

Symbolic trace analysis of cryptographic protocols*

Michele Boreale
Università di Firenze
Dipartimento di Sistemi e Informatica

Abstract

A cryptographic protocol can be described as a system of concurrent processes, and analysis of the traces generated by this system can be used to verify authentication and secrecy properties of the protocol. However, this approach suffers from a state-explosion problem that causes the set of states and traces to be typically infinite or very large. In this paper, starting from a process language inspired by the spi-calculus, we propose a symbolic operational semantics that relies on unification and leads to compact models of protocols. We prove that the symbolic and the conventional semantics are in full agreement, and then give a method by which trace analysis can be carried out directly on the symbolic model. The method is proven to be complete for the considered class of properties and is amenable to automatic checking.

Keywords: spi-calculus, concurrency, formal methods for security protocols.

Author's Address: Dipartimento di Sistemi e Informatica, Università di Firenze, Via Lombroso 6/17, Firenze.

Phone: +39-55-4796769

Fax: +39-55-4796730

E-mail: boreale@dsi.unifi.it

*A preliminary version of this paper has been circulated as [5].

1 Introduction

In recent years, formal methods have proven useful in the analysis of cryptographic protocols, often revealing previously unknown attacks. A popular approach is that of modelling a protocol as a system of concurrent processes, described using an appropriate language, like CSP [13, 20, 22] or the spi-calculus [1] – the latter an extension of the π -calculus [18]. In this setting, Abadi and Gordon advocate the use of observational equivalences to formalize and verify protocol properties [2, 7]. Here, in the vein of [3, 4, 12, 13, 16, 20, 22], we analyze the sequences of actions (traces) that a given process may execute. As an example, a *secrecy* property like “protocol P never leaks the datum d ”, might be verified by adding to the description of P an ‘error’ action to be performed as soon as the environment learns d (the way this is done depends on the specific formalism, see e.g. [3]), and then checking that P never performs that ‘error’ action.

The main drawback of trace analysis is that the execution of a protocol typically generates infinitely many traces. The reason lies in the modelling of the environment, whose behaviour is largely unpredictable. Rather than trying to describe this behaviour as a specific process, it is sensible to simply assume that the communication network is totally under the control of the environment. The latter can store, duplicate, hide or replace messages that travel on the network. It can also operate according to the rules followed by honest participants and synthesize new messages by pairing, decryption, encryption and creation of fresh nonces and keys, or by arbitrary combinations of these operations (this approach seems to date back to Dolev and Yao [11]). Thus, an agent waiting for an input at a given moment may expect any of the infinitely many messages the environment can produce and send on the network. This leads to a state explosion that makes the protocol model, typically a state-transition graph, infinite (more precisely, infinite-branching). In practice, those approaches that rely on *model checking* [13, 20, 22] cut down the model to a convenient finite size by imposing upper-bounds to the critical parameters (number of keys, number of pairing and encryption in messages, ...). Exhaustive exploration of the state-space is then possible by standard techniques. However, this approach makes correctness in the general case (completeness) very difficult to establish – though some progress has recently been made [15, 21]. Furthermore, even when those upper-bounds can be justified and the model is finite, the branching factor of input actions may cause the number of states and traces to explode as larger systems are considered.

In this paper we explore an alternative approach to trace analysis of cryptographic protocols. As a base language, we consider a variant of the spi-calculus, but this choice is not critical for the development of the theory. The idea is to replace the infinitely many transitions arising from an input action by a single *symbolic* transition, and to represent the received message as a variable. Constraints on this variable are accumulated as the execution proceeds. Let us see this in more detail. In the variant of the spi-calculus we use, the receiver of a message is written as $\mathbf{a}(x).R$, where \mathbf{a} is an arbitrary label, x is the input variable and R is the continuation. The conventional (‘concrete’) operational semantics of the language requires x to be instantiated with each message that can possibly be received from the network, and this causes the state explosion. In our symbolic semantics, x is not instantiated immediately, rather constraints on its value are added as needed. These constraints take the form of *most general unifiers*. As an example, suppose that a process P , after receiving a message x , tries decryption of x using key k , and, if this succeeds, calls y the result and proceeds like P' . This is written as $P \stackrel{\text{def}}{=} \mathbf{a}(x).\text{case } x \text{ of } \{y\}_k \text{ in } P'$. We represent a state of the protocol as a pair (σ, Q) , where σ is the trace of process’ past actions and Q is a process term. The two initial symbolic transitions of (ϵ, P) will be:

$$(\epsilon, P) \longrightarrow_s (\mathbf{a}\langle x \rangle, \text{case } x \text{ of } \{y\}_k \text{ in } P') \longrightarrow_s (\mathbf{a}\langle \{y\}_k \rangle, P'[\{y\}_k/x])$$

where $\{y\}_k/x$ is the most general unifier for x and $\{y\}_k$. In general, protocols not using replication/recursion will generate finitely many symbolic transitions. The resulting model is rather

compact: sequential processes will just exhibit a single complete symbolic trace, while, for parallel compositions, traces will be obtained as usual by interleaving. We prove that the symbolic and the conventional semantics are in full agreement, and then give a method by which trace analysis can be carried out directly on the symbolic traces. We focus our attention on a specific class of properties, those of the form “in every execution of the protocol, action α happens prior to action β ”, for given α and β . As we shall see, this scheme is flexible enough to express interesting forms of authentication and secrecy. The method is proven to be complete for the language we consider, and is easily mechanizable; this immediately yields decidability of trace analysis for the considered language. When a protocol does not satisfy a property, the method also gives an easy way to compute an attack, i.e. a trace that violates the property. A prototype implementation of the method is already available [6].

The language we consider does not contain replication/recursion operators, that would make trace analysis undecidable (see e.g. [12, 10]). Thus we consider only protocols with a bounded number of participants. However, recent work by Lowe [15] and Stoller [23] indicates that, in some cases, it is possible to reduce the analysis of an unbounded protocol to the analysis of a protocol with a bounded, statically determined number of participants.

The symbolic approach to the analysis of security protocols has been explored by other authors, including Amadio and Lugiez [4] and Huima [12]. Discussion with these and other related work can be found in the concluding section.

The rest of the paper is organized as follows. The language, a variant of the spi-calculus with shared-key cryptography, is introduced in Section 2, along with its conventional operational semantics. Section 3 introduces trace analysis. The symbolic operational semantics, and its agreement with the conventional semantics, are discussed in Section 4. Section 5 presents the symbolic method for trace analysis, at the core of which is the concept of *refinement*. For the sake of presentation, the language of Section 2 does not include the restriction operator, whose treatment is postponed to Section 6. The symbolic trace analysis method is applied to an example, the Otway-Rees protocol, in Section 7. Section 8 contains discussion on related work and a few concluding remarks. A separate appendix contains the proof of a major theorem.

2 The language

Syntax (Table 1) We presuppose three countable disjoint sets: \mathcal{L} , \mathcal{N} and \mathcal{V} . The set \mathcal{L} of *labels* is ranged over by $\mathbf{a}, \mathbf{b}, \dots$. The set of \mathcal{N} of *names* is partitioned into two countable sets, a set \mathcal{LN} of *local names* a, b, \dots and a set \mathcal{EN} of *environmental names* $\underline{a}, \underline{b}, \dots$: these sets represent the basic data (keys, nonces, ...) initially known to the process and to the environment, respectively. The set \mathcal{V} of *variables* is ranged over by x, y, \dots . The set $\mathcal{N} \cup \mathcal{V}$ is ranged over by letters u, v, \dots . Names and variables can be used to build compound *messages*, in \mathcal{M} , via shared-key encryption and pairing. In particular, $\{M\}_k$ represents the message obtained by encrypting M (the *argument*) under name k (the *key*), using a shared-key encryption system. We allow pairing and encryptions to be arbitrarily nested, but only permit atomic keys. *Terms* are obtained by closing messages under substitutions (they are just ‘garbage’ that can be generated at run-time, with no semantical significance).

