

Simulation based security in the applied pi calculus*

Stéphanie Delaune¹, Steve Kremer¹, and Olivier Pereira²

¹LSV, ENS Cachan & CNRS & INRIA, France

²Université catholique de Louvain, B-1348 Belgium

ABSTRACT. We present a symbolic framework for refinement and composition of security protocols. The framework uses the notion of ideal functionalities. These are abstract systems which are secure by construction and which can be combined into larger systems. They can be separately refined in order to obtain concrete protocols implementing them. Our work builds on ideas from the “trusted party paradigm” used in computational cryptography models. The underlying language we use is the applied pi calculus which is a general language for specifying security protocols. In our framework we can express the different standard flavours of simulation-based security which happen to all coincide. We illustrate our framework on an authentication functionality which can be realized using the Needham-Schroeder-Lowe protocol. For this we need to define an ideal functionality for asymmetric encryption and its realization. We show a joint state result for this functionality which allows composition (even though the same key material is reused) using a tagging mechanism.

1 Introduction

Symbolic techniques showed to be a very useful approach for the modeling and analysis of security protocols: for twenty years, they improved our understanding of security protocols, allowed discovering flaws [17], and provided support for protocol design [9]. These techniques also resulted in the creation of powerful automated analysis tools (e.g. [3]), and impacted on several protocol standards used every day, e.g., [8].

Until now, symbolic techniques mostly concentrated on specifying and proving confidentiality and correspondence properties, i.e., showing which symbols are kept secret, and on which session parameters participants agree when a protocol session completes. However, such specifications do not provide any information about the behavior of protocols when they are used in composition with other protocols, and surprising behaviors are well known to happen in such contexts [7]. Moreover, protocols are often expected to provide more sophisticated security guarantees, which may be difficult to formalize.

In this paper, we present a symbolic framework for refinement and composition of security protocols, in which protocols are defined in terms of the behavior of trusted parties, or ideal functionalities, following the general outline of simulation-based security [5, 12, 4]. A lower-level protocol is said to securely emulate a higher-level protocol, or ideal functionality, if any behavior that can be observed from the interaction of an adversary with the lower-level protocol can also be observed from the interaction of another adversary (called

*This work has been supported by the ANR-07-SESU-002 project. O. Pereira is a Research Associate of the Belgian Fund for Scientific Research FRS-FNRS. Part of this work was done when he was visiting ENS Cachan.

the simulator) with the higher-level protocol. As a result, ideal functionalities can be successively refined into more concrete protocols, but also composed to build more complex protocols. Functionalities have been proposed for a wide range of protocol tasks, including general secure multi-party computation [5]. In the spi-calculus [2], Abadi and Gordon also present the idea of a protocol being equivalent to an idealized version. This is however more restrictive as they do not have the notion of a simulator.

Simulation-based security frameworks can typically be decomposed into two “layers”: (i) a foundational layer that provides a general model for concurrent computation, and (ii) a security layer that provides general security definitions, most importantly the notion of secure protocol emulation to be used. While the security layer is essentially common to all frameworks [4, 5, 6, 14, 18], including this paper, the foundational layer varies widely. Those variations typically lie in the concurrency model (from the most common token-passing mechanism to the use of schedulers with various powers) and in the definition of computational bounds. These differences typically result in incomparable security notions.

Defining simulation-based security while choosing the applied pi calculus [1] as the foundational layer brings the main benefits of this approach into the symbolic world:

- it provides a powerful machinery that can be used to specify a wide range of sophisticated protocol tasks in terms of the behavior of functionalities, and
- general composition theorems guarantee that protocols keep behaving as expected when executed in arbitrary contexts.

While we tried to stick to the common definitions from the security layer of simulation-based security frameworks, the use of the applied pi calculus as foundational layer raised interesting challenges.

First, at the most fundamental level, one has to adopt a notion of indistinguishability of processes. While the symmetric notions of computational indistinguishability and observational equivalence are commonly used in the cryptographic and symbolic worlds, the symmetry of such relations appeared to be too restrictive for our purpose. For instance, a symmetric equivalence relation makes the addition of an adversary that simply acts as a relay visible. The resulting undesired behaviors motivate the introduction of new notions of observational preorder and labelled simulation relations in the applied pi calculus.

