} \uplus \mathcal{C}_0)$.

2.2 Process Algebra

We consider processes that may receive and send messages. We assume that each process communicates on a dedicated public channel. In practice, IP addresses and sessions identifiers are typically used to desambiguate which message is addressed to whom and for which session. Of course, these channels may be freely manipulated by the attacker. Since we consider equivalence properties, distinct (public) channels provide more abilities for the adversary to distinguish between protocols. Formally, given a set Ch of channels, we consider the fragment of simple processes without replication built on basic processes as defined *e.g.* in [10].

Definition 1. A basic processes is defined as follows:

$$P, Q := 0 \mid \text{in}(c, u_1).P \mid \text{out}(c, u_2).P \mid i:P$$

with $u_1, u_2 \in \mathcal{T}_0(\Sigma_c, \mathcal{C}_0 \uplus \mathcal{N} \uplus \mathcal{X})$, $c \in Ch$, and increasing phase numbers. A simple process is a multiset of basic processes on pairwise distinct channels. A protocol is a simple process such that all its variables are in the scope of an input.

The process 0 does nothing and we often omit it. The process “ $\text{in}(c, u_1).P$ ” expects a message m of the form u_1 on channel c and then behaves like $P\sigma$ where σ is a substitution such that $m = u_1\sigma$. Note that checking whether a received message has the expected form is done through pattern-matching instead of explicit tests. The process “ $\text{out}(c, u_2).P$ ” emits u_2 on c , and then behaves like P . Our calculus also has a phase instruction, in the spirit of [6], denoted $i:P$. This instruction is useful to model security requirements, for example in case the attacker interacts with the protocol before being given some secret.

Example 1. As an illustrative example, we consider a simplified version of the Denning-Sacco protocol which is a key distribution protocol relying on asymmetric encryption and signature. Informally, the protocol is as follows.

$$A \rightarrow B : \text{aenc}(\text{sign}(\langle A, B, K_{ab} \rangle), \text{prv}(A)), \text{pub}(B))$$

The agents A and B aim at authenticating each other and establishing a fresh session key K_{ab} . We model this protocol in our formalism through the simple process $\mathcal{P}_{\text{DS}} = \{P_A; P_B\}$ where $P_A = \text{out}(c_A, \text{aenc}(\text{sign}(\langle a, b, k_{ab} \rangle_3, sk_a), \text{pub}(sk_b))) \cdot 0$ and $P_B = \text{in}(c_B, \text{aenc}(\text{sign}(\langle a, b, x \rangle_3, sk_a), \text{pub}(sk_b))) \cdot 0$ where sk_a, sk_b , and k_{ab} are names, a and b are constants, and x is a variable.

The operational semantics of a process is defined using a relation over configurations. A *configuration* is a tuple $(\mathcal{P}; \phi; \sigma; i)$ with $i \in \mathbb{N}$ and such that:

- \mathcal{P} is a multiset of processes (not necessarily ground);
- $\phi = \{w_1 \triangleright m_1, \dots, w_n \triangleright m_n\}$ is a *frame*, *i.e.* a substitution where w_1, \dots, w_n are variables in \mathcal{W} , and m_1, \dots, m_n are messages;
- σ is a substitution such that $fv(\mathcal{P}) \subseteq \text{dom}(\sigma)$, and $\text{img}(\sigma)$ are messages.

IN	$(i:\text{in}(c, u).P \cup \mathcal{P}; \phi; \sigma; i)$	$\xrightarrow{\text{in}(c,R)}$	$(i:P \cup \mathcal{P}; \phi; \sigma \uplus \sigma_0; i)$	where R is a recipe such that $R\phi\downarrow$ is a message, and $R\phi\downarrow = (u\sigma)\sigma_0$ for σ_0 with $\text{dom}(\sigma_0) = \text{vars}(u\sigma)$.
OUT	$(i:\text{out}(c, u).P \cup \mathcal{P}; \phi; \sigma; i)$	$\xrightarrow{\text{out}(c,w)}$	$(i:P \cup \mathcal{P}; \phi \cup \{w \triangleright u\sigma\}; \sigma; i)$	with w a fresh variable from \mathcal{W} , and $u\sigma$ is a message.
MOVE	$(\mathcal{P}; \phi; \sigma; i)$	$\xrightarrow{\text{phase } i'}$	$(\mathcal{P}; \phi; \sigma; i')$	with $i' > i$.
PHASE	$(i':i'':P \cup \mathcal{P}; \phi; \sigma; i)$	$\xrightarrow{\text{tr}}$	$(i'':P \cup \mathcal{P}; \phi; \sigma; i)$	

Fig. 1. Semantics for processes

A configuration is said to be *initial* when $\sigma = \emptyset$. Intuitively, \mathcal{P} represents the processes that still remain to be executed; ϕ represents the sequence of messages that have been learnt so far by the attacker, and σ stores the value of the variables that have already been instantiated. We write P instead of $0:P$ and $P \uplus \mathcal{P}$ instead of $\{P\} \uplus \mathcal{P}$. Given a protocol \mathcal{P} , we also often write \mathcal{P} instead of $(\mathcal{P}; \emptyset; \emptyset; 0)$. The operational semantics is induced by the relation $\xrightarrow{\alpha}$ over configurations defined in Fig. 1. For example, the IN rule defines how messages can be input on a (public) channel: the adversary may send any message, provided she can construct it through a recipe R applied on her previous knowledge ϕ . Note that only messages can be received (and sent). The relation $\xrightarrow{\text{tr}}$ between configurations (where tr is a possibly empty sequence of actions) is defined in the usual way. Given a configuration \mathcal{K} , we write:

$$\text{trace}(\mathcal{K}) = \{(\text{tr}, \phi) \mid \mathcal{K} \xrightarrow{\text{tr}} (P'; \phi; \sigma; i) \text{ for some configuration } (P'; \phi; \sigma; i)\}.$$

Example 2. Continuing Example 1, let $\mathcal{K}_{\text{DS}} = (\{P_A; P_B; P_{B'}\}; \phi_0; \emptyset; 0)$ where $P_{B'}$ models an additional session of the role B obtained by simply renaming c_B and x with c'_B and x' . The frame $\phi_0 = \{w_a \triangleright \text{vk}(sk_a), w_b \triangleright \text{pub}(sk_b)\}$ models the fact that the attacker initially knows the public key of \mathbf{b} and the verification key of \mathbf{a} . We consider a simple scenario without dishonest participant. The trace $\text{tr}_0 = \text{out}(c_A, w_1).\text{in}(c_B, w_1).\text{in}(c'_B, w_1)$ is executable from \mathcal{K}_{DS} , and yields $\phi = \phi_0 \uplus \{w_1 \triangleright \text{aenc}(\text{sign}(\langle \mathbf{a}, \mathbf{b}, k_{ab} \rangle_3, sk_a), \text{pub}(sk_b))\}$, i.e. $(\text{tr}_0, \phi) \in \text{trace}(\mathcal{K}_{\text{DS}})$.

2.3 Type-Compliance

We present here our main assumption on protocols. Intuitively, we assume that ciphertexts cannot be confused, and we rely for this on a notion of typing system.

Definition 2. A typing system is a pair $(\mathcal{T}_{\text{init}}, \delta)$ where $\mathcal{T}_{\text{init}}$ is a set of elements called initial types, and δ is a function mapping data in $\mathcal{C}_0 \uplus \mathcal{N} \uplus \mathcal{X}$ to types τ :

$$\tau, \tau_1, \tau_2 = \tau_0 \mid \mathbf{f}(\tau_1, \dots, \tau_n) \text{ with } \mathbf{f} \in \Sigma_{\mathcal{C}} \text{ and } \tau_0 \in \mathcal{T}_{\text{init}}$$

Then, δ is extended to constructor terms as follows:

$$\delta(\mathbf{f}(t_1, \dots, t_n)) = \mathbf{f}(\delta(t_1), \dots, \delta(t_n)) \text{ with } \mathbf{f} \in \Sigma_{\mathcal{C}}.$$

A configuration is type-compliant if two unifiable encrypted subterms have the same type. We write $Est(t)$ for the set of *encrypted subterms* of t , i.e. $Est(t) = \{u \in St(t) \mid u \text{ is of the form } f(u_1, \dots, u_n) \text{ and } f \neq \langle \ \rangle_i\}$.