The syntax of *agent expressions*, in \mathcal{A} , is taken essentially from the spi-calculus [1]. Agent **case** $\{M\}_h$ of $\{y\}_k$ in A tries decryption of $\{M\}_h$ using k as a key: if this is possible (that is, if $k = h$), the result of the decryption, M , is bound to y , and the agent proceeds like A , otherwise, the whole expression is stuck. Similarly, **pair** M of $\langle x, y \rangle$ in A tries to split M into two components and call them x and y . A difference from spi/ π -calculus is that, in our case, input and output labels ($\mathbf{a}, \mathbf{b}, \dots$) must not be regarded as a channels (as already noted, we assume just one public network), but rather as ‘tags’ attached to process actions for ease of reference. Also, the only

m, n, \dots	names \mathcal{N}	x, y, \dots	variables \mathcal{V}
a, b, \dots, h, k, \dots	local names \mathcal{LN}	$\underline{a}, \underline{b}, \dots, \underline{h}, \underline{k}, \dots$	environmental names \mathcal{EN}
u, v, \dots	variables or names $\mathcal{V} \cup \mathcal{N}$	$\mathbf{a}, \mathbf{b}, \dots$	labels \mathcal{L}
$M, N ::=$	$u \mid \{M\}_u \mid \langle M, N \rangle$	messages \mathcal{M}	
$\eta, \zeta ::=$	$u \mid \{\zeta\}_\eta \mid \langle \zeta, \eta \rangle$	terms \mathcal{Z}	
$A, B ::=$		agents \mathcal{A}	
	$\mathbf{0}$	(null)	
	$\mathbf{a}(x). A$	(input)	
	$\overline{\mathbf{a}}(\zeta). A$	(output)	
	$\mathbf{case} \zeta \text{ of } \{y\}_\eta \text{ in } A$	(decryption)	
	$\mathbf{pair} \zeta \text{ of } \langle x, y \rangle \text{ in } A$	(selection)	
	$[\zeta = \eta]A$	(matching)	
	$A \parallel B$	(parallel composition)	
The occurrences of variables x and y in (input), (decryption) and (selection) are bound.			

Table 1: Syntax

useful cases for output and decryption are when ζ is a message and η is a name; otherwise the whole agent is stuck.

Given the presence of binders for variables, notions of *free variables*, $\mathbf{v}(A) \subseteq \mathcal{V}$ and *alpha-equivalence* arise as expected. We shall identify alpha-equivalent agent expressions. For any M and u , $[M/u]$ denotes the operation of substituting the free occurrences of u by M . An agent expression A is said to be *closed* or a *process* if $\mathbf{v}(A) = \emptyset$; the set of processes \mathcal{P} is ranged over by P, Q, \dots . Local names and environmental names occurring in A are denoted by $\text{ln}(A)$ and $\text{en}(A)$, respectively. A process P is *initial* if $\text{en}(P) = \emptyset$. These notations are extended to terms, messages, and tuples/string/sets of such objects, component-wise. We shall also use such abbreviations as $\text{ln}(M, P, Q)$ to mean $\text{ln}(M) \cup \text{ln}(P) \cup \text{ln}(Q)$.

Example 2.1 (the wide-mouthed frog protocol, WMF [8]) Principals A and B share two secret keys, k_{AS} and k_{BS} respectively, with a server S . The purpose of the protocol is that of establishing a new secret key k between A and B , which A may use to send a confidential datum d to B . For the sake of simplicity, we suppose that the protocol is always started by A . The protocol and its translation in spi-calculus, WMF , are described below:

$$\begin{aligned}
1. A \longrightarrow S : \{k\}_{k_{AS}} & \quad A \stackrel{\text{def}}{=} \overline{\mathbf{a}}1\langle\{k\}_{k_{AS}}\rangle. \overline{\mathbf{a}}2\langle\{d\}_k\rangle. A' \\
2. S \longrightarrow B : \{k\}_{k_{BS}} & \quad S \stackrel{\text{def}}{=} \mathbf{s}1(x). \mathbf{case} x \text{ of } \{x'\}_{k_{AS}} \text{ in } \overline{\mathbf{s}}2\langle\{x'\}_{k_{BS}}\rangle. \mathbf{0} \\
3. A \longrightarrow B : \{d\}_k & \quad B \stackrel{\text{def}}{=} \mathbf{b}1(y). \mathbf{case} y \text{ of } \{y'\}_{k_{BS}} \text{ in } \mathbf{b}2(z). \mathbf{case} z \text{ of } \{z'\}_{y'} \text{ in } B' \\
WMF & \stackrel{\text{def}}{=} A \parallel S \parallel B.
\end{aligned}$$

Agents A' and B' represent the behaviour of A and B , respectively, after the protocol has been completed. The above description just accounts for a single instance of the protocol. The case of $n > 1$ instances is described by composing n copies of A, S and B in parallel and then appropriately renaming the instance-dependent quantities (k, d, n_A and n_B above).

Operational semantics The semantics of the calculus is given in terms of a transition relation \longrightarrow , which we will sometimes refer to as ‘concrete’ (as opposed to the ‘symbolic’ one we shall introduce later on). We will find it convenient to model a state of the system as a pair (s, P) , where, s records the current environment’s knowledge (i.e. the sequence of messages the environment has “seen” travelling on the network up to that moment) and P is a process. Similarly to [3, 4, 9], we characterize the messages that the environment can *produce* (or deduce) at a given moment,

$\text{(Ax)} \frac{}{S \vdash M} M \in S$	$\text{(ENV)} \frac{}{S \vdash \underline{a}} \underline{a} \in \mathcal{EN}$	
$\text{(PROJ}_1\text{)} \frac{S \vdash \langle M, N \rangle}{S \vdash M}$	$\text{(PROJ}_2\text{)} \frac{S \vdash \langle M, N \rangle}{S \vdash N}$	$\text{(PAIR)} \frac{S \vdash M \quad S \vdash N}{S \vdash \langle M, N \rangle}$
$\text{(DEC)} \frac{S \vdash \{M\}_u \quad S \vdash u}{S \vdash M}$	$\text{(ENC)} \frac{S \vdash M \quad S \vdash u}{S \vdash \{M\}_u}$	

Table 2: Deductive system (\vdash)

starting from the current knowledge s , via a deductive system. These concepts are formalized below.

Definition 2.2 (the deductive system) *Let $S \subseteq_{\text{fin}} \mathcal{M}$ and M a message. We let \vdash be the least binary relation generated by the deductive system in Table 2. If $S \vdash M$ we say that S can produce M .*

Note that, whatever S , the set of messages that S can produce is infinite, due to rules ENV, PAIR and ENC. An *action* is a term of the form $\mathbf{a}\langle M \rangle$ (*input action*) or $\bar{\mathbf{a}}\langle M \rangle$ (*output action*), for \mathbf{a} a label and M a message. The set of actions Act is ranged over by α, β, \dots , while the set Act^* of strings of actions is ranged over by s, s', \dots . String concatenation is written ‘ \cdot ’. We denote by $\text{act}(s)$ and $\text{msg}(s)$ the set of actions and messages, respectively, appearing in s . A string s is *closed* if $\text{v}(s) = \emptyset$ (note that s does not contain binders) and *initial* if $\text{en}(s) = \emptyset$. In what follows, we write $s \vdash M$ for $\text{msg}(s) \vdash M$. We are now set to define *traces*, that is sequences of actions that may result from the interaction between a process and its environment. In traces, each message received by the process (input message) is deducible from the knowledge the environment has previously acquired. In *configurations*, the latter is explicitly recorded.

Definition 2.3 (traces and configurations) *A trace is a closed string $s \in Act^*$ such that for each s_1, s_2 and $\mathbf{a}\langle M \rangle$, if $s = s_1 \cdot \mathbf{a}\langle M \rangle \cdot s_2$ then $s_1 \vdash M$.*

A configuration, written as $\langle s, P \rangle$, is a pair consisting of a trace s and a process P . A configuration is initial if $\text{en}(s, P) = \emptyset$. Configurations are ranged over by $\mathcal{C}, \mathcal{C}', \dots$ \diamond

The concrete transition relation on configurations is defined by the rules in Table 3. Each action taken by the process is recorded in the first component of the configuration (thus a ‘labelled’ transition relation is not needed). Rule (INP) makes the transition relation infinite-branching, as M ranges over the infinite set $\{M : s \vdash M, M \text{ closed}\}$. Another point worth to notice is that decryption with key k is achieved by matching the term to decrypt against a pattern $\{y\}_k$, where y is a fresh variable (rule (CASE)). Finally, no handshake communication is provided (rule (PAR)): all messages go through the environment.