Next, our attempts at translating ideal functionalities from the computational world into the symbolic world showed to be a non immediate task. For instance, traditional ideal functionalities for asymmetric encryption produce ciphertexts by encrypting random strings. An association table (plaintext/ciphertext) is then necessary to perform decryption. In our symbolic setting we avoid such a table by using two layers of encryption and a secure key.

Eventually, we investigate the statement of general composition theorems, and of a specific joint state composition theorem for our asymmetric encryption functionality, as this functionality is typically expected to be used in several protocol sessions. While these theorems appear to be the natural counterpart of their computational versions [4, 5, 6, 16], the joint state composition theorem brings message tagging constraints that, interestingly, are consistent with those obtained by using a completely different symbolic approach (e.g. [13]).

Because of lack of space the proofs are omitted, but can be found in [10].

2 The applied pi calculus

2.1 Syntax and informal semantics

To describe processes, one starts with a set of *names* (which are used to name communication channels or other atomic data), a set of *variables*, and a *signature* Σ which consists of the *function symbols* which will be used to define *terms*. In the case of security protocols, typical function symbols will include *enc* for encryption, and *dec* for decryption. Terms are defined as names, variables, and function symbols applied to other terms. Terms and function symbols are sorted. While the details of the sort systems are not essential it is important to distinguish sorts of base types and sorts of channel type. Function symbols can only be applied and return terms of base type. By the means of an equational theory E we describe the equations which hold on terms. We denote $=_E$ the equivalence relation induced by E .

Example 1 *In the equational theory $\{\text{dec}(\text{enc}(x, k), k) = x, \text{test}(\text{enc}(x, y), y) = \text{ok}\}$, we have that $\text{test}(\text{dec}(\text{enc}(\text{enc}(n, k_1), k_2), k_2), k_1) =_E \text{ok}$.*

In the applied pi calculus, one has *plain processes* and *extended processes*. Plain processes (P, Q, R) are built up in a similar way to processes in the pi calculus, except that messages can contain terms (rather than just names). Extended processes add *active substitutions* and restriction on variables. Below, M is a term, n is a name, and x a variable.

$$A, B, C := P \mid A \mid B \mid \nu n.A \mid \nu x.A \mid \{M/x\}$$

Active substitutions generalise “let”. The process $\nu x.(\{M/x\} \mid P)$ corresponds exactly to the process “let $x = M$ in P ”. As usual, names and variables have scopes, which are delimited by restrictions and by inputs. We write $fv(A)$, $bv(A)$, $fn(A)$ and $bn(A)$ for the sets of free and bound variables and free and bound names of A , respectively. We also assume that, in an extended process, there is at most one substitution for each variable, and there is exactly one when the variable is restricted. We say that an extended process is *closed* if all its variables are either bound or defined by an active substitution.

Active substitutions are useful because they allow us to map an extended process A to its *frame* $\phi(A)$ by replacing every plain process in A with 0 . A frame is an extended process built up from 0 and active substitutions by parallel composition and restriction. The frame $\phi(A)$ can be viewed as an approximation of A that accounts for the static knowledge A exposes to its environment, but not A 's dynamic behaviour. The *domain* of a frame ϕ , denoted by $\text{dom}(\phi)$, is the set of variables for which ϕ defines a substitution (those variables that are not under a restriction). An *evaluation context* $C[_]$ is an extended process with a hole instead of an extended process. A context $C[_]$ *closes* A when $C[A]$ is closed.

2.2 Semantics

The operational semantics of processes in the applied pi calculus is defined by structural rules defining two relations: *structural equivalence*, denoted \equiv , and *internal reduction*, denoted \rightarrow . Structural equivalence takes into account some basic structural rules, e.g. associativity and commutativity of the parallel operator. Internal reduction \rightarrow is the smallest relation on extended processes closed under structural equivalence and application of evaluation contexts such that:

COMM	$\text{out}(a, M).P \mid \text{in}(a, x).Q \rightarrow P \mid Q\{M/x\}$	
THEN	$\text{if } M = N \text{ then } P \text{ else } Q \rightarrow P$	where $M =_{\text{E}} N$
ELSE	$\text{if } M = N \text{ then } P \text{ else } Q \rightarrow Q$	

for any terms M and N without variable such that $M \neq_{\text{E}} N$

The operational semantics is extended by a *labelled* operational semantics enabling us to reason about processes that interact with their environment. Labelled operational semantics defines the relation $\xrightarrow{\alpha}$ where α is either an input $\text{in}(a, M)$ (a is a channel name and M is a term that can contain names and variables), or $\nu x.\text{out}(a, x)$ (x is a variable of base type), or $\text{out}(a, c)$ or $\nu c.\text{out}(a, c)$ (c is a channel name).

IN	$\text{in}(a, x).P \xrightarrow{\text{in}(a, M)} P\{M/x\}$	SCOPE	$\frac{A \xrightarrow{\alpha} A' \quad u \text{ does not occur in } \alpha}{\nu u.A \xrightarrow{\alpha} \nu u.A'}$
OUT-CH	$\text{out}(a, c).P \xrightarrow{\text{out}(a, c)} P$		
OPEN-CH	$\frac{A \xrightarrow{\text{out}(a, c)} A' \quad c \neq a}{\nu c.A \xrightarrow{\nu c.\text{out}(a, c)} A'}$	PAR	$\frac{\begin{array}{l} \text{bn}(\alpha) \cap \text{fn}(B) = \emptyset \\ \text{bv}(\alpha) \cap \text{fv}(B) = \emptyset \end{array}}{A \xrightarrow{\alpha} A' \quad B \xrightarrow{\alpha} B' \quad A' \equiv B'}{A \mid B \xrightarrow{\alpha} A' \mid B}$
OUT-T	$\text{out}(a, M).P \xrightarrow{\nu x.\text{out}(a, x)} P \mid \{M/x\}$ $x \notin \text{fv}(P) \cup \text{fv}(M)$	STRUCT	$\frac{A \equiv B \quad B \xrightarrow{\alpha} B' \quad A' \equiv B'}{A \xrightarrow{\alpha} A'}$

Our rules differ slightly from those described in [1]. It is proved in [11] that labelled bisimulation (see Section 2.3) in our system coincides with labelled bisimulation in [1].

Example 2 Consider the following process P :

$\nu k. (\text{in}(io_1, x).\text{out}(\text{net}, \text{enc}(x, k)) \mid \text{in}(\text{net}, y). \text{if } \text{test}(y, k) = \text{ok} \text{ then } \text{out}(io_2, \text{dec}(y, k)) \text{ else } 0)$.

The first component receives a message x on the channel io_1 and sends its encryption with the key k on the channel net . The second one is waiting for an input y on net , uses the secret key k to decrypt it. If the decryption succeeds, then it forwards the resulting plaintext on io_2 . We have that:

$$P \xrightarrow{\text{in}(io_1, s)} \nu k. (\text{out}(\text{net}, \text{enc}(s, k)) \mid \text{in}(\text{net}, y). \text{if } \text{test}(y, k) = \text{ok} \text{ then } \text{out}(io_2, \text{dec}(y, k)) \text{ else } 0) \\ \longrightarrow^* \nu k. \text{out}(io_2, s) \xrightarrow{\nu x.\text{out}(io_2, x)} \nu k. \{s/x\}$$

Let A be the resulting process. We have that $\phi(A) \equiv \nu k. \{s/x\}$.

2.3 Indistinguishability relations

In [1], it is shown that observational equivalence coincides with labelled bisimilarity. This relation is like the usual definition of bisimilarity, except that at each step one additionally requires that the processes are statically equivalent.

DEFINITION 1. Two terms M and N are equal in the frame ϕ , written $(M =_{\text{E}} N)\phi$, if, and only if there exists \tilde{n} and a substitution σ such that $\phi \equiv \nu \tilde{n}.\sigma$, $M\sigma =_{\text{E}} N\sigma$, and $\tilde{n} \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$. Two frames ϕ_1 and ϕ_2 are statically equivalent, $\phi_1 \approx_s \phi_2$, when $\text{dom}(\phi_1) = \text{dom}(\phi_2)$, and for all terms M, N we have that $(M =_{\text{E}} N)\phi_1$ if and only if $(M =_{\text{E}} N)\phi_2$.