Definition 3. *An initial configuration \mathcal{K} is type-compliant w.r.t. a typing system $(\mathcal{T}_{init}, \delta)$ if for every $t, t' \in Est(\mathcal{K})$ we have that t and t' unifiable implies that $\delta(t) = \delta(t')$.*

Example 3. Continuing our running example, we consider the typing system generated from $\mathcal{T}_{DS} = \{\tau_a, \tau_b, \tau_k, \tau_{sk}\}$ of initial types, and the function δ_{DS} that associates the expected type to each constant/name ($\delta_{DS}(a) = \tau_a$, $\delta_{DS}(k_{ab}) = \tau_k$, etc.), and such that $\delta_{DS}(x) = \delta_{DS}(x') = \tau_k$. We have that \mathcal{K}_{DS} is type-compliant w.r.t. $(\mathcal{T}_{DS}, \delta_{DS})$: unifiable encrypted subterms occurring in the configuration have the same type since $\delta_{DS}(x) = \delta_{DS}(x') = \delta_{DS}(k_{ab})$.

Type-compliant protocols have the property that, when looking for attacks, it is sufficient to consider well-typed execution: $\mathcal{K} \xrightarrow{tr} (\mathcal{P}; \phi; \sigma; i)$ is *well-typed* w.r.t. a typing system $(\mathcal{T}_{init}, \delta)$, if σ is a well-typed substitution, i.e. every variable of its domain has the same type as its image.

2.4 Trace Equivalence

Many privacy properties such as vote-privacy or untraceability are expressed as trace equivalence [2, 21]. Intuitively, two configurations are trace equivalent if an attacker cannot tell with which of the two configurations she is interacting. We first introduce a notion of equivalence (actually, inclusion) between frames.

Definition 4. *Two frames ϕ_1 and ϕ_2 are in static inclusion, written $\phi_1 \sqsubseteq_s \phi_2$, when $dom(\phi_1) = dom(\phi_2)$, and:*

- for any recipe R , we have that $R\phi_1\downarrow$ is a message implies that $R\phi_2\downarrow$ is a message;
- for any recipes R, R' such that $R\phi_1\downarrow, R'\phi_1\downarrow$ are messages, we have that: $R\phi_1\downarrow = R'\phi_1\downarrow$ implies $R\phi_2\downarrow = R'\phi_2\downarrow$.

Intuitively, ϕ_1 is included in ϕ_2 if any recipe producing a message in ϕ_1 also produces a message in ϕ_2 and if any equality satisfied in ϕ_1 is also satisfied in ϕ_2 .

Example 4. We consider $\phi_1 = \phi \uplus \{w_2 \triangleright \text{senc}(m_1, k_{ab}), w'_2 \triangleright \text{senc}(m_1, k_{ab})\}$, and $\phi_2 = \phi \uplus \{w_2 \triangleright \text{senc}(m_2, k), w'_2 \triangleright \text{senc}(m_2, k')\}$ where $m_1, m_2 \in \mathcal{C}_0$. We have that $w_2\phi_1\downarrow = w'_2\phi_1\downarrow$ whereas this equality does not hold in ϕ_2 . Hence $\phi_1 \not\sqsubseteq_s \phi_2$.

Trace inclusion is the active counterpart of static inclusion. Two configurations are in trace inclusion if, however the attacker behaves, the resulting sequences of messages observed by the attacker are in static inclusion.

Definition 5. *Let \mathcal{K} and \mathcal{K}' be two configurations. We have that $\mathcal{K} \sqsubseteq_t \mathcal{K}'$, if for every $(tr, \phi) \in \text{trace}(\mathcal{K})$, there exists $(tr, \phi') \in \text{trace}(\mathcal{K}')$ such that $\phi \sqsubseteq_s \phi'$.*

We easily derive a notion of trace equivalence: two configurations \mathcal{K} and \mathcal{K}' are trace equivalence, denoted $\mathcal{K} \approx_t \mathcal{K}'$, if $\mathcal{K} \sqsubseteq_t \mathcal{K}'$ and $\mathcal{K}' \sqsubseteq_t \mathcal{K}$. This notion of trace equivalence slightly differs from the one used in e.g. [12] but they actually coincide on the class of protocols we consider in this paper [8].

Example 5. To model secrecy of the key k_{ab} , we define *strong secrecy* of k_{ab} by requiring that k_{ab} is indistinguishable from a fresh value. Formally, we consider P_B^1 (resp. $P_{B'}^1$) obtained by replacing the process 0 with $1:\text{out}(c_B, \text{senc}(m_1, x))$ (resp. $1:\text{out}(c'_B, \text{senc}(m_1, x'))$). On the other side of the equivalence, we consider P_B^2 and $P_{B'}^2$, obtained by replacing the process 0 with $1:\text{out}(c_B, \text{senc}(m_2, k))$ (resp. $1:\text{out}(c'_B, \text{senc}(m_2, k'))$) with fresh names k and k' .

$$\mathcal{K}_{\text{DS}}^1 = (\{P_A; P_B^1; P_{B'}^1\}; \phi_0) \text{ and } \mathcal{K}_{\text{DS}}^2 = (\{P_A; P_B^2; P_{B'}^2\}; \phi_0).$$

Then, we can show that $\mathcal{K}_{\text{DS}}^1 \not\sqsubseteq_t \mathcal{K}_{\text{DS}}^2$ since k_{ab} is not strongly secret. An attacker can replay the message sent by A due to lack of freshness. This is exemplified by the trace $\text{tr}_0.\text{out}(c_B, w_2).\text{out}(c'_B, w'_2)$ and the test given in Example 4.

3 From Static Inclusion to Planning

The overall objective of this paper is to provide a practical algorithm for deciding trace inclusion (and thus trace equivalence) relying on graph planning and SAT solving. We start here by explaining how to build a planning problem from two frames such that the planning problem has a solution if, and only if, the two corresponding frames are not in static inclusion.

3.1 Planning Problems

We first recall the definition of a planning problem, slightly simplified from [15]. Intuitively, a planning system defines a transition system from sets of facts to sets of facts. New facts may be produced and some old facts may be deleted.

Definition 6. A planning system is tuple $\langle \text{Fact}, \text{Init}, \text{Rule} \rangle$ where Fact is a set of ground formulas called facts, $\text{Init}_0 \subseteq \text{Fact}$ is a set of facts representing the initial state, and Rule is a set of rules of the form $\text{Pre} \rightarrow \text{Add}; \text{Del}$ where $\text{Pre}, \text{Add}, \text{Del}$ are finite sets of facts such that $\text{Add} \cap \text{Del} = \emptyset, \text{Del} \subseteq \text{Pre}$. We write $\text{Pre} \rightarrow \text{Add}$ when $\text{Del} = \emptyset$.

Given a rule $r \in \text{Rule}$ of the form $\text{Pre} \rightarrow \text{Add}; \text{Del}$, we denote $\text{Pre}(r) = \text{Pre}, \text{Add}(r) = \text{Add}$, and $\text{Del}(r) = \text{Del}$. If $S \subseteq \text{Fact}$ are such that $\text{Pre}(r) \subseteq S$, then we say that the rule is *applicable* in S , denoted $S \xrightarrow{r} S'$, and the state $S' = (S \setminus \text{Del}) \cup \text{Add}$ is the state resulting from the application of r to S . We allow some rules to be applied in parallel when no facts are deleted. Given $S \subseteq \text{Fact}$, and a set of rules $\{r_1, \dots, r_k\}$ such that $\text{Del}(r_i) = \emptyset$ and $\text{Pre}(r_i) \subseteq S$ for any $i \in \{1, \dots, k\}$, $\{r_1, \dots, r_k\}$ is *applicable* in S , denoted $S \xrightarrow{\{r_1, \dots, r_k\}} S'$, and the state $S' = \bigcup_{i=1}^k \text{Add}(r_i) \cup S$ is the state resulting from the application of $\{r_1, \dots, r_k\}$ to S .