3 Trace analysis

Given a configuration $\langle s, P \rangle$ and a trace s' , we say that $\langle s, P \rangle$ *generates* s' , written $\langle s, P \rangle \searrow s'$, if $\langle s, P \rangle \rightarrow^* \langle s', P' \rangle$ for some P' . Given a string of actions $s \in Act^*$, and actions α and β , we say that α *occurs prior to β in s* if whenever $s = s' \cdot \beta \cdot s''$ then $\alpha \in \text{act}(s')$. We let ρ range

(INP)	$\langle s, \mathbf{a}(x).P \rangle \longrightarrow \langle s \cdot \mathbf{a}\langle M \rangle, P[M/x] \rangle \quad s \vdash M, M \text{ closed}$
(OUT)	$\langle s, \overline{\mathbf{a}}\langle M \rangle.P \rangle \longrightarrow \langle s \cdot \overline{\mathbf{a}}\langle M \rangle, P \rangle$
(CASE)	$\langle s, \mathbf{case} \{ \zeta \}_k \mathbf{of} \{ y \}_k \mathbf{in} P \rangle \longrightarrow \langle s, P[\zeta/y] \rangle$
(SELECT)	$\langle s, \mathbf{pair} \langle \zeta, \eta \rangle \mathbf{of} \langle x, y \rangle \mathbf{in} P \rangle \longrightarrow \langle s, P[\zeta/x, \eta/y] \rangle$
(MATCH)	$\langle s, [\zeta = \zeta]P \rangle \longrightarrow \langle s, P \rangle$
(PAR)	$\frac{\langle s, P \rangle \longrightarrow \langle s', P' \rangle}{\langle s, P \parallel Q \rangle \longrightarrow \langle s', P' \parallel Q \rangle}$
<i>plus symmetric version of (PAR).</i>	

Table 3: Transition relation (\longrightarrow).

over ground substitutions, i.e. finite maps from a set $\text{dom}(\rho) \subseteq \mathcal{V}$ to closed messages; $t\rho$ denotes the result of replacing each $x \in \text{v}(t) \cap \text{dom}(\rho)$ by $\rho(x)$. The properties of configurations we are interested in are defined below.

Definition 3.1 (properties and satisfaction relation) *Let α and β be actions, with $\text{v}(\alpha) \subseteq \text{v}(\beta)$, and let s be a trace. We write $s \models \alpha \leftrightarrow \beta$, or s satisfies $\alpha \leftrightarrow \beta$, if for each ground substitution ρ it holds that $\alpha\rho$ occurs prior to $\beta\rho$ in s . We say that a configuration \mathcal{C} satisfies $\alpha \leftrightarrow \beta$, and write $\mathcal{C} \models \alpha \leftrightarrow \beta$, if all traces generated by \mathcal{C} satisfy $\alpha \leftrightarrow \beta$. \diamond*

Note that the variables in α and β can be thought of as being universally quantified (so any consistent renaming of these variables does not change the set of traces and configurations that satisfy $\alpha \leftrightarrow \beta$). In practice, the scheme $\alpha \leftrightarrow \beta$ permits formalizing all forms of authentication in Lowe’s hierarchy [14], except the most demanding one (but can be easily modified to include this one as well). As we shall see, the scheme also permits expressing secrecy as a reachability property, in the style of [3]. To this purpose, it is convenient to assume a fixed ‘absurd’ action \perp that is nowhere used in agent expressions. Thus the formula $\perp \leftrightarrow \alpha$ expresses that action α should never take place, and can be used to encode reachability.

Example 3.2 (authentication and secrecy in WMF) We discuss the use of properties on the simple protocol WMF (Example 2.1), but our considerations are indeed quite general.

A property of WMF that one would like to check is the following: if B accepts as ‘good’ a datum d encrypted under key k (step 3), then this message has actually been sent by A . This is a form of *authentication*. In order to formalize this particular property, we make B explicitly declare if, at the end of the protocol, a particular message has been accepted. That is, we consider the process $B_{auth} \stackrel{\text{def}}{=} \mathbf{b1}(y). \mathbf{case} y \mathbf{of} \{ y' \}_{k_{BS}} \mathbf{in} \mathbf{b2}(z). \mathbf{case} z \mathbf{of} \{ z' \}_{y'} \mathbf{in} (B' \parallel \overline{\mathbf{accept}}\langle z \rangle). \mathbf{0}$ instead of B , and $WMF_{auth} \stackrel{\text{def}}{=} A \parallel S \parallel B_{auth}$ instead of WMF . We have to check that, in every trace of WMF_{auth} , every $\overline{\mathbf{accept}}$ is preceded by the corresponding $\overline{\mathbf{a2}}$. More formally, we have to check that $\langle \epsilon, WMF_{auth} \rangle \models \overline{\mathbf{a2}}\langle t \rangle \leftrightarrow \overline{\mathbf{accept}}\langle t \rangle$ (where t is any variable).

Another important property is *secrecy*: the environment should never learn the confidential datum d . Following [3], we can formalize this property by considering a version of WMF that also includes a ‘guardian’ listening to the public network: $WMF_{secr} \stackrel{\text{def}}{=} WMF \parallel \mathbf{guard}(x). \mathbf{0}$. Evidently, WMF generates a trace s s.t. $s \vdash d$ (i.e. the environment learns d) if and only if $\langle \epsilon, WMF_{secr} \rangle$ generates a trace containing action $\mathbf{guard}\langle d \rangle$. Thus, we have to check that action $\mathbf{guard}\langle d \rangle$ never

takes place, i.e. that $\langle \epsilon, WMF_{secr} \rangle \models \perp \leftrightarrow \mathbf{guard}\langle d \rangle$. Similarly, we can verify that the session key k is never compromised by checking that $\langle \epsilon, WMF_{secr} \rangle \models \perp \leftrightarrow \mathbf{guard}\langle k \rangle$. Using small transformations of this kind, we can, e.g., check that compromise of the session key k does not compromise the long-term keys k_{AS}, k_{BS} .

4 Symbolic semantics

‘Concrete’ traces and configurations can be given a symbolic counterpart, which may contain free variables.

Definition 4.1 (symbolic traces and configurations) *A symbolic trace is string $\sigma \in Act^*$ such that: (a) $\text{en}(\sigma) = \emptyset$, and (b) for each $\sigma_1, \sigma_2, \alpha$ and x , if $\sigma = \sigma_1 \cdot \alpha \cdot \sigma_2$ and $x \in \text{v}(\alpha) - \text{v}(\sigma_1)$ then α is an input action. Symbolic traces are ranged over by σ, σ', \dots*

A symbolic configuration, written $\langle \sigma, A \rangle_s$, is a pair composed by a symbolic trace σ and an agent A , such that $\text{en}(A) = \emptyset$ and $\text{v}(A) \subseteq \text{v}(\sigma)$. \diamond

Note that, due to condition (b) in the definition, e.g. $\bar{\mathbf{a}}\langle x \rangle \cdot \mathbf{a}\langle \{h\}_x \rangle$ is not a symbolic trace, while $\mathbf{a}\langle \{h\}_x \rangle \cdot \bar{\mathbf{a}}\langle x \rangle$ is. Let us now recall some standard terminology about substitutions. A substitution θ is a finite partial map from \mathcal{V} to \mathcal{M} and, for any object (i.e. variable, message, process, trace, ...) t , we denote by $t\theta$ the result of applying θ to t . A substitution θ is a *unifier of t_1 and t_2* if $t_1\theta = t_2\theta$. We denote by $\text{mgu}(t_1, t_2)$ a chosen *most general unifier (mgu)* of t_1 and t_2 , that is a unifier θ such that any other unifier can be written as a composition of substitutions $\theta\theta'$ for some θ' .