Example 3 Let $\phi_0 = \nu s. \{\text{enc}(s, k)/x\}$ and $\phi_1 = \nu r. \{r/x\}$ where k, s and r are names and E be the theory given in Example 1. We have that $(\text{test}(x, k) =_{\text{E}} \text{ok})\phi_0$ but not $(\text{test}(x, k) =_{\text{E}} \text{ok})\phi_1$, thus $\phi_0 \not\approx_s \phi_1$. However, we have that $\nu k. \phi_0 \approx_s \phi_1$.

Now, we introduce the notion of a *barb*. Given an extended process A and a channel name a , we write $A \Downarrow a$ when $A \rightarrow^* C[\text{out}(a, M).P]$ for some term M , plain process P , and

evaluation context $C[_]$ that does not bind a .

DEFINITION 2. Observational preorder (\preceq) (resp. equivalence (\approx)) is the largest (resp. largest symmetric) relation on extended processes with same domain s.t. $A \mathcal{R} B$ implies

1. if $A \Downarrow a$ then $B \Downarrow a$;
2. if $A \rightarrow^* A'$, then $B \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' ;
3. $C[A] \mathcal{R} C[B]$ for all closing evaluation contexts $C[_]$.

DEFINITION 3. A relation \mathcal{R} on closed extended processes is a simulation if $A \mathcal{R} B$ implies

1. $\phi(A) \approx_s \phi(B)$,
2. if $A \rightarrow A'$, then $B \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' ,
3. if $A \xrightarrow{\alpha} A'$ and $fv(\alpha) \subseteq \text{dom}(A)$ and $bn(\alpha) \cap fn(B) = \emptyset$, then $B \rightarrow^* \xrightarrow{\alpha} \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' .

If \mathcal{R} and \mathcal{R}^{-1} are both simulations we say that \mathcal{R} is a bisimulation. Labelled similarity (\preceq_ℓ), resp. labelled bisimilarity (\approx_ℓ), is the largest simulation, resp. bisimulation relation.

Observational preorder and similarity were not introduced in [1]. However, these definitions seem natural for our purposes. Obviously we have that $\approx \subseteq \preceq$ and $\approx_\ell \subseteq \preceq_\ell$. We now show that labelled similarity is a precongruence.

PROPOSITION 4. Let A and B be two extended processes such that $A \preceq_\ell B$. We have that $C[A] \preceq_\ell C[B]$ for all closing evaluation context $C[_]$.

From this proposition it follows that $\preceq_\ell \subseteq \preceq$. Hence, we can use labelled similarity as a convenient proof technique for observational preorder. We actually expect the two relations to coincide but did not prove it as we did not need it. We have also the following lemma:

LEMMA 5. Let P and Q be two closed plain processes. We have that: (i) if $P \preceq_\ell Q$ then $!P \preceq_\ell !Q$; (ii) $!(P \mid Q) \preceq_\ell !P \mid !Q$ and $!P \mid !Q \preceq_\ell !(P \mid Q)$.

3 Simulation based security

3.1 Basic definitions

The simulation-based security approach classically distinguishes between *input-output channels*, which yield the internal interface of a protocol or *functionality* to its environment and *network channels*, which allow the adversary to interact with the functionality. We suppose that all channels are of one of these two sorts: IO or NET. Moreover the sort system ensures that names of sort NET can never be conveyed as data on a channel, i.e. these channels can never be transmitted. We write $\text{fn}(P)$ for $fn(P) \cap \text{NET}$.

DEFINITION 6. A functionality \mathcal{F} is a closed plain process. An adversary for \mathcal{F} is an evaluation context $\mathcal{A}[_]$ of the form: $v\widetilde{\text{net}}_1.(A_1 \mid v\widetilde{\text{net}}_2.(A_2 \mid \dots \mid v\widetilde{\text{net}}_k.(A_k \mid _)) \dots)$ where each A_i ($1 \leq i \leq k$) is a closed plain process, $\text{fn}(\mathcal{F}) \subseteq \bigcup_{1 \leq i \leq k} \widetilde{\text{net}}_i \subseteq \text{NET}$, and $fn(\mathcal{A}[_]) \cap \text{IO} = \emptyset$.