A *planning path* from $S_0 \subseteq \mathcal{Fact}$ to $S_n \subseteq \mathcal{Fact}$ is a sequence r_1, \dots, r_n made of rules or sets of rules in \mathcal{Rule} such that $S_0 \xrightarrow{r_1} S_1 \xrightarrow{r_2} \dots S_{n-1} \xrightarrow{r_n} S_n$ for some states $S_1, \dots, S_{n-1} \subseteq \mathcal{Fact}$. A *planning problem* for a system $\Theta = \langle \mathcal{Fact}, \mathit{Init}, \mathcal{Rule} \rangle$ is a pair $\Pi = \langle \Theta, S_f \rangle$ where $S_f \subseteq \mathcal{F}$ represents the target facts. A solution to $\Pi = \langle \Theta, S_f \rangle$, called a *plan*, is a planning path from Init to a state S_n such that $S_f \subseteq S_n$.

A transition $S \xrightarrow{\{r_1, \dots, r_k\}} S'$ can be mimicked by $S \xrightarrow{r_1} S_1 \xrightarrow{r_2} \dots \xrightarrow{r_k} S'$, thus the possibility of applying set of rules in a single step does not change the set of reachable states from a given state S . However, this allows us to consider plans of smaller length and will be useful later on to derive a tight bound and ensure the termination of our algorithm.

In this section, we explain the translation of static inclusion into a planning problem. We consider an (infinite) set \mathcal{Fact}_0 of facts that represent the attacker's knowledge, i.e. formulas of the form $\mathbf{att}(u_P, u_Q)$ where u_P and u_Q are messages, plus a special symbol \mathbf{bad} . Intuitively, $\mathbf{att}(u_P, u_Q)$ means that the attacker knows u_P in the "left" frame, while she knows u_Q in the "right" one.

3.2 Attacker Analysis Rules

Following [16], we first describe the planning rules that correspond to the analysis part of the attacker behaviours. We start by describing a set of abstract rules \mathbf{R}_{Ana} that will be instantiated later on, yielding a (concrete) planning system.

$$\begin{aligned} \mathbf{att}(\langle x_1, \dots, x_k \rangle_k, \langle y_1, \dots, y_k \rangle_k) &\rightarrow \mathbf{att}(x_i, y_i) \text{ with } i \leq k \\ \mathbf{att}(\mathbf{senc}(x_1, x_2), \mathbf{senc}(y_1, y_2)), \mathbf{att}(x_2, y_2) &\rightarrow \mathbf{att}(x_1, y_1) \\ \mathbf{att}(\mathbf{aenc}(x_1, \mathbf{pub}(x_2)), \mathbf{aenc}(y_1, \mathbf{pub}(y_2))), \mathbf{att}(x_2, y_2) &\rightarrow \mathbf{att}(x_1, y_1) \\ \mathbf{att}(\mathbf{sign}(x_1, x_2), \mathbf{sign}(y_1, y_2)) &\rightarrow \mathbf{att}(x_1, y_1) \end{aligned}$$

These rules correspond to the attacker's ability to project, decrypt, and retrieve messages from their signature. There is no \mathbf{Del} since the attacker never forgets. Given a rule $r \in \mathbf{R}_{\text{Ana}}$, we explain how to compute its concretization denoted $\mathbf{Concrete}(r)$. Formally, we have that $\mathbf{Concrete}(r) = \mathbf{Concrete}^+(r) \cup \mathbf{Concrete}^-(r)$.

$\mathbf{Concrete}^+(r)$. The positive concretizations of r consist of instantiating r such that the resulting terms are messages. More formally, we have:

$$\mathbf{Concrete}^+(r) = \{r\sigma \mid \sigma \text{ substitution such that } r\sigma \text{ only involve messages.}\}$$

$\mathbf{Concrete}^-(r)$. We say that a sequence of ground facts $\mathbf{att}(u_1, v_1), \dots, \mathbf{att}(u_k, v_k)$ left-unifies with a sequence $\mathbf{att}(u'_1, v'_1), \dots, \mathbf{att}(u'_k, v'_k)$ if there exists σ such that $u'_1\sigma = u_1, \dots, u'_k\sigma = u_k$ (and symmetrically for right-unification). Given an abstract attacker rule $r = \mathbf{Pre} \rightarrow \mathbf{Add}$, we define $\mathbf{Concrete}^-(r)$ as the set containing $f_1, \dots, f_k \rightarrow \mathbf{bad}$ for any sequence of facts $f_1, \dots, f_k \in \mathcal{Fact}_0$ such that f_1, \dots, f_k left-unifies with \mathbf{Pre} , whereas f_1, \dots, f_k does not right-unify with \mathbf{Pre} .

Example 6. The negative concretizations of the abstract rule corresponding to asymmetric decryption are all the concrete rules of the form

$$\mathbf{att}(\mathbf{aenc}(u_1, \mathbf{pub}(u_2)), v), \mathbf{att}(u_2, v') \rightarrow \mathbf{bad}$$

where $u_1, u_2, v, v', \mathbf{aenc}(u_1, \mathbf{pub}(u_2))$ are messages, whereas $\mathbf{adec}(v, v')\downarrow$ is not.

3.3 Static Inclusion

According to Definition 4, to break static inclusion, an attacker may build new terms (using both analysis and synthesis rules) but also check for equalities and computation failures. To encode static inclusion using planning in an efficient way, we need to strictly control the terms that an attacker has to synthesize.

We say that R is *destructor-only* if $R \in \mathcal{T}(\Sigma_d, \mathcal{C}_0 \cup \mathcal{W})$. It is *simple* if there exists destructor-only recipes R_1, \dots, R_k , and a context C made of constructors such that $R = C[R_1, \dots, R_k]$.

Definition 7. Let ϕ, ψ be such that $\text{dom}(\phi) = \text{dom}(\psi)$. We write $\phi \sqsubseteq_s^{\text{simple}} \psi$ if:

1. For each destructor-only recipe R such that $R\phi\downarrow$ is a (resp. atomic) message, $R\psi\downarrow$ is a (resp. atomic) message.
2. For each simple recipe R and destructor-only recipe R' such that $R\phi\downarrow, R'\phi\downarrow$ are messages and $R\phi\downarrow = R'\phi\downarrow$, we have that $R\psi\downarrow = R'\psi\downarrow$.
3. For each destructor-only recipes R, R' , if $R\phi\downarrow = \mathbf{sign}(t, s)$, and $R'\phi\downarrow = \mathbf{vk}(s)$ for some term t and atom s , then $R\psi\downarrow = \mathbf{sign}(t', s')$, and $R'\psi\downarrow = \mathbf{vk}(s')$ for some term t' and atom s' .
4. For each destructor-only recipe R , such that $R\phi\downarrow = \mathbf{pub}(s)$ for some atom s , $R\psi\downarrow = \mathbf{pub}(s')$ for some atom s' .

We write $\phi \sqsubseteq_s^{\text{simple}^+} \psi$ when the test described at item 2 is only performed when (i) either R is destructor-only; (ii) or $\text{root}(R) \notin \{\mathbf{senc}\} \cup \{\langle \rangle_k \mid 2 \leq k \leq n\}$, and $\text{root}(R') \neq \mathbf{adec}$.

This notion of static inclusion is equivalent to the original one.