The transition relation on symbolic configurations, \longrightarrow_s , is defined by the rules in Table 4. There, a function $\text{new}_v(\cdot)$ is assumed such that, for any given $V \subseteq_{\text{fin}} \mathcal{V}$, $\text{new}_v(V)$ is a variable not in V . Note that, differently from the concrete semantics, input variables are not instantiated immediately in the input rule (INP_s). Rather, constraints on these variables are added as soon as they are needed, and recorded via mgu’s. This may occur due to rules (CASE_s), (SELECT_s) and (MATCH_s). In the following example, after the first step, variable x gets instantiated to name b due to a (MATCH_s)-reduction:

$$\langle \epsilon, \mathbf{a}\langle x \rangle.[x = b]P \rangle_s \longrightarrow_s \langle \mathbf{a}\langle x \rangle, [x = b]P \rangle_s \longrightarrow_s \langle \mathbf{a}\langle b \rangle, P[b/x] \rangle_s .$$

The side condition on B' in (PAR_s) ensures that constraints are propagated across parallel components sharing variables, like in the following (MATCH_s)-reduction: $\langle \sigma, [x = M]A \parallel \bar{\mathbf{a}}\langle x \rangle.B \rangle_s \longrightarrow_s \langle \sigma[M/x], A[M/x] \parallel \bar{\mathbf{a}}\langle M \rangle.B[M/x] \rangle_s$.

Whenever $\langle \sigma, A \rangle_s \xrightarrow{s}^* \langle \sigma', A' \rangle_s$ for some A' , we say that $\langle \sigma, A \rangle_s$ *symbolically generates* σ' , and write $\langle \sigma, A \rangle_s \searrow_s \sigma'$. The relation \longrightarrow_s is finite-branching. This implies that each configuration generates a finite number of symbolic traces. It is important to stress that many symbolic traces are in fact ‘garbage’ – jumbled sequences of actions that cannot be instantiated to give a concrete trace. This is the case, e.g., for the trace $\mathbf{a}\langle \{x\}_k \rangle \cdot \bar{\mathbf{a}}\langle x \rangle$, which is symbolically generated by $\langle \epsilon, P \rangle_s$, where $P \stackrel{\text{def}}{=} \mathbf{a}\langle y \rangle.\mathbf{case} y \text{ of } \{x\}_k \text{ in } \bar{\mathbf{a}}\langle x \rangle.\mathbf{0}$. To state soundness and completeness of \longrightarrow_s w.r.t. \longrightarrow , we need a notion of consistency for symbolic traces, given below.

Definition 4.2 (solutions of symbolic traces) *Given a symbolic trace σ and a ground substitution ρ , we say that ρ satisfies σ if $\sigma\rho$ is a trace. In this case, we also say that $\sigma\rho$ is a solution of σ , and that σ is consistent. \diamond*

Theorem 4.3 (soundness and completeness) *Let \mathcal{C} be an initial configuration and s a trace. Then $\mathcal{C} \searrow_s s$ if and only if there is σ s.t. $\mathcal{C} \searrow_s \sigma$ and s is a solution of σ .*

(INP _s)	$\langle \sigma, \mathbf{a}(x).A \rangle_s \longrightarrow_s \langle \sigma \cdot \mathbf{a}(x), A \rangle_s$	
(OUT _s)	$\langle \sigma, \bar{\mathbf{a}}(M).A \rangle_s \longrightarrow_s \langle \sigma \cdot \bar{\mathbf{a}}(M), A \rangle_s$	
(CASE _s)	$\langle \sigma, \text{case } \{\zeta\}_v \text{ of } \{x\}_u \text{ in } A \rangle_s \longrightarrow_s \langle \sigma\theta, A\theta \rangle_s$	$\theta = \text{mgu}(\{\zeta\}_v, \{x\}_u)$
(SELECT _s)	$\langle \sigma, \text{pair } \zeta \text{ of } \langle x, y \rangle \text{ in } A \rangle_s \longrightarrow_s \langle \sigma\theta, A\theta \rangle_s$	$\theta = \text{mgu}(\zeta, \langle x, y \rangle)$
(MATCH _s)	$\langle \sigma, [\zeta = \eta]A \rangle_s \longrightarrow_s \langle \sigma\theta, A\theta \rangle_s$	$\theta = \text{mgu}(\zeta, \eta)$
(PAR _s)	$\frac{\langle \sigma, A \rangle_s \longrightarrow_s \langle \sigma', A' \rangle_s}{\langle \sigma, A \parallel B \rangle_s \longrightarrow_s \langle \sigma', A' \parallel B' \rangle_s}$	$B' = B\theta \text{ if } \sigma' = \sigma\theta, B' = B \text{ if } \sigma' = \sigma$
<p><i>plus symmetric version of (PAR_s). In the rules, it is assumed that:</i></p> <ul style="list-style-type: none"> (i) $x = \text{new}_v(V)$ – where V is the set of free variables in the source configuration, (ii) $y = \text{new}_v(V \cup \{x\})$ and (iii) $\text{msg}(\sigma)\theta \subseteq \mathcal{M}$. 		

Table 4: Symbolic transition relation (\longrightarrow_s)

PROOF: By transition induction on \longrightarrow and \longrightarrow_s , and then by induction on the length of traces. \square

Any given configuration generates only finitely many symbolic traces. Thus, by the previous theorem, the task of checking $\mathcal{C} \models \alpha \leftrightarrow \beta$ is reduced to analysing each of these symbolic traces in turn. To do this, we need at least a method to tell whether any given symbolic trace is consistent or not. More precisely, the previous theorem reduces the problem $\mathcal{C} \models \alpha \leftrightarrow \beta$ to the following, that will be faced in the next section.

Symbolic trace analysis problem (STAP) Given actions α, β ($v(\alpha) \subseteq v(\beta)$) and a symbolic trace σ , check whether or not each solution s of σ satisfies $\alpha \leftrightarrow \beta$.

We write $\sigma \models \alpha \leftrightarrow \beta$ if the answer to STAP with σ, α and β is ‘yes’. Note that STAP is a non-trivial problem: one has to consider *every* solution of σ , and there may be infinitely many of them.

5 Refinement

As a first step towards devising a method for STAP, let us consider the simpler problem of checking consistency of a symbolic trace. Existence of solutions (i.e. ground instances that are traces) of a symbolic trace σ depends on the form of input actions in σ . For example, the symbolic trace $\sigma_0 = \bar{\mathbf{a}}(h) \cdot \mathbf{b}(\{x\}_k)$ ($h \neq k$) has no solution, because no matter which m is substituted for x , we have $\{h\} \not\vdash \{m\}_k$. On the contrary, in $\sigma_1 = \bar{\mathbf{c}}(k) \cdot \mathbf{c}(\{x\}_k)$, instantiating x with any environmental name $\underline{a} \in \mathcal{EN}$ will give a solution, because $\{k\} \vdash \{\underline{a}\}_k$. Yet a different case is $\sigma_2 = \bar{\mathbf{c}}(\{a\}_k) \cdot \bar{\mathbf{c}}(\{b\}_k) \cdot \mathbf{c}(\{x\}_k)$. Since $\{\{a\}_k, \{b\}_k\} \not\vdash k$, there are only two ways of getting a solution of σ_2 : to unify $\{x\}_k$ with either $\{a\}_k$ or $\{b\}_k$. These examples suggest that it should be possible to check consistency of a symbolic trace by gradually instantiating it until a trace is obtained. We will call this process *refinement*. In order to formalize this concept, we first need to lift the definition of ‘trace’ to the non-ground case. This requires a few more notations and concepts.