One may note that the nested form of the adversary process allows to express any arbitrary context while expliciting the restricted names whose scope ranges on the hole. We also note that if $\mathcal{A}[_]$ is an adversary for \mathcal{F} then $\text{fn}(\mathcal{A}[_]) = \text{fn}(\mathcal{A}[\mathcal{F}])$.

LEMMA 7. *Let \mathcal{F} be a functionality and $\mathcal{A}_1[-]$ be an adversary for \mathcal{F} . Then $\mathcal{A}_1[\mathcal{F}]$ is a functionality. If $\mathcal{A}_2[-]$ is an adversary for $\mathcal{A}_1[\mathcal{F}]$, then $\mathcal{A}_2[\mathcal{A}_1[-]]$ is an adversary for \mathcal{F} .*

While adversaries control the communication of functionalities on NET channels, IO contexts model the environment of functionalities, providing inputs and collecting outputs.

DEFINITION 8. *An IO context is an evaluation context $C_{io}[-]$ of the form $v\tilde{i}o_1.(C_1 \mid v\tilde{i}o_2.(C_2 \mid \dots \mid v\tilde{i}o_k.(C_k \mid -) \dots))$ where each C_i ($1 \leq i \leq k$) is a closed plain process, and $\bigcup_{1 \leq i \leq k} \tilde{i}o_i \subseteq \text{IO}$*

Note that if \mathcal{F} is a functionality and $C_{io}[-]$ is an IO context, then $C_{io}[\mathcal{F}]$ is a functionality.

3.2 Strong simulatability

The notion of strong simulatability [15], which is probably the simplest secure emulation notion used in simulation-based security, can be formulated in our setting.

DEFINITION 9. *Let \mathcal{F}_1 and \mathcal{F}_2 be two functionalities. \mathcal{F}_1 emulates \mathcal{F}_2 in the sense of strong simulatability, written $\mathcal{F}_1 \leq^{SS} \mathcal{F}_2$, if there exists an adversary \mathcal{S} for \mathcal{F}_2 (the simulator) such that $\text{fnet}(\mathcal{F}_1) = \text{fnet}(\mathcal{S}[\mathcal{F}_2])$ and $\mathcal{F}_1 \preceq \mathcal{S}[\mathcal{F}_2]$.*

The definition ensures that any behavior of \mathcal{F}_1 can be matched by \mathcal{F}_2 executed in the presence of a specific adversary \mathcal{S} . Hence, there are no more attacks on \mathcal{F}_1 than attacks on \mathcal{F}_2 . Moreover, the presence of \mathcal{S} allows abstract definitions of higher-level functionalities, which are independent of a specific realization. One may also note that $0 \leq^{SS} \mathcal{F}$ for any functionality \mathcal{F} . This seems natural in a simulation based framework which only aims at preserving security. Non-triviality conditions may be imposed independently [5].

Example 4 *Let $\mathcal{F}_{cc} = \text{in}(io_1, x_s).\text{out}(\text{net}_{cc}, \text{ok}).\text{in}(\text{net}_{cc}, x)$. if $x = \text{ok}$ then $\text{out}(io_2, x_s)$. The functionality models a confidential channel and takes a potentially secret value s as input on channel io_1 . The adversary is notified via channel net_{cc} that this value is to be transmitted. If the adversary agrees the value is output on channel io_2 . This functionality can be realized by the process described in Example 2. Let $\mathcal{S} = v\text{net}_{cc}.\text{in}(\text{net}_{cc}, x).\text{vr}.\text{out}(\text{net}, r).\text{in}(\text{net}, x)$. if $x = r$ then $\text{out}(\text{net}_{cc}, \text{ok}) \mid -$. We indeed have that $P \preceq_{\ell} \mathcal{S}[\mathcal{F}_{cc}]$ and $\text{fnet}(P) = \text{fnet}(\mathcal{S}[\mathcal{F}_{cc}])$.*

In order to examine the properties of strong simulatability in our specific setting, we now define a particular adversary $D[-]$ which is called a *dummy adversary*: intuitively, it acts as a relay which forwards all messages. The formal definition is technical because $D[-]$ needs to both restrict all names in $\text{fnet}(\mathcal{F})$ and ensure that $\text{fnet}(\mathcal{F}) = \text{fnet}(D[\mathcal{F}])$. It therefore relies on two internal channels $\text{sim}_j^{i/o}$ for inputs, resp. outputs, for each channel in $\text{fnet}(\mathcal{F})$.