Lemma 1. Let ϕ and ψ be two frames having the same domain. We have that:

$$\phi \sqsubseteq_s \psi \Leftrightarrow \phi \sqsubseteq_s^{\text{simple}} \psi \Leftrightarrow \phi \sqsubseteq_s^{\text{simple}^+} \psi.$$

From this new characterisation of static inclusion $\sqsubseteq_s^{\text{simple}}$, we derive the planning rules that capture all the cases of failures with those in $\text{Concrete}^-(\mathbf{R}_{\text{Ana}})$.

$$\mathcal{R}_{\text{fail}}^{\text{atom}} = \{\mathbf{att}(u, v) \rightarrow \mathbf{bad} \mid u \text{ is an atom but } v \text{ is not}\}$$

$$\mathcal{R}_{\text{fail}}^{\text{pub}} = \{\mathbf{att}(\mathbf{pub}(u), v) \rightarrow \mathbf{bad} \mid v \text{ is not of the form } \mathbf{pub}(v')\}$$

$$\mathcal{R}_{\text{fail}}^{\text{check}} = \left\{ \begin{array}{l} \mathbf{att}(\mathbf{sign}(u_1, u_2), v_1) \rightarrow \mathbf{bad} \mid \text{check}(v_1, v_2)\downarrow \text{ is not a message} \\ \mathbf{att}(\mathbf{vk}(u_2), v_2) \end{array} \right\}$$

$$\mathcal{R}_{\text{fail}}^{\text{test}} = \left\{ \begin{array}{l} \mathbf{att}(u_1, v_1), \dots, \mathbf{att}(u_k, v_k) \rightarrow \mathbf{bad} \mid C \text{ is a constructor context,} \\ \mathbf{att}(C[u_1, \dots, u_k], v) \rightarrow \mathbf{bad} \mid C[u_1, \dots, u_k] \in \text{St}(\phi) \cup \mathcal{C}_0 \\ v \neq C[v_1, \dots, v_k]. \end{array} \right\}$$

Actually, not all subterms of $St(\phi)$ need to be considered. Therefore, we consider an optimised version that captures only the terms that may not be reconstructed from their subterms. Formally, $St_{\text{opti}}(t)$ is defined as follows.

- $St_{\text{opti}}(\langle t_1, t_2 \rangle) = St_{\text{opti}}(t_1) \cup St_{\text{opti}}(t_2)$;
- $St_{\text{opti}}(\text{senc}(t_1, t_2)) = St_{\text{opti}}(t_1)$;
- $St_{\text{opti}}(\text{aenc}(t_1, t_2)) = \{\text{aenc}(t_1, t_2)\} \cup (St_{\text{opti}}(t_1) \setminus \{t_1\})$
- $St_{\text{opti}}(\text{sign}(t_1, t_2)) = \{\text{sign}(t_1, t_2)\} \cup St_{\text{opti}}(t_1)$;
- $St_{\text{opti}}(f(t)) = \{f(t)\}$ with $f \in \{\text{hash}, \text{pub}, \text{vk}\}$.

Thanks to the fact that $\sqsubseteq_s^{\text{simple}^+}$ is equivalent to static inclusion, we may only consider simple recipes which evaluation yields a term in $St_{\text{opti}}(\phi)$.

Lemma 2. *Let ϕ be a frame, $R = C[R_1, \dots, R_k]$ be a simple recipe such that $\text{root}(R) \notin \{\text{senc}\} \cup \{\langle \rangle_k \mid 2 \leq k \leq n\}$, and R' be a destructor-only recipe such that $\text{root}(R') \neq \text{adec}$. Assume that $R\phi \downarrow$ and $R'\phi \downarrow$ are both messages such that $R\phi \downarrow = R'\phi \downarrow$. We have that either C is the empty context, or $R\phi \downarrow \in St_{\text{opti}}(\phi) \cup \mathcal{C}_0$.*

Therefore, $\mathcal{R}_{\text{fail}}^{\text{test}}$ can be replaced by the following (smaller) set of rules:

$$\mathcal{R}_{\text{fail}}^{\text{test}_1} = \{\text{att}(u_1, v_1), \text{att}(u_1, v_2) \rightarrow \text{bad} \mid v_1 \neq v_2\}$$

$$\mathcal{R}_{\text{fail}}^{\text{test}_2} = \{\text{att}(u_1, v_1), \dots, \text{att}(u_k, v_k), \text{att}(C[u_1, \dots, u_k], v) \rightarrow \text{bad} \mid C \text{ is a non-empty constructor context, } C[u_1, \dots, u_k] \in St_{\text{opti}}(\phi) \cup \mathcal{C}_0, \text{ and } v \neq C[v_1, \dots, v_k]\}$$

Let ϕ and ψ be two frames with $\text{dom}(\phi) = \text{dom}(\psi)$ and built using constants from $\mathcal{C} \subseteq \mathcal{C}_0$. The set of facts associated to ϕ and ψ is defined as follows:

$$\text{Fact}_{\mathcal{C}}(\phi, \psi) = \{\text{att}(a, a) \mid a \in \mathcal{C}\} \cup \{\text{att}(w\phi, w\psi) \mid w \in \text{dom}(\phi)\}$$

Two frames are in static inclusion if, and only if, the corresponding planning system has no solution. Actually, when the frames are not in static inclusion, we provide a bound on the length of the (minimal) plan witnessing this fact.

Proposition 1. *Let ϕ and ψ be two frames with $\text{dom}(\phi) = \text{dom}(\psi)$, and $\Theta = \langle \text{Fact}_0, \text{Fact}_{\mathcal{C}_0}(\phi, \psi), \mathcal{R} \rangle$ where*

$$\mathcal{R} = \text{Concrete}(\mathcal{R}_{\text{Ana}}) \cup \mathcal{R}_{\text{fail}}^{\text{test}_1} \cup \mathcal{R}_{\text{fail}}^{\text{test}_2} \cup \mathcal{R}_{\text{fail}}^{\text{atom}} \cup \mathcal{R}_{\text{fail}}^{\text{check}} \cup \mathcal{R}_{\text{fail}}^{\text{pub}}.$$

Let $\Pi = \langle \Theta, \{\text{bad}\} \rangle$. We have that $\phi \not\sqsubseteq_s \psi$ if, and only if, Π has a solution of length at most $(N + 1) \times \text{depth}(\phi) + 1$ where N is the number of names n occurring in ϕ at a key position, i.e. such that n (resp. $\text{pub}(n)$) occurs in key position of an encryption in ϕ .

Intuitively, once all needed keys are derived, the minimal plan witnessing non-inclusion contains at most $\text{depth}(\phi)$ rules where $\text{depth}(\phi)$ is the maximal depth of a term occurring ϕ . Then we may need $\text{depth}(\phi)$ rule to derive each deducible key, hence the bound.

4 From Trace Inclusion to Planning

We are now ready for the active case. Given two configurations, we show how to build a planning problem such that the planning problem has a solution if, and only if, the two corresponding configurations are not in trace inclusion.

In several places of this section, we will consider three special constants, namely c_0^* and c_1^* of sort `atom`, and c_+^* of sort `bitstring`. These three constants have a special type, denoted τ_* .

4.1 Abstract Protocol Rules

We first define the abstract rules describing the protocol behaviour. We denote $\mathcal{C}_{\mathcal{P}}$ (resp. $\mathcal{C}_{\mathcal{Q}}$) the constants from \mathcal{C}_0 occurring in \mathcal{P} (resp. \mathcal{Q}), and we consider $\mathcal{C}^* = (\mathcal{C}_{\mathcal{P}} \cup \mathcal{C}_{\mathcal{Q}}) \uplus \{c_0^*, c_1^*, c_+^*\}$. For simplicity we assume that variables of \mathcal{P} and \mathcal{Q} are disjoint. In addition to the facts of the form $\text{att}(u, v)$ used to represent attacker's knowledge, we also consider:

- facts of the form $\text{Phase}(i)$ with $i \in \mathbb{N}$ to represent phases; and
- facts of the form $\text{St}(P, Q) = \text{state}_{P, Q}^c(id_P, id_Q)$ where P, Q are two basic processes on channel c , and id_P (resp. id_Q) is the identity substitution of domain $fv(P)$ (resp. $fv(Q)$).

Therefore, in this section, we consider the infinite set of facts \mathcal{Fact}_0 that consists of all the ground facts of this form, plus the special symbol `bad`.