In refinement, we shall consider both ordinary variables and *marked* variables; roughly, the latter can only be instantiated to messages that the environment can produce. This is made precise in the sequel. We consider a new set $\widehat{\mathcal{V}}$ of marked variables, which is in bijection with \mathcal{V} via a mapping $\hat{\cdot}$: thus variables x, y, z, \dots have marked versions $\hat{x}, \hat{y}, \hat{z}, \dots$. Marked messages are messages that may also contain marked variables, and marked symbolic traces are defined similarly. The deduction relation \vdash is extended to marked messages by adding the new axiom

$$\text{(MVAR)} \quad \frac{}{S \vdash \hat{x}} \quad \hat{x} \in \widehat{\mathcal{V}}$$

to the system of Table 2. For any \hat{x} and any sequence σ , we denote by $\sigma \setminus \hat{x}$ the longest prefix of σ not containing \hat{x} . The satisfaction relation is extended to marked symbolic traces as follows:

Definition 5.1 *Let σ be a marked symbolic trace and ρ be a ground substitution. We say that ρ satisfies σ if $\sigma\rho$ is a trace and, for each $\hat{x} \in \text{v}(\sigma)$, it holds that $(\sigma \setminus \hat{x})\rho \vdash \rho(\hat{x})$. We also say that $\sigma\rho$ is a solution of σ , and that σ is consistent.* \diamond

The terminology introduced above agrees with Definition 4.2 when σ does not contain marked variables. We can give now the definition of *solved form*, that lifts the definition of trace to the non-ground case (note that this definition is formally the same as Def. 2.3)

Definition 5.2 (solved forms) *Let σ be a marked symbolic trace. We say σ is in solved form (sf) if for every $\sigma_1, \mathbf{a}\langle M \rangle$ and σ_2 s.t. $\sigma = \sigma_1 \cdot \mathbf{a}\langle M \rangle \cdot \sigma_2$ it holds that $\sigma_1 \vdash M$.* \diamond

Solved forms are consistent: the next lemma gives us a specific way to instantiate a solved form so as to get a trace.

Lemma 5.3 *Let σ be in solved form and let ρ be any substitution from $\text{v}(\sigma)$ to \mathcal{EN} . Then ρ satisfies σ .*

A key concept of refinement is that of decomposing a message into its irreducible components, those that cannot be further split or decrypted using the knowledge of a given σ .

Definition 5.4 (decomposition of messages) *Let σ be a marked symbolic trace.*

- $\mathbf{I}(\sigma) \stackrel{\text{def}}{=} \{M : \sigma \vdash M \text{ and either } M \in \mathcal{LN} \cup \mathcal{V} \text{ or } M = \{N\}_u \text{ for some } u \text{ s.t. } \sigma \not\vdash u\}$.
- The set $[M]_\sigma$ is defined by induction on M as follows:

$$\begin{aligned} [u]_\sigma &= \{u\} - (\widehat{\mathcal{V}} \cup \mathcal{EN}) \\ [(M, N)]_\sigma &= [M]_\sigma \cup [N]_\sigma \\ [\{M\}_u]_\sigma &= \begin{cases} \{\{M\}_u\} & \text{if } \sigma \not\vdash u \\ [M]_\sigma & \text{if } \sigma \vdash u \end{cases} \end{aligned}$$

\diamond

The irreducible components of σ , $\mathbf{I}(\sigma)$, are the building blocks of messages that can be produced by σ . This is the content of the next proposition, that makes the relationship among $\mathbf{I}(\sigma)$, $[M]_\sigma$ and \vdash precise.

Proposition 5.5 *Let σ be a marked symbolic trace. Then $\sigma \vdash M$ if and only if $[M]_\sigma \subseteq \mathbf{I}(\sigma)$.*

Let σ be a marked symbolic trace, and assume $\sigma = \sigma' \cdot \mathbf{a}\langle M \rangle \cdot \sigma''$, where σ' is the longest prefix of σ that is in solved form. Assume $N \in [M]_{\sigma'} - \mathbf{I}(\sigma')$.

(REF ₁)	$\frac{N \notin \mathcal{V}, \quad N \in_{\theta} \mathbf{I}(\sigma'), \quad \tilde{y} = \{x \mid \hat{x} \in v(\sigma) \text{ and } (\sigma\theta) \setminus \hat{x} \text{ is shorter than } \sigma \setminus \hat{x}\}}{\sigma \succ \sigma\theta[\tilde{y}/\hat{y}]}$
(REF ₂)	$\frac{N = x \text{ or } N = \{N'\}_x}{\sigma \succ \sigma[\hat{x}/x]}$

Table 5: Refinement (\succ)

There are two points worth noting with respect to the above proposition. First, $\mathbf{I}(\sigma)$ is finite and can be easily computed by an iterative procedure, thus the proposition gives us an effective method to decide $\sigma \vdash M$; this also implies that the set of solved forms is decidable. Second, the proposition suggests a strategy for refining a generic σ to a solved form: for any input message M in σ , one tries to make the condition $[M]_{\sigma'} \subseteq \mathbf{I}(\sigma')$ true, for the appropriate prefix σ' of σ . We are now ready to define refinement formally. In the sequel, we shall use the following notations. We write ‘ $t \in_{\theta} S$ ’ for: there is $t' \in S$ s.t. $\theta = \text{mgu}(t, t')$; when \tilde{y} is a set of variables, we denote by $[\tilde{y}/\hat{y}]$ the substitution that for each $x \in \tilde{y}$ maps \hat{x} to x .

Definition 5.6 (refinement) *We let refinement, written \succ , be the least binary relation over marked symbolic traces generated by the two rules in Table 5. \diamond*

Rule (REF₁) implements the basic step: an element N in the decomposition of M gets instantiated, via θ , to an irreducible component of some past message (to be found in $\mathbf{I}(\sigma')$). E.g., consider again the above $\sigma_2 = \bar{c}\langle\{a\}_k\rangle \cdot \bar{c}\langle\{b\}_k\rangle \cdot c\langle\{x\}_k\rangle$: its possible refinements are $\sigma_2 \succ \sigma_2[a/x]$ and $\sigma_2 \succ \sigma_2[b/x]$, and the refined traces are in sf. By rule (REF₂), one may deem a variable marked. Sometimes marked variables need to be ‘unmarked’ back to variables, and this is achieved via the renaming $[\tilde{y}/\hat{y}]$ in (REF₁).¹

Refinement is repeatedly applied until some solved form is reached. It is important to realize that the reflexive and transitive closure $(\succ)^*$ is a non-deterministic relation, and that not all sequences of refinement lead to a solved form. However, the set of possible solved forms reachable from σ completely characterizes the set of solutions of σ . Formally, for any symbolic trace σ , we let $\mathbf{SF}(\sigma) \stackrel{\text{def}}{=} \{\sigma' \mid \sigma (\succ)^* \sigma' \text{ and } \sigma' \text{ is in sf}\}$. Then we have the following theorem (The proof is reported in the appendix).

Theorem 5.7 (characterization of solutions) *Let σ be a symbolic trace and s a trace. Then s is a solution of σ if and only if s is a solution of some $\sigma' \in \mathbf{SF}(\sigma)$.*

By the above theorem and Lemma 5.3, we obtain:

Corollary 5.8 *A symbolic trace σ is consistent if and only if $\mathbf{SF}(\sigma) \neq \emptyset$.*

Note that $\mathbf{SF}(\sigma)$ can be effectively computed, and is always finite, as: (a) \succ is finitely-branching relation, and (b) infinite sequences of refinement steps cannot arise. As to the latter point, note that, since each (REF₁)-step eliminates at least one variable, any sequence of refinement steps can

¹Unmarking of \hat{x} occurs in a (REF₁)-step if the prefix $\sigma \setminus \hat{x}$ gets shorter, like in: $\bar{a}\langle\{a\}_k\rangle \cdot a\langle\hat{z}\rangle \cdot \bar{a}\langle\{\hat{z}\}_h\rangle \cdot \bar{a}\langle k \rangle \cdot a\langle\{\hat{x}\}_g\rangle \cdot a\langle\{\{\hat{x}\}_g\}_h\rangle \succ \bar{a}\langle\{a\}_k\rangle \cdot a\langle\{x\}_y\rangle \cdot \bar{a}\langle\{\{x\}_y\}_h\rangle \cdot \bar{a}\langle k \rangle \cdot a\langle\{x\}_y\rangle \cdot a\langle\{\{x\}_y\}_h\rangle$, where $\{\{\hat{x}\}_g\}_h$ gets unified with $\{\hat{z}\}_h$.

contain only finitely many (REF₁)-steps, after the last of which rule (REF₂) can only be applied a finite number of times. Thus, computing $\mathbf{SF}(\sigma)$ is a method to decide consistency of symbolic traces. This also suggests a method for solving STAP. As an example, suppose that we want to check the property $\sigma \models \perp \leftrightarrow \alpha$, that is, no solution of σ contains an instance of α . Then we can proceed as follows: for each action γ in σ , we check whether there is a mgu θ that unifies γ and α ; if such a θ does not exist, or if it exists but $\sigma\theta$ is not consistent (i.e. $\mathbf{SF}(\sigma\theta) = \emptyset$, Corollary 5.8), then the property is true, otherwise it is not. Considering the general case $\alpha \leftrightarrow \beta$ leads us to the next theorem, which gives us an effective method to check $\sigma \models \alpha \leftrightarrow \beta$. Its proof relies on Theorem 5.7 and on Lemma 5.3, plus routine calculations on mgu's.