DEFINITION 10. *Let \mathcal{F} be a functionality. The dummy adversary for \mathcal{F} is the adversary $D[-] = v\tilde{\text{sim}}.(D_1 \mid v\tilde{\text{net}}.(D_2 \mid -))$ where $\tilde{\text{net}} = \text{fnet}(\mathcal{F}) = \{\text{net}_1, \dots, \text{net}_n\}$; $\tilde{\text{sim}} = \{\text{sim}_1^i, \dots, \text{sim}_n^i, \text{sim}_1^o, \dots, \text{sim}_n^o\} \subseteq \text{NET}$; and*

- $D_1 = !\text{in}(\text{net}_1, x).\text{out}(\text{sim}_1^i, x) \mid \dots \mid !\text{in}(\text{net}_n, x).\text{out}(\text{sim}_n^i, x) \mid !\text{in}(\text{sim}_1^o, x).\text{out}(\text{net}_1, x) \mid \dots \mid !\text{in}(\text{sim}_n^o, x).\text{out}(\text{net}_n, x);$
- $D_2 = !\text{in}(\text{sim}_1^i, x).\text{out}(\text{net}_1, x) \mid \dots \mid !\text{in}(\text{sim}_n^i, x).\text{out}(\text{net}_n, x) \mid !\text{in}(\text{net}_1, x).\text{out}(\text{sim}_1^o, x) \mid \dots \mid !\text{in}(\text{net}_n, x).\text{out}(\text{sim}_n^o, x).$

By construction we have that $\text{fnet}(D[\mathcal{F}]) = \text{fnet}(\mathcal{F})$.

LEMMA 11. *Let \mathcal{F} be a functionality and $D[-]$ be the dummy adversary for \mathcal{F} : $\mathcal{F} \preceq D[\mathcal{F}]$.*

However, we do not have that $\mathcal{F} \approx D[\mathcal{F}]$, since $D[\mathcal{F}]$ has more non-determinism than \mathcal{F} . A direct consequence of this lemma is that $\mathcal{F}_1 \preceq \mathcal{F}_2$ and $\text{fnet}(\mathcal{F}_1) = \text{fnet}(\mathcal{F}_2)$ implies that $\mathcal{F}_1 \leq^{\text{SS}} \mathcal{F}_2$: as $\mathcal{F}_2 \preceq D[\mathcal{F}_2]$ we have by transitivity that $\mathcal{F}_1 \preceq D[\mathcal{F}_2]$. We use these observations to show that \leq^{SS} is a preorder (Lemma 12), which is closed under application of IO contexts (Proposition 13) and parallel composition (Proposition 14).

LEMMA 12. (i) $\mathcal{F}_1 \leq^{\text{SS}} \mathcal{F}_1$; (ii) $\mathcal{F}_1 \leq^{\text{SS}} \mathcal{F}_2$ and $\mathcal{F}_2 \leq^{\text{SS}} \mathcal{F}_3 \Rightarrow \mathcal{F}_1 \leq^{\text{SS}} \mathcal{F}_3$.

PROPOSITION 13. Let $\mathcal{F}_1, \mathcal{F}_2$ be two functionalities and C_{io} be an IO context.

$$\mathcal{F}_1 \leq^{\text{SS}} \mathcal{F}_2 \Longrightarrow C_{io}[\mathcal{F}_1] \leq^{\text{SS}} C_{io}[\mathcal{F}_2].$$

PROPOSITION 14. Let $\mathcal{F}_1, \mathcal{F}_2$ and \mathcal{F}_3 be three functionalities. We have that:

(i) $\mathcal{F}_1 \leq^{\text{SS}} \mathcal{F}_2 \Rightarrow \mathcal{F}_1 \mid \mathcal{F}_3 \leq^{\text{SS}} \mathcal{F}_2 \mid \mathcal{F}_3$; and (ii) $\mathcal{F}_1 \leq^{\text{SS}} \mathcal{F}_2 \Rightarrow !\mathcal{F}_1 \leq^{\text{SS}} !\mathcal{F}_2$.