To deal with phases, we mimic the `PHASE` rule by considering basic processes in normal form w.r.t. the rule $i:j:P \rightarrow j:P$. Then, the transformation $\text{Rule}(P; Q)$ from basic processes (in normal form) to abstract planning rules is defined by $\text{Rule}(P; Q) = \emptyset$ when $P = i:0$, and otherwise:

1. Case output: i.e. if $P = i:\text{out}(c, u).P'$.
 - $\{\text{St}(P, Q), \text{Phase}(i) \rightarrow \text{att}(u, v), \text{St}(P', Q'); \text{St}(P, Q)\} \cup \text{Rule}(i:P'; i:Q')$
when if $Q = i:\text{out}(c, v).Q'$
 - $\{\text{St}(P, Q), \text{Phase}(i) \rightarrow \text{att}(u, c_0^*), \text{bad}\}$ otherwise.
2. Case input: i.e. $P = i:\text{in}(c, u).P'$.
 - $\{\text{St}(P, Q), \text{att}(u, v), \text{Phase}(i) \rightarrow \text{St}(P', Q'); \text{St}(P, Q)\} \cup \text{Rule}(i:P'; i:Q')$
when $Q = i:\text{in}(c, v).Q'$
 - $\{\text{St}(P, Q), \text{att}(u, x), \text{Phase}(i) \rightarrow \text{bad}\}$ otherwise (with x fresh).

Intuitively, abstract rules simply try to mimic each step of P by a similar step in Q . Clearly, if Q cannot follow P , the two processes are not in trace equivalence, which is modelled here by the `bad` state. Note that, in case $P = i:\text{out}(c, u).P'$ whereas Q is not ready to perform an output, `bad` will be triggered only if the sent term is indeed a message. This transformation is then extended to protocols in a natural way considering in addition planning rule to model phase changes. We consider $\mathcal{P} = \{P_1, \dots, P_n\}$ and $\mathcal{Q} = \{Q_1, \dots, Q_n\}$, and we assume w.l.o.g. that P_i and Q_i are basic processes on channel c_i . We define:

- $\text{Rule}(\mathcal{P}, \mathcal{Q}) = \text{Rule}(P_1, Q_1) \cup \dots \cup \text{Rule}(P_n, Q_n)$.
- $\mathcal{R}^{\text{phase}} = \{\text{Phase}(i) \rightarrow \text{Phase}(i+1); \text{Phase}(i) \mid i \in \mathbb{N}\}$.

4.2 Concrete Protocol Rules

To derive concrete rules from the abstract ones, we could instantiate them with arbitrary terms. However, this would not allow us to derive a decision procedure. Moreover, we would like our algorithm to have good performance. To achieve this, we first show that only three constants need to be considered (and no nonces), in addition to those explicitly mentioned in the protocol.

Given a protocol \mathcal{P} that is type-compliant w.r.t. to a typing system $(\mathcal{T}_{\mathcal{P}}, \delta_{\mathcal{P}})$ (and such that τ_* does not occur in $\delta_{\mathcal{P}}(\mathcal{P})$), an execution $\mathcal{P} \xrightarrow{\text{tr}} (\mathcal{P}'; \phi'; \sigma'; i')$ is *quasi-well-typed* if $\delta_{\mathcal{P}}(x\sigma') \preceq \delta_{\mathcal{P}}(x)$ for every variable $x \in \text{dom}(\sigma')$ where \preceq is the smallest relation on types defined as follows:

- $\tau_* \preceq \tau$ and $\tau \preceq \tau$ for any type τ (initial or not);
- $f(\tau_1, \dots, \tau_k) \preceq f(\tau'_1, \dots, \tau'_k)$ when $\tau_1 \preceq \tau'_1, \dots, \tau_k \preceq \tau'_k$, and $f \in \Sigma_c$.

The attacker needs at most the constants c_0^*, c_1^*, c_+^* to mount an attack.

Theorem 1. *Let $\mathcal{K}_{\mathcal{P}}$ be an initial configuration type-compliant w.r.t. $(\mathcal{T}_{\mathcal{P}}, \delta_{\mathcal{P}})$ and $\mathcal{K}_{\mathcal{Q}}$ be another initial \mathcal{C}_0 -configuration. We have that $\mathcal{K}_{\mathcal{P}} \not\sqsubseteq_t \mathcal{K}_{\mathcal{Q}}$ if, and only if, there exists a witness $(\text{tr}, \phi) \in \text{trace}(\mathcal{K}_{\mathcal{P}})$ of this non-inclusion which only involves constants from \mathcal{C}^* , simple recipes, and with a quasi-well-typed underlying execution.*

The existence of a quasi well-typed witness comes from [14] with some extra work to guarantee that we can consider simple recipes. The reduction to three constants extends the previous reduction [16] to asymmetric primitives.

Flattening. In terms of efficiency, one key step of our algorithm is to avoid composition rules from the attacker. For static inclusion, we only consider specific contexts, hence very specific synthesis rules, guided by the form of the underlying frames. For the active case, we transform protocol rules in order to pre-compute all necessary composition steps. This flattening step was already used in *e.g.* [3, 16], and is quite intuitive.

Example 7. Consider our Denning Sacco protocol presented in Example 1. Agent B expects a message of the form $u = \{\text{sign}(\langle \mathbf{a}, \mathbf{b}, x \rangle_3, sk_a)\}_{\text{pub}(sk_b)}$. Either the attacker obtains a message m of the expected form, or the attacker obtains several components of it and forges the whole message. For example, it is sufficient for him to obtain m_1 of the form $u_1 = \text{sign}(\langle \mathbf{a}, \mathbf{b}, x \rangle_3, sk_a)$ and m_2 of the form $u_2 = \text{pk}(sk_b)$. Therefore, in addition to the (informal) protocol rule $u \rightarrow \dots$, we also consider the rule $u_1, u_2 \rightarrow \dots$. Similarly, we also need to consider the rules $\mathbf{a}, \mathbf{b}, x, sk_a, \text{pk}(sk_b) \rightarrow \dots$ and $\mathbf{a}, \mathbf{b}, x, sk_a, sk_b \rightarrow \dots$.

More generally, given an abstract protocol rule r , we now define $\text{Flat}(r)$ the set of rules obtained by performing flattening on each fact. To decompose a term, we follow its structure, and the structure of a variable is given by its type. Moreover, when the other side of the process is not able to follow the decomposition, this leads us to a failure rule.

Definition 8. Given a term $u \in \mathcal{T}_0(\Sigma_c, \mathcal{C}_0 \uplus \mathcal{N} \uplus \mathcal{X})$, we say that u is decomposable when either $u \in \mathcal{X}$ and $\delta_{\mathcal{P}}(u)$ is not an initial type; or $u \notin \mathcal{C}_0 \uplus \mathcal{N} \uplus \mathcal{X}$.

A variable of non initial type is decomposable since it may be instantiated by a non atomic term which, in turns, may have been obtained by composition. Given $\text{att}(u, v)$ with u decomposable, and let $f \in \Sigma_c$ be such that $\delta_{\mathcal{P}}(u) = f(\tau_1, \dots, \tau_k)$, $\text{split}(\text{att}(u, v)) = (f; \{\text{att}(x_1, y_1), \dots, \text{att}(x_k, y_k)\}; \sigma_{\mathcal{P}}; \sigma_{\mathcal{Q}})$ where

- x_1, \dots, x_k are fresh variables of type τ_1, \dots, τ_k , $\sigma_{\mathcal{P}} = \text{mgu}(u, f(x_1, \dots, x_k))$;
- y_1, \dots, y_k are fresh variables, $\sigma_{\mathcal{Q}} = \text{mgu}(v, f(y_1, \dots, y_k))$.

Note that $\sigma_{\mathcal{P}}$ exists and is necessarily a quasi-well-typed substitution. By convention, we assume that $\text{mgu}(u, u') = \perp$ when u and u' are not unifiable.

Let r be an abstract rule of the form $\text{Pre} \rightarrow \text{Add}; \text{Del}$ with $f = \text{att}(u, v) \in \text{Pre}$ such that u is decomposable and $\text{split}(f) = (f, S, \sigma_{\mathcal{P}}, \sigma_{\mathcal{Q}})$. The decomposition of r w.r.t. f , denoted $\text{decom}(r, f)$, is defined as follows:

1. $((\text{Pre} \setminus f) \cup S \rightarrow \text{bad})\sigma_{\mathcal{P}}$ in case $\sigma_{\mathcal{Q}} = \perp$;
2. $((\text{Pre} \setminus f) \cup S \rightarrow \text{Add}; \text{Del})(\sigma_{\mathcal{P}} \uplus \sigma_{\mathcal{Q}})$ otherwise.