Theorem 5.9 (a method for STAP) *Let σ be a symbolic trace and let $pr = \alpha \leftrightarrow \beta$, where $v(\alpha) \subseteq v(\beta)$ and $v(\beta) \cap v(\sigma) = \emptyset$. Then $\sigma \models pr$, if and only if the following is true: for each θ such that $\alpha \in_{\theta} \text{act}(\sigma)$ and for each $\sigma' \in \mathbf{SF}(\sigma\theta)$, say $\sigma' = \sigma\theta\theta'$, it holds that $\alpha\theta\theta'$ occurs prior to $\beta\theta\theta'$ in σ' .*

The above theorem immediately yields decidability of STAP, because there are finitely many mgu's θ to consider (at most one for each action in σ), and $\mathbf{SF}(\sigma)$ can be effectively computed. This result lifts of course to configurations.

Corollary 5.10 (decidability) *Let \mathcal{C} be an initial configuration and $\alpha \leftrightarrow \beta$ be a property. It is decidable whether $\mathcal{C} \models \alpha \leftrightarrow \beta$ or not.*

PROOF: Compute $\{\sigma \mid \mathcal{C} \text{ symbolically generates } \sigma\}$, which is finite, and then check whether or not for each σ in this set it is the case that $\sigma \models \alpha \leftrightarrow \beta$, which can be effectively done. The thesis is a consequence of Theorem 4.3. \square

In a practical implementation, rather than generating the whole set of symbolic traces of a given configuration and then check the property, it is more convenient to generate the symbolic traces ‘on-the-fly’: every last symbolic action γ taken by the configuration is compared against action β of the property $\alpha \leftrightarrow \beta$; only when β and γ are unifiable the refinement procedure $\mathbf{SF}(\cdot)$ is invoked. This also the way our implementation works [6]. If $\langle \sigma, P \rangle \not\models \alpha \leftrightarrow \beta$, this method allows one to compute a trace violating the property, that is, an attack on $\langle \sigma, P \rangle$. We will see a detailed example of this in Section 7.

6 Restriction

We consider extending the base language via the *restriction* operator $(\mathbf{new} \ a)A$, where $a \in \mathcal{LN}$ and A an agent; $(\mathbf{new} \ a)$ is binder for name a . The intended meaning of $(\mathbf{new} \ a)A$ is that a new name a is created, which is private to A . The concrete and symbolic rules for restriction are given below. A function $\mathbf{new}_{\mathcal{LN}}(\cdot)$ is assumed that, for any set of names $V \subseteq_{\text{fin}} \mathcal{LN}$, yields a local name $a \notin V$.

$$\begin{aligned} (\text{NEW}) \quad & \langle s, (\mathbf{new} \ a)P \rangle \longrightarrow \langle s, P \rangle \quad a = \mathbf{new}_{\mathcal{LN}}(V) \\ (\text{NEW}_s) \quad & \langle \sigma, (\mathbf{new} \ a)A \rangle_s \longrightarrow_s \langle \sigma, A \rangle_s \quad a = \mathbf{new}_{\mathcal{LN}}(V) \end{aligned}$$

In both rules, V is the set of local names occurring free in the source configuration. Note that the side-condition on name a is always met modulo alpha-renaming. A change is required in the rules for parallel composition $A \parallel B$, both in the concrete and in the symbolic case: in the conclusion, an additional renaming $[b/a]$ (where $a = \mathbf{new}_{\mathcal{LN}}(\text{fn}(\sigma, A))$ and $b = \mathbf{new}_{\mathcal{LN}}(\text{fn}(\sigma, A \parallel B))$) is applied onto the target configuration: this prevents a new name a possibly created by a (NEW)-transition of A from clashing with free occurrences of a in B (this is just the side-condition of the rule (PAR) of the π -calculus [18] rephrased in our language).

We remark that, in the absence of recursion/replication, restriction is a handy renaming device, but has little semantical relevance, as its effect can be achieved by a statical renaming of names. On the other hand, adding recursion/replication makes trace analysis undecidable, as already noted.

7 An example

We consider a variant of the Otway-Rees protocol [8], which is designed to let any two agents A and B establish a private session key k , with the help of an authentication server S . The protocol may be described informally as follows:

1. $A \rightarrow B$: $n_A, A, B, \{n_A, A, B\}_{k_A}$
2. $B \rightarrow S$: $n_A, A, B, \{n_A, A, B\}_{k_A}, n_B, \{n_A, A, B\}_{k_B}$
3. $S \rightarrow B$: $n_A, \{k, n_A\}_{k_A}, \{k, n_B\}_{k_B}$
4. $B \rightarrow A$: $n_A, \{k, n_A\}_{k_A}$.

Here, k_A and k_B represent the long-term keys that A and B , respectively, share with the server S , while n_A and n_B are fresh nonces. The formal description of the protocol is in Table 6. We assume three participants named A , B and I . I is a malicious insider: in other words, the hostile environment has registered itself as a legitimate participant having name I and long-term key k_I . The identity of each participant is known to the others, and each participants may act as the initiator or as the responder of a given round. We have added an action $\overline{\text{accept}}_{in.re}$ that the initiator in performs when he believes to have successfully completed a round with a responder re . To make the description more readable some obvious meta-notation is used. In particular, we abbreviate $A_1 \parallel \dots \parallel A_n$ as $\prod_{i \in \{1, \dots, n\}} A_i$. We have also eliminated parentheses in pairs, and used abbreviations like $\mathbf{a}(\{M, N\}_u)$ in place of $\mathbf{a}(x)$. $\text{case } x \text{ of } \{x'\}_u \text{ in (pair } x' \text{ of } \langle y, z \rangle \text{ in } ([y = M][z = N]A))$. Subscripts in input variables are meant to remind the reader their expected values, bearing in mind that ‘ in ’ stands for ‘initiator’, ‘ re ’ for ‘responder’, ‘ n ’ for ‘nonce’, ‘ c ’ for ‘certificate’ (messages of the form $\{\dots\}_{K_x}$), ‘ k ’ for ‘key’ and ‘ f ’ for ‘forward’. For instance, the expected value for $y_{re.n}$ is the responder’s nonce (n_B in the informal description above). The initial environment’s knowledge is given by the message in α_0 .