While, (i) is a direct consequence of Proposition 13 (notice that $_ \mid \mathcal{F}_3$ is an IO-context) the proof of (ii) is more involved and given in [10].

3.3 Other notions of simulation based security

Several other notions of simulation based security appear in the literature. We present them, and show that they all coincide in our setting. This coincidence is regarded as highly desirable [15, 14], while it does not hold in most simulation-based security frameworks [5, 4].

DEFINITION 15. Let \mathcal{F}_1 and \mathcal{F}_2 be two functionalities and \mathcal{A} be any adversary for \mathcal{F}_1 .

- \mathcal{F}_1 emulates \mathcal{F}_2 in the sense of black box simulatability, $\mathcal{F}_1 \leq^{\text{BB}} \mathcal{F}_2$, iff $\exists \mathcal{S}. \forall \mathcal{A}. \mathcal{A}[\mathcal{F}_1] \preceq \mathcal{A}[\mathcal{S}[\mathcal{F}_2]]$ where \mathcal{S} is an adversary for \mathcal{F}_2 with $\text{fnet}(\mathcal{S}[\mathcal{F}_2]) = \text{fnet}(\mathcal{F}_1)$.
- \mathcal{F}_1 emulates \mathcal{F}_2 in the sense of universally composable simulatability, $\mathcal{F}_1 \leq^{\text{UC}} \mathcal{F}_2$, iff $\forall \mathcal{A}. \exists \mathcal{S}. \mathcal{A}[\mathcal{F}_1] \preceq \mathcal{S}[\mathcal{F}_2]$ where \mathcal{S} is an adversary for \mathcal{F}_2 s.t. $\text{fnet}(\mathcal{A}[\mathcal{F}_1]) = \text{fnet}(\mathcal{S}[\mathcal{F}_2])$.
- \mathcal{F}_1 emulates \mathcal{F}_2 in the sense of universally composable simulatability with dummy adversary, $\mathcal{F}_1 \leq^{\text{UCDA}} \mathcal{F}_2$, iff $\exists \mathcal{S}. D[\mathcal{F}_1] \preceq \mathcal{S}[\mathcal{F}_2]$ where D is the dummy adversary for \mathcal{F}_1 and \mathcal{S} is an adversary for \mathcal{F}_2 such that $\text{fnet}(\mathcal{S}[\mathcal{F}_2]) = \text{fnet}(D[\mathcal{F}_1])$.

THEOREM 16. We have that $\leq^{\text{SS}} = \leq^{\text{BB}} = \leq^{\text{UC}} = \leq^{\text{UCDA}}$.

The above security notions can also be defined replacing observational preorder by observational equivalence denoted $\leq^{\text{SS}}_{\approx}, \leq^{\text{BB}}_{\approx}, \leq^{\text{UC}}_{\approx}$ and $\leq^{\text{UCDA}}_{\approx}$. Surprisingly, the use of observational equivalence turns out to be too strong, ruling out natural secure emulation cases: for instance, the $\leq^{\text{SS}}_{\approx}$ relation is not reflexive, due to the extra non-determinism that the simulator introduces. While symbolic analysis techniques typically rely on bisimulation relations, this is however consistent with the definitions proposed in the task-PIOA framework [6], which also allows non-deterministic executions for simulation based security.