Then, decomposition is applied recursively on each rule.

$$\text{Flat}(r) = \{r\} \cup \text{Flat}(\{\text{decom}(r, f) \mid f = \text{att}(u, v) \in \text{Pre}(r) \text{ with } u \text{ decomposable}\})$$

Concretization. Given an abstract rule r , we denote $\text{vars}_{\text{left}}(r)$ the variables occurring on the left (first parameter) of a predicate occurring in r , i.e.

$$\text{vars}_{\text{left}}(\text{att}(u, v)) = \text{vars}(u); \text{ and } \text{vars}_{\text{left}}(\text{state}_{P,Q}^c(\sigma_{\mathcal{P}}, \sigma_{\mathcal{Q}})) = \text{vars}(\text{img}(\sigma_{\mathcal{P}})).$$

Given a substitution σ grounding for r , the application of σ on an abstract state is the concrete state obtained by simply composing the substitutions, i.e.

$$\text{st}_{P,Q}^c(\sigma_{\mathcal{P}}, \sigma_{\mathcal{Q}})\sigma = \text{st}_{P,Q}^c(\sigma \circ \sigma_{\mathcal{P}}, \sigma \circ \sigma_{\mathcal{Q}}).$$

Given an abstract protocol rule r , its concretizations $\text{Concrete}(r)$ simply consist in all its positive and negative concretizations. The positive concretizations are all its instantiations that are quasi-well-typed w.r.t. the left side of the rule.

$$\text{Concrete}^+(r) = \{r\sigma \mid \sigma \text{ substitution such that } r\sigma \text{ only involves messages with constants in } \mathcal{C}^* \text{ and } \delta_{\mathcal{P}}(x\sigma) \preceq \delta_{\mathcal{P}}(x) \text{ for any } x \in \text{vars}_{\text{left}}(r)\}$$

Similarly to the static case, we need to make sure that we can detect when P and Q are *not* in trace inclusion, and we therefore consider some additional rules. Given an abstract protocol rule $r = \text{Pre} \rightarrow \text{Add}; \text{Del}$, $\text{Concrete}^-(r)$ is the set of planning rules that contains: $f_1, \dots, f_k \rightarrow \text{bad}$ for any sequence of facts f_1, \dots, f_k such that f_1, \dots, f_k left-unify with Pre with substitution σ_L and $u \in \mathcal{T}_0(\Sigma_c, \mathcal{N} \cup \mathcal{C}^*)$ for any $\text{att}(u, v) \in \text{Add}\sigma_L$, and such that one of the following conditions holds:

- f_1, \dots, f_k does not right-unify with Pre ;
- f_1, \dots, f_k right-unify with Pre with substitution σ_R but $v \notin \mathcal{T}_0(\Sigma_c, \mathcal{N} \cup \mathcal{C}^*)$ for some $\text{att}(u, v) \in \text{Add}\sigma_R$.

Main Result. Our main technical result states that our encoding is sound and complete: two protocols are in trace inclusion if, and only if, the corresponding planning system has a solution. Moreover, when a witness of non-inclusion exists, we are able to bound the length of the resulting plan. Below, $\text{nb}_{\text{in}}(\mathcal{P})$ (resp. $\text{nb}_{\text{out}}(\mathcal{P})$) denotes the number of inputs (resp. outputs) occurring in \mathcal{P} whereas $\text{max}_{\text{phase}}(\mathcal{P})$ is the maximal integer occurring in a phase instruction in \mathcal{P} .

Theorem 2. *Let \mathcal{P} a protocol type-compliant w.r.t. $(\mathcal{T}_{\mathcal{P}}, \delta_{\mathcal{P}})$, and \mathcal{Q} be another protocol. We consider the following set \mathcal{R} of concrete rules:*

$$\text{Concrete}(\text{R}_{\text{Ana}} \cup \text{Flat}(\text{Rule}(\mathcal{P}, \mathcal{Q}))) \cup \mathcal{R}^{\text{phase}} \cup \mathcal{R}_{\text{fail}}^{\text{test}_1} \cup \mathcal{R}_{\text{fail}}^{\text{test}_2} \cup \mathcal{R}_{\text{fail}}^{\text{atom}} \cup \mathcal{R}_{\text{fail}}^{\text{check}} \cup \mathcal{R}_{\text{fail}}^{\text{pub}}$$

Let $\Theta = \langle \text{Fact}_0, \text{Fact}_{\mathcal{C}^*}(\mathcal{P}, \mathcal{Q}), \mathcal{R} \rangle$ and $\Pi = \langle \Theta, \{\text{bad}\} \rangle$. We have that $\mathcal{P} \not\sqsubseteq_t \mathcal{Q}$ if, and only if, Π has a solution of length

$$1 + \text{nb}_{\text{in}}(\mathcal{P}) + \text{nb}_{\text{out}}(\mathcal{P}) + \text{max}_{\text{phase}}(\mathcal{P}) + \text{depth}(\delta_{\mathcal{P}}(\mathcal{P})) \times [1 + \text{nb}_{\text{in}}(\mathcal{P}) + N]$$

where N is the number of names occurring in \mathcal{P} having a key type, i.e. such that $\delta_{\mathcal{P}}(n)$ (resp. $\text{pub}(\delta_{\mathcal{P}}(n))$) occurs in key position of an encryption in $\delta_{\mathcal{P}}(\mathcal{P})$.

Proof (Sketch). It is rather easy to establish that a solution to the planning problem defines a witness of non trace inclusion. Conversely, thanks to Theorem 1, if $\mathcal{P} \not\sqsubseteq_t \mathcal{Q}$, then there exists a quasi well-typed witness of non trace inclusion, that uses at most three constants (besides the constants of \mathcal{P} and \mathcal{Q}). This witness guides the definition of a plan of Π . Establishing a not too coarse bound on its length requires some care. It relies on the flattening of the protocol and the fact that the plan can mimic the computation of several messages in parallel. \square

5 Algorithm

Similarly to the algorithm presented in [16], we decide trace inclusion by applying graph planning and SAT-solving techniques to the planning problem that encodes trace inclusion (thanks to Theorem 2). Given a protocol \mathcal{P} , type-compliant w.r.t. $(\mathcal{T}_{\mathcal{P}}, \delta_{\mathcal{P}})$, and a protocol \mathcal{Q} , our algorithm proceeds as follows.

1. It first computes the corresponding abstract rules, namely $\text{Flat}(\text{Rule}(\mathcal{P}; \mathcal{Q})) \cup \text{R}_{\text{Ana}}$ and the initial state $\text{Fact}(\mathcal{P}, \mathcal{Q})$.
2. It then applies a planning graph algorithm, a standard technique to solve planning problems (see e.g. [7]). The only difference is that, for efficiency reasons, we do not construct the planning problem Π *a priori* but instead, we compute it “on the fly”, while building the associated planning graph. This planning graph over-approximates the possible solutions by executing several actions in parallel, even if they may be incompatible. Some incompatibilities are recorded and propagated through so-called mutex. The planning graph is deemed to capture all possible plans. More precisely, the planning graph built until depth k captures all possible plans of length at most k .

3. In case no fact **bad** has been reached while building the planning graph, we can immediately conclude that $\mathcal{P} \sqsubseteq_t \mathcal{Q}$. Otherwise, since the planning graph over-approximates the possible executions, we need to check that **bad** is truly reachable. This is done by encoding each path leading to **bad** as a SAT formula. We then call the SAT solver mini-SAT to decide its satisfiability. In case **bad** is indeed reachable, mini-SAT provides a solution that is translated back to a witness of non-inclusion. To improve termination, we check accessibility of a state containing **bad** as soon as it appears in the graph, even if the construction of the graph is not completed yet.