This version of the procol is subject to a subtle form of attack [19]. A message accepted by A acting as the initiator at step 4 should originate from B , if B is the intended responder. Formally, we wish that $\mathcal{C} \models \overline{\mathbf{b}4}\langle x \rangle \leftrightarrow \overline{\text{accept}}_{A.B}\langle x \rangle$ holds true. However, this is not the case: the intruder I can place itself between A and the rest of the world and eventually convince A to have completed a run of the protocol with B . As a result, A believes he shares a session key with B , while the session key is actually shared with the intruder. The attack involves two interleaved executions of the protocol, one in which A acts as the initiator and one in which A acts as the responder. In more detail, \mathcal{C} generates the symbolic trace $\sigma = \alpha_0 \cdot \alpha_1 \dots \alpha_7$, where:

$$\begin{aligned}
\alpha_1 &= \overline{\mathbf{a}1} \langle n_A, A, B, \{n_A, A, B\}_{k_A} \rangle & \alpha_5 &= \overline{\mathbf{s}2} \langle z_{in.n}, \{k, z_{in.n}\}_{k_I}, \{k, z_{re.n}\}_{k_A} \rangle \\
\alpha_2 &= \mathbf{a}'1 \langle y_{in.n}, y_{in}, A, y_{in.c} \rangle & \alpha_6 &= \mathbf{a}2 \langle n_A, \{x_k, n_A\}_{k_A} \rangle \\
\alpha_3 &= \overline{\mathbf{a}'2} \langle y_{in.n}, y_{in}, A, y_{in.c}, n'_A, \{y_{in.n}, y_{in}, A\}_{k_A} \rangle & \alpha_7 &= \overline{\text{accept}}_{A,B} \langle n_A, \{x_k, n_A\}_{k_A} \rangle \\
\alpha_4 &= \mathbf{s}1 \langle z_{in.n}, I, A, \{z_{in.n}, I, A\}_{k_I}, z_{re.n}, \{z_{in.n}, I, A\}_{k_A} \rangle.
\end{aligned}$$

To see that $\sigma \not\models \overline{\mathbf{b}4}\langle x \rangle \leftrightarrow \overline{\text{accept}}_{A.B}\langle x \rangle$, it is sufficient inspect the solved form σ' obtained by refining σ : the key steps are unifying $\{z_{in.n}, I, A\}_{k_A}$ in α_4 with $\{y_{in.n}, y_{in}, A\}_{k_A}$ in α_3 , and $\{x_k, n_A\}_{k_A}$ in α_6 with $\{k, z_{re.n}\}_{k_A}$ in α_5 . By instantiating the variables remaining in σ' to an arbitrary $\underline{a} \in \mathcal{EN}$ we can get a ground trace s generated by \mathcal{C} in the concrete semantics

$$\begin{aligned}
s &= \overline{\mathbf{e}}\langle A, B, I, k_I \rangle \cdot \overline{\mathbf{a}1}\langle n_A, A, B, \{n_A, A, B\}_{k_A} \rangle \cdot \mathbf{a}'1\langle \underline{a}, I, A, \underline{a} \rangle \cdot \overline{\mathbf{a}'2}\langle \underline{a}, I, A, \underline{a}, n'_A, \{\underline{a}, I, A\}_{k_A} \rangle \cdot \\
&\quad \mathbf{s}1\langle \underline{a}, I, A, \{\underline{a}, I, A\}_{k_I}, n_A, \{\underline{a}, I, A\}_{k_A} \rangle \cdot \overline{\mathbf{s}2}\langle \underline{a}, \{k, \underline{a}\}_{k_I}, \{k, n_A\}_{k_A} \rangle \cdot \mathbf{a}2\langle n_A, \{k, n_A\}_{k_A} \rangle \cdot \\
&\quad \overline{\text{accept}}_{A,B}\langle n_A, \{k, n_A\}_{k_A} \rangle
\end{aligned}$$

that violates the property and shows what has happened: the environment uses the nonce n_A obtained from action $\overline{\mathbf{a}1}$, and the certificate $\{\underline{a}, I, A\}_{k_A}$ from $\overline{\mathbf{a}'2}$ to make S produce a session key

A	$\stackrel{\text{def}}{=} \Pi_{X \in \{I, B\}} (\text{new } n_A) \left(\overline{\mathbf{a}1} \langle n_A, A, X, \{n_A, A, X\}_{k_A} \rangle. \mathbf{a}2 \langle n_A, \{x_k, n_A\}_{k_A} \rangle. \right.$ $\left. \overline{\text{accept}}_{A.X} \langle n_A, \{n_A, x_k\}_{k_A} \rangle. \mathbf{0} \right)$ \parallel $(\text{new } n'_A) \left(\mathbf{a}'1 \langle y_{in.n}, y_{in}, A, y_{in.c} \rangle. \overline{\mathbf{a}'2} \langle y_{in.n}, y_{in}, A, y_{in.c}, n'_A, \{y_{in.n}, y_{in}, A\}_{k_A} \rangle. \right.$ $\left. \mathbf{a}'3 \langle y_{in.n}, y_f, \{y_k, n'_A\}_{k_A} \rangle. \overline{\mathbf{a}'4} \langle y_{in.n}, y_f \rangle. \mathbf{0} \right)$
B	is like A , but ‘ b ’ replaces ‘ a ’ in labels, and the roles of ‘ B ’ and ‘ A ’ are swapped.
S	$\stackrel{\text{def}}{=} \mathbf{s}1 \langle z_{in.n}, z_{in}, z_{re}, z_{in.c}, z_{re.n}, z_{re.c} \rangle. \Pi_{X \in \{A, B, I\}} [z_{in} = X] \Pi_{Y \in \{A, B, I\} - \{X\}} [z_{re} = Y]$ $\left(\text{case } z_{in.c} \text{ of } \{z_{in.n}, X, Y\}_{k_X} \text{ in case } z_{re.c} \text{ of } \{z_{in.n}, X, Y\}_{k_Y} \text{ in} \right.$ $\left. (\text{new } k) \overline{\mathbf{s}2} \langle z_{in.n}, \{k, z_{in.n}\}_{k_X}, \{k, z_{in.n}\}_{k_Y} \rangle. \mathbf{0} \right)$
OR	$\stackrel{\text{def}}{=}} (\text{new } k_A, k_B) (A \parallel S \parallel B)$
\mathcal{C}	$\stackrel{\text{def}}{=}} \langle \alpha_0, OR \rangle$ where $\alpha_0 \stackrel{\text{def}}{=}} \overline{\mathbf{e}} \langle A, B, I, k_I \rangle$ (A, B, I and k_I distinct local names)

Table 6: The Otway-Rees protocol

k , which A believes is in response to his message $\overline{\mathbf{a}1}$ (hence shared with B), but which is actually shared with the environment. We have found the symbolic trace σ above manually, being guided by previous knowledge of the protocol and of the attack (see e.g. [19]). However, automatic exhaustive search of the *symbolic* traces generated by \mathcal{C} is perfectly feasible. When fed with this example, our tool found the attack in a few seconds [6].

8 Conclusions

We have presented a symbolic method for analysing cryptographic protocols. The method is well suited for an efficient mechanization. The word ‘efficient’ should be taken with a grain of salt here, because the trace analysis problem is NP-hard [4, 10], thus pathological examples are unavoidable. However, we expect the method to perform well in practical cases; this is further discussed below. Experiments conducted with a preliminary, non optimized implementation have given encouraging results [6]. Developments in the near future include an optimized implementation of the method and extension of the present results to other cryptographic primitives, like public key and hashing: this should not present conceptual difficulties, though we have not checked the details yet.

Approaches based on symbolic analysis have also been explored by Huima in [12] and Amadio and Lugiez [4]. In [12], Huima presents a symbolic semantics by which the execution of a protocol generates a set of equational constraints; only an informal description is given of the kind of equational rewriting needed to solve these constraints. Amadio and Lugiez in [4] consider a variant of the spi-calculus equipped with a symbolic semantics. Similarly to Huima’s, their symbolic semantics generates equational constraints of a special form, rather than unifiers. The (rather complex) constraint-solving procedure is embedded into symbolic execution, and uses a brute-force method to resolve variables in key position (all possible instantiations of variables to names that are around are tried). These factors have a relevant impact on the size of the symbolic model. On the contrary, in our case symbolic execution and consistency check are kept separate, and this permits to keep the size of the model to a minimum. As explained in Section 5, the consistency check procedure (refinement) is invoked only when necessary, and does not use brute-force instantiation. Finally, Amadio and Lugiez encode authentication via reachability: this may add to the complexity of their method.

Model checking [9, 13, 20, 22] and theorem proving [17, 19] seem to be among the most successful approaches to the formal analysis of security protocols. As Paulson has pointed out [19], theorem proving is intuitive, but, within it, verification is not fully automated and general completeness

results are difficult to establish. On the contrary, model checking is automatic, but suffers from the state explosion problem, which requires the model to be cut down to a convenient finite size. Our paper might be regarded as an attempt at bridging the two approaches. We extract the unification mechanism underlying theorem proving and bring it on the ground of a process language (a variant of the spi-calculus) that naturally supports a notion of variable binding. This allows us to obtain precise completeness results for trace analysis.