4 Applications

We illustrate our framework by showing the secure emulation of a mutual authentication functionality by the Needham-Shroeder-Lowe (NSL) protocol [17]. As the NSL protocol uses public key encryption we first introduce in Section 4.1 functionalities for asymmetric

$$\begin{aligned}
\mathcal{P}_{\text{pke}} &:= \text{in}(io_{\text{pke}}, io_{\text{pke}}^1).vsk.\text{out}(io_{\text{pke}}^1, \langle \text{KEY}, \text{pk}(sk) \rangle). \\
&\quad (\text{let } io_{\text{pke}}^i = io_{\text{pke}}^1 \text{ in } !\mathcal{P}_{\text{enc}} \mid \text{let } io_{\text{pke}}^i = io_{\text{pke}}^2 \text{ in } !\mathcal{P}_{\text{enc}} \mid !\mathcal{P}_{\text{dec}}) \\
\mathcal{P}_{\text{enc}} &:= \text{in}(io_{\text{pke}}^i, \langle = \text{ENC}, m \rangle). \\
&\quad vr_2. \text{let } m_{\text{enc}} = \text{aenc}(\langle \text{TAG}_0, m \rangle, \text{pk}(sk), r_2) \text{ in } \text{out}(io_{\text{pke}}^i, \langle \text{CIPHER}, m_{\text{enc}} \rangle) \\
\mathcal{P}_{\text{dec}} &:= \text{in}(io_{\text{pke}}^1, \langle = \text{DEC}, m \rangle). \\
&\quad \text{let } \langle = \text{TAG}_0, m_1 \rangle = \text{adec}(m, sk) \text{ in } \text{out}(io_{\text{pke}}^1, \langle \text{PLAIN}, m \rangle)
\end{aligned}$$

Figure 1: Real encryption functionality

encryption. Then, we briefly present the mutual authentication functionality and its realization through the NSL protocol. Finally we use the joint state composition result in Section 4.3 to obtain a result for an unbounded number of concurrent sessions.

4.1 Asymmetric encryption with joint state

We introduce a functionality for asymmetric encryption together with a *joint state composition result* which is crucial for composition of protocols that share key material. Even though encryption in a Dolev-Yao model is already idealized we will see that by introducing an *ideal functionality* for encryption we are able to obtain a joint state composition result. Throughout this section we rely on the following equational theory allowing us to model randomized asymmetric encryption:

$$\text{adec}(\text{aenc}(x, \text{pk}(y), z), y) = x \quad \text{testdec}(\text{aenc}(x, \text{pk}(y), z), y) = \text{ok}.$$

Real encryption. This functionality (described in Figure 1) receives a channel name io_{pke}^1 , which will be used for all sensitive information exchanges. A fresh private key sk is generated and the corresponding public key, i.e. $\text{pk}(sk)$, is sent on io_{pke}^1 . Then the process is ready to receive encryption or decryption requests. Note that encryption requests can be sent on the sensitive channel io_{pke}^1 or on the public channel io_{pke}^2 which is the channel the environment will typically use. Decryption requests are only available through the sensitive channel io_{pke}^1 and thus will not be used by the attacker. Each time a decryption request is received on the channel io_{pke}^1 , it tries to decrypt the ciphertext and checks whether the tag is TAG_0 . If so, it outputs the plaintext on the channel io_{pke}^1 . Otherwise, it does nothing.

Ideal functionality. We now propose, in Figure 2, an idealized version \mathcal{F}_{pke} of the real encryption functionality, which guarantees that the confidentiality of messages is preserved independently of any cryptanalytic effort that could be performed on ciphertexts from the knowledge of public keys. In various cryptographic settings [4, 5, 16], this is achieved by computing ciphertexts as the encryption of random messages instead of the actual plaintext. To be able to perform decryption, a table for plaintext/ciphertext associations is maintained. The burden of this association table is avoided in our symbolic specification by using two layers of encryption: messages are first encrypted using a secure key $\text{pk}(ssk)$, then tagged and encrypted with the public key $\text{pk}(sk)$ that is published during the initialization step. We stress that neither $\text{pk}(ssk)$ nor ssk are ever transmitted by \mathcal{F}_{pke} , guaranteeing that it is impossible to decrypt such a ciphertext outside the functionality, even if the key sk is adversarially chosen, which will be a crucial feature for our joint state composition theorem.

$$\begin{aligned}
\mathcal{F}_{\text{pke}} &:= \text{in}(io_{\text{pke}}, io_{\text{pke}}^1). \text{out}(net, \text{INIT}). \text{in}(net, \langle = \text{ALGO}, sk, tag \rangle). \text{out}(io_{\text{pke}}^1, \langle \text{KEY}, \text{pk}