Termination. The algorithm defined above may not terminate. The planning graph contains facts of the form $\mathbf{att}(u, v)$ where u must be (quasi) well-typed. There is therefore only a finite number of such u . However, the planning graph construction may introduce several facts of the form $\mathbf{att}(u, v_1), \dots, \mathbf{att}(u, v_k)$, where the v_i get arbitrarily large. We exhibit some (contrived) examples where the algorithm does not terminate (see Appendix A). [16] suggests that termination could be enforced by checking at each step (thanks to the SAT-solver) that each node of the planning graph is indeed reachable. This would however not be practical. Instead, we can enforce termination thanks to the bound provided in Theorem 2 that also bounds the maximal depth of the planning graph that needs to be considered. Indeed, it is sufficient to simply stop the construction of the planning graph as soon as the bound is reached. The interest of this approach is that we guarantee termination at no cost (computing the bound is immediate). In practice, the planning graph is typically much smaller than this bound.

SAT-Equiv. We have implemented our new algorithm in the tool SAT-Equiv, extending it to protocols with phases and all the standard cryptographic primitives and guaranteeing termination. Moreover, we significantly improve its efficiency by rewriting parts of the codes and modifying the data structure.

6 Experiments

In this section, we analyse several protocols of the literature and compare the results obtained using different tools. We ran our experiments on a single Intel 3.1 GHz Xeon. We limit the memory to 128 Go (MO stands for memory out) and the execution time to 24 h (TO stands for time out).

For all the considered protocols, we analyse strong secrecy of the exchanged key or nonce, as for Example 5, except for the passport protocol (PA), for which we prove anonymity as in [2]. We progressively increase the number of sessions in order to consider a *semi complete scenario*, where Alice's role is instantiated by honest a talking to honest b or dishonest c and Bob's role is instantiated by b talking to a or c . This typically corresponds to 7 sessions in the case of a symmetric key protocol (with 3 roles).

	Spec	Akiss	Deepsec	CSF'17	Sat-Eq
Denning-Sacco	7	10	35	98	> 210 (4h)
Needham-Schroeder sym	6	6	21	21	94* (20h30)
Wide Mouth Frog	7	12	28	84	> 210 (6min)
Yahalom-Paulson	6	6	12	7	> 28 (7h)
Passive Authentication	6	8	46	–	> 400 (98s)
Active Authentication	6	8	50	–	> 400 (78s)
Needham-Schroeder-Lowe	4	6	16	–	> 64 (11min)
Denning-Sacco signature	8	8	18	–	> 64 (100s)

Fig. 2. Comparison of SAT-Equiv with the other tools. We indicate the number of sessions for which the tool fails (time out, memory out, or other issues). When we did not reach the limit of the tool, we write $>k$ to indicate that the tool can analyse more than k sessions, and we indicate the analysis time for k . *See Sect. 6.2

6.1 Comparison with the Other Tools

Our experiments show a significant speed-up w.r.t. the original version of SAT-Equiv [16]. Our new is 100 faster in average, allowing to analyse about twice more sessions, as exemplified in Fig. 2. We compare SAT-Equiv with other tools of the literature that decide equivalence for a bounded number of sessions, namely Spec [25], Akiss [8] and Deepsec [11]. We did not include APTE in our study [9] as it is now subsumed by Deepsec. For each protocol, we progressively increased the number of sessions until we reached a time out. The overall results of our experiments are summarized in Fig. 2. They show a significant speed-up even w.r.t. the very recent Deepsec tool. Note however that Deepsec covers more protocols (with else branches, or not type compliant), except if they include phases. Deepsec can also be parallelized thus the analysis time can be divided by the number of available cores. The detailed results for the Denning-Sacco protocol are below.

Denning-Sacco	Spec	Akiss	Deepsec	CSF'17	SAT-Equiv	
3	12 s	0.08 s	<0.01 s	0.3 s	0.07 s	42
6	5 h	9 s	<0.01 s	1 s	0.1 s	64
7	MO	75 s	<0.01 s	2 s	0.2 s	74
10		MO	0.01 s	4 s	0.3 s	114
21			18 s	60 s	1.3 s	216
35			TO	9 min	6 s	344
84				13 h	164 s	792
98				TO	6 min	920
210					4 h 20	1942

The 2nd column for SAT-Equiv indicates the theoretical bound on the length of the planning graph, as given by Theorem 2. This illustrates that this bound remains reasonable although our tool actually terminates before reaching it.

6.2 Towards an Unbounded Number of Sessions

Although equivalence is undecidable in general for an unbounded number of sessions, [13] exhibits a decidability result, for type-compliant protocols that have an *acyclic* dependency graph. Intuitively, the dependency graph captures how a message expected as input may be built (and therefore may depend) from messages sent as output of the protocol. Decidability is proven by showing that a (minimal) attack trace may be mapped to this dependency graph. Looking at the dependency graphs of the Denning-Sacco and the Needham-Schroeder symmetric key protocols, we deduce that it is sufficient to analyse respectively 42 and 94 sessions. Thanks to the efficiency of SAT-Equiv, we can easily analyse 42 sessions of Denning-Sacco (in 10 s). We can therefore deduce from [13] that the protocol remains secure even if the considered sessions are arbitrarily replicated. The case of the Needham-Schroeder protocol requires a bit more work as 94 sessions is slightly out of reach of SAT-Equiv. However, we noticed that, according to [13], we do not need to analyse 94 full sessions. Instead, some of them may be truncated (a minimal attack will use only the first step for example). Since SAT-Equiv can prove equivalence of these refined 94 sessions (in 20h30min), we can again deduce from [13] that the protocol remains secure even if the considered sessions are arbitrarily replicated.

As future work, we plan to optimize the bound on sessions induced by [13] and automatically generate the desired scenario, in order to extend SAT-Equiv to proofs of equivalence for an unbounded number of sessions.

A Examples of Non Termination

We exhibit two examples on which the original SAT-Equiv algorithm does not terminate. Given a channel c , consider $P(c)$ and $Q(c)$ defined as follows.

$$\begin{aligned} P(c) &:= \text{in}(c, \langle x, \mathbf{a} \rangle). \text{out}(c, \langle x, \mathbf{a} \rangle) \\ Q(c) &:= \text{in}(c, \langle x, \mathbf{a} \rangle). \text{out}(c, \langle \langle x, x \rangle, \mathbf{a} \rangle) \end{aligned}$$

where \mathbf{a} is a public constant and x a variable. We consider $\mathcal{K}_P = \{P(c_1); P(c_2)\}$ and $\mathcal{K}_Q = \{Q(c_1); Q(c_2)\}$ for some public channel names c_1, c_2 . Starting with $\text{att}(\mathbf{b}, \mathbf{b})$ (with \mathbf{b} a public constant in the initial knowledge of the attacker), the following facts will be successively added when computing the planning graph:

$$\text{att}(\langle \mathbf{b}, \mathbf{a} \rangle, \langle \mathbf{b}, \mathbf{a} \rangle), \text{att}(\langle \mathbf{b}, \mathbf{a} \rangle, \langle \langle \mathbf{b}, \mathbf{b} \rangle, \mathbf{a} \rangle), \text{att}(\langle \mathbf{b}, \mathbf{a} \rangle, \langle \langle \langle \mathbf{b}, \mathbf{b} \rangle, \langle \mathbf{b}, \mathbf{b} \rangle \rangle, \mathbf{a} \rangle), \dots$$

Actually, $\text{att}(\langle \mathbf{b}, \mathbf{a} \rangle, \langle \langle \mathbf{b}, \mathbf{b} \rangle, \mathbf{a} \rangle)$ can be added in two different ways: either considering the output on c_1 , or the one on c_2 . Therefore this fact will not be put in mutex with the other ones. In particular, the fact $\text{att}(\langle \mathbf{b}, \mathbf{a} \rangle, \langle \langle \mathbf{b}, \mathbf{b} \rangle, \mathbf{a} \rangle)$

and the state fact indicating that the process on channel c_1 has not yet started are not in mutex , and can be used to trigger the planning rules leading to $\text{att}(\langle \mathbf{b}, \mathbf{a} \rangle, \langle \langle \mathbf{b}, \mathbf{b} \rangle, \langle \mathbf{b}, \mathbf{b} \rangle \rangle, \mathbf{a})$. Since the term computed on the Q's side grows at each step, this computation is endless.