References

- [1] M. Abadi, A.D. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148(1):1-70, 1999.
- [2] M. Abadi, A.D. Gordon. A Bisimulation Method for Cryptographic Protocols. *Nordic Journal of Computing*, 5(4):267-303, 1998.
- [3] R. Amadio, S. Prasad. The game of the name in cryptographic tables. RR 3733 INRIA Sophia Antipolis. In *Proc. of Asian'00*, LNCS, 2000.
- [4] R.M. Amadio, S. Lugiez. On the reachability problem in cryptographic protocols. In *Proc. of Concur'00*, LNCS, 2000. Full version: RR 3915 Inria Sophia Antipolis.
- [5] M. Boreale. Symbolic analysis of cryptographic protocols in the spi-calculus. Manuscript, 2000. Available at <http://www.dsi.unifi.it/boreale/papers.html>.
- [6] M. Boreale. STA: a tool for trace analysis of cryptographic protocols. ML object code and examples, 2001. Available at <http://www.dsi.unifi.it/boreale/tool.html>.
- [7] M. Boreale, R. De Nicola, R. Pugliese. Proof Techniques for Cryptographic Processes. In *Proc. of LICS'99*, IEEE Computer Society Press, 1999.
- [8] M. Burrows, M. Abadi, R.M. Needham. A logic of authentication. *Proc. of the Royal Society of London*, 426:233-271, 1989.
- [9] E.M. Clarke, S. Jha, W. Marrero. Using state exploration and a natural deduction style message derivation engine to verify security protocols. In *Proc. of the IFIP Working Conference on Programming Concepts and Methods (PROCOMET)*, 1998.
- [10] N. Durgin, P. Lincoln, J. Mitchell, A. Scedrov. Undecidability of bounded security protocols. In *Proc. of Workshop on Formal Methods and Security Protocols*, Trento, 1999.
- [11] D. Dolev, A.C. Yao. On the security of public key protocols. In *IEEE Transactions on Information Theory* 29(2):198-208, 1983.
- [12] A. Huima. Efficient infinite-state analysis of security protocols. In *Proc. of Workshop on Formal Methods and Security Protocols*, Trento, 1999.
- [13] G. Lowe. Breaking and Fixing the Needham-Schroeder Public-Key Protocol Using FDR. In *Proceedings of TACAS'96*, (T. Margaria, B. Steffen, Eds.), LNCS 1055, pp. 147-166, Springer-Verlag, 1996.
- [14] G. Lowe. A Hierarchy of Authentication Specifications. In *10th IEEE Computer Security Foundations Workshop*, IEEE Computer Society Press, 1997.
- [15] G. Lowe. Towards a completeness result for model checking of security protocols. In *11th Computer Security Foundations Workshop*, IEEE Computer Society Press, 1998.
- [16] D. Marchignoli, F. Martinelli. Automatic verification of cryptographic protocols through compositional analysis techniques. In *Proc. of TACAS99*, LNCS 1579:148-163, 1999.
- [17] C. Meadows. The NRL protocol analyzer: An overview. In *Proc. of the 2nd Conference on the Practical Applications of Prolog*, 1994.
- [18] R. Milner, J. Parrow, D. Walker. A calculus of mobile processes, (Part I and II). *Information and Computation*, 100:1-77, 1992.

- [19] L.C. Paulson. Proving Security Protocols Correct. In *Proc. of LICS'99*, IEEE Computer Society Press, 1999.
- [20] A.W. Roscoe. Modelling and verifying key-exchange using CSP and FDR. In *8th Computer Security Foundations Workshop*, IEEE Computer Society Press, 1995.
- [21] A.W. Roscoe. Proving security protocols with model checkers by data independent techniques. In *11th Computer Security Foundations Workshop*, IEEE Computer Society Press, 1998.
- [22] S. Schneider. Verifying Authentication Protocols in CSP. *IEEE Transactions on Software Engineering*, 24(8):743-758, 1998.
- [23] S. Stoller. A reduction for automated verification of security protocols. In *Proc. of Workshop on Formal Methods and Security Protocols*, Trento, 1999.

A Proof of Theorem 5.7

Lemma A.1 *Let σ be a marked symbolic trace and suppose that ρ satisfies σ . Then either σ is in sf, or there are σ' and ρ' such that $\sigma \succ \sigma'$, $\sigma\rho = \sigma'\rho'$ and ρ' satisfies σ' .*

PROOF: The proof is based on the following three facts, whose proof is relatively straightforward and omitted. Let σ be in solved form and suppose ρ satisfies σ . Then:

- (a) for each name k , $\sigma\rho \vdash k$ iff $\sigma \vdash k$;
- (b) $\mathbf{I}(\sigma\rho) = \mathbf{I}(\sigma)\rho$;
- (c) if $\sigma\rho \vdash M\rho$ and $N \in [M]_\sigma$ then $\sigma\rho \vdash N\rho$.

Now, suppose that σ is *not* in sf, and let σ_1 be the longest prefix of σ which is in sf. This means $\sigma = \sigma_1 \cdot \mathbf{a}\langle M \rangle \cdot \sigma_2$, for some M s.t. $\sigma_1 \not\vdash M$. By virtue of Proposition 5.5, we can pick some $N \in ([M]_{\sigma_1} - \mathbf{I}(\sigma_1))$.

Since ρ satisfies σ , we must have $\sigma_1\rho \vdash M\rho$ (as $\sigma\rho$ is a trace), which, by virtue of fact (c) above, implies $\sigma_1\rho \vdash N\rho$. Now, depending on the form of N , there are three cases:

- $N = \{N'\}_k$, for some N' and k s.t. $\sigma_1 \not\vdash k$. By virtue of (a) above, $\sigma_1\rho \not\vdash k$ as well. Since $\sigma_1\rho \vdash N\rho$, we deduce that $N\rho \in \mathbf{I}(\sigma_1\rho) = \mathbf{I}(\sigma_1)\rho$ (fact (b)). Thus there is $\{N''\}_u \in \mathbf{I}(\sigma_1)$ s.t. $(\{N''\}_u)\rho = (\{N'\}_k)\rho$. Therefore, there exists $\theta = \text{mgu}(\{N''\}_u, \{N'\}_k)$ and ρ is an instance of θ as a substitution. By rule (REF₁), $\sigma \succ \sigma' \stackrel{\text{def}}{=} \sigma\theta[\tilde{y}/\tilde{y}]$, for an appropriate renaming $[\tilde{y}/\tilde{y}]$. Furthermore, $\rho = \theta[\tilde{y}/\tilde{y}]\rho'$, for some ground substitution ρ' . Also note that, thanks to the renaming $[\tilde{y}/\tilde{y}]$, for each $\hat{x} \in \text{v}(\sigma')$, $\sigma' \setminus \hat{x}$ is not longer than $\sigma \setminus \hat{x}$, and this guarantess that ρ' satisfies σ' . Thus we have found σ' and ρ' as required by the statement of the lemma.
- $N = \{N'\}_x$, for some x and N' . If $\sigma_1\rho \vdash \rho(x)$, then we can apply (REF₂), and define $\sigma' \stackrel{\text{def}}{=} \sigma[\hat{x}/x]$ and $\rho' \stackrel{\text{def}}{=} [x/\hat{x}]\rho$, and the thesis will follow. If $\sigma_1\rho \not\vdash \rho(x) = k$, then we can proceed like in the previous case.
- $N = x$. Then $\sigma_1\rho \vdash \rho(x)$, and we can apply (REF₂).

□

Theorem A.2 (Theorem 5.7) *Let σ be a symbolic trace and s a trace. Then s is a solution of σ if and only if s is a solution of some $\sigma' \in \mathbf{SF}(\sigma)$.*

PROOF: Suppose $s = \sigma'\rho$ is a solution of σ' : then obviously s is a solution of σ , as $\sigma' = \sigma\theta$, for some θ . On the converse, suppose $s = \sigma\rho$ is a solution of σ . By repeated application of the previous lemma, we find that there is σ' in solved form and ρ' s.t. that $\sigma \succ \sigma'$, $s = \sigma\rho = \sigma'\rho'$ and ρ' satisfies σ' . Hence, by definition, s is a solution of σ' . □