Here, \mathcal{K}_P is not trace included in \mathcal{K}_Q : an attacker can distinguish between \mathbf{b} and $\langle \mathbf{b}, \mathbf{b} \rangle$. So, as soon as a message is outputted, the resulting frames are not in static inclusion. Therefore, termination can be retrieved by enforcing SAT-Equiv to stop the exploration of the planning graph as soon as an attack is found.

We can turn this example into a more complex one on which the original SAT-Equiv will not terminate even if we decide to stop the exploration of the planning graph as soon as an attack is found. Consider the processes $P_0(c)$, $P_1(c)$ and $Q_1(c)$ given below. We assume that k is name representing a symmetric secret key, whereas $\mathbf{a}, \mathbf{b}, \mathbf{c}$ are public constants.

$$\begin{aligned} P_0(c) &= \text{in}(c, x).\text{out}(c, \text{senc}(x, k)) \\ P_1(c) &= \text{in}(c, \langle \text{senc}(\mathbf{a}, k), \text{senc}(\mathbf{b}, k), \text{senc}(\mathbf{c}, k) \rangle_3).P(c) \\ Q_1(c) &= \text{in}(c, \langle \text{senc}(\mathbf{a}, k), \text{senc}(\mathbf{b}, k), \text{senc}(\mathbf{c}, k) \rangle_3).Q(c) \end{aligned}$$

We consider the configurations $\mathcal{K}'_P = \{P_0(c_0); P_0(c_1); P_1(c_2); P_1(c_3)\}$ and $\mathcal{K}'_Q = \{P_0(c_0); P_0(c_1); Q_1(c_2); Q_1(c_3)\}$ where c_0, c_1, c_2, c_3 are public channel names. Processes P_0 on channels c_0 and c_1 are used as oracles. Roughly, we can get two ciphertexts among the three ciphertexts: $\text{senc}(\mathbf{a}, k)$, $\text{senc}(\mathbf{b}, k)$, and $\text{senc}(\mathbf{c}, k)$. It is however not possible to get the three of them. Noticing this, it is then easy to see that \mathcal{K}_P and \mathcal{K}_Q are trace included.

However, as in the previous example, the planning graph is not precise enough to detect that it is not possible to obtain these three ciphertexts. Once the inputs on channel c_2 and c_3 are executed, we reach a situation similar to the one discussed in the previous example. Each time bad will be added into the planning graph, our SAT encoding will tell us that this state is not truly reachable (but only exists in the over-approximation). Thus, we will continue to explore the planning graph for ever since no attack will be found (the protocols are trace-equivalent).

References

1. Abadi, M., Fournet, C.: Mobile values, new names, and secure communication. In: Proceedings of 28th ACM Symposium on Principles of Programming Languages, POPL 2001, pp. 104–115. ACM (2001)
2. Arapinis, M., Chothia, T., Ritter, E., Ryan, M.: Analysing unlinkability and anonymity using the applied pi calculus. In: Proceedings of 23rd Computer Security Foundations Symposium (CSF 2010), pp. 107–121. IEEE Computer Society Press (2010)
3. Armando, A., Compagna, L.: Sat-based model-checking for security protocols analysis. *Int. J. Inf. Secur.* **7**, 3–32 (2008)
4. Bhargavan, K., Blanchet, B., Kobeissi, N.: Verified models and reference implementations for the TLS 1.3 standard candidate. In: Proceedings of 38th IEEE Symposium on Security and Privacy (S&P 2017). IEEE Computer Society Press (2017)

5. Blanchet, B.: Symbolic and computational mechanized verification of the ARINC823 avionic protocols. In: Proceedings of 30th IEEE Computer Security Foundations Symposium (CSF 2017), pp. 68–82. IEEE Computer Society Press (2017)
6. Blanchet, B., Abadi, M., Fournet, C.: Automated verification of selected equivalences for security protocols. *J. Logic Algebr. Program.* **75**(1), 3–51 (2008)
7. Blum, A., Furst, M.: Fast planning through planning graph analysis. *Artif. Intell.* **90**, 281–300 (1997)
8. Chadha, R., Ciobăcă, Ș., Kremer, S.: Automated verification of equivalence properties of cryptographic protocols. In: Seidl, H. (ed.) ESOP 2012. LNCS, vol. 7211, pp. 108–127. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-28869-2_6
9. Cheval, V.: APTE: an algorithm for proving trace equivalence. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014. LNCS, vol. 8413, pp. 587–592. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-642-54862-8_50
10. Cheval, V., Cortier, V., Delaune, S.: Deciding equivalence-based properties using constraint solving. *Theoret. Comput. Sci.* **492**, 1–39 (2013)
11. Cheval, V., Kremer, S., Rakotonirina, I.: DEEPSEC: deciding equivalence properties in security protocols - theory and practice. In: Proceedings of 39th IEEE Symposium on Security and Privacy (S&P 2018), pp. 525–542. IEEE Computer Society Press (2018)
12. Chrétien, R., Cortier, V., Delaune, S.: Typing messages for free in security protocols: the case of equivalence properties. In: Baldan, P., Gorla, D. (eds.) CONCUR 2014. LNCS, vol. 8704, pp. 372–386. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-44584-6_26
13. Chrétien, R., Cortier, V., Delaune, S.: Decidability of trace equivalence for protocols with nonces. In: Proceedings of the 28th IEEE Computer Security Foundations Symposium (CSF 2015). IEEE Computer Society Press (2015)
14. Chrétien, R., Cortier, V., Dallon, A., Delaune, S.: Typing messages for free in security protocols. Technical report (2018)
15. Compagna, L.: SAT-based model-checking of security protocols. Ph.D. thesis, Università degli Studi di Genova and the University of Edinburgh (joint programme), September 2005
16. Cortier, V., Dallon, A., Delaune, S.: SAT-equiv: an efficient tool for equivalence properties. In: Proceedings of the 30th IEEE Computer Security Foundations Symposium (CSF 2017). IEEE Computer Society Press, August 2017
17. Cortier, V., Dallon, A., Delaune, S.: Efficiently deciding equivalence for standard primitives and phases. Research report, June 2018. <https://hal.archives-ouvertes.fr/hal-01819366>
18. Cortier, V., Grimm, N., Lallemand, J., Maffei, M.: A type system for privacy properties. In: Proceedings of 24th ACM Conference on Computer and Communications Security (CCS 2017), pp. 409–423. ACM (2017)
19. Cortier, V., Grimm, N., Lallemand, J., Maffei, M.: Equivalence properties by typing in cryptographic branching protocols. In: Bauer, L., Küsters, R. (eds.) POST 2018. LNCS, vol. 10804, pp. 160–187. Springer, Cham (2018). https://doi.org/10.1007/978-3-319-89722-6_7
20. Cremers, C., Horvat, M., Hoyland, J., Scott, S., van der Merwe, T.: A comprehensive symbolic analysis of TLS 1.3. In: Proceedings of 24th ACM Conference on Computer and Communications Security (CCS 2017), pp. 1773–1788. ACM (2017)
21. Delaune, S., Kremer, S., Ryan, M.: Verifying privacy-type properties of electronic voting protocols. *J. Comput. Secur.* **4**, 435–487 (2008)

22. Escobar, S., Meadows, C., Meseguer, J.: A rewriting-based inference system for the NRL protocol analyzer and its meta-logical properties. *Theoret. Comput. Sci.* **367**(1–2), 162–202 (2006)
23. Kautz, H., Selman, B.: Planning as satisfiability. In: *Proceedings 10th European Conference on Artificial Intelligence (ECAI 1992)*, pp. 359–363 (1992)
24. Meier, S., Schmidt, B., Cremers, C., Basin, D.: The TAMARIN prover for the symbolic analysis of security protocols. In: Sharygina, N., Veith, H. (eds.) *CAV 2013*. LNCS, vol. 8044, pp. 696–701. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39799-8_48
25. Tiu, A., Dawson, J.: Automating open bisimulation checking for the spi calculus. In: *Proceedings of 23rd IEEE Computer Security Foundations Symposium (CSF 2010)*, pp. 307–321. IEEE Computer Society Press (2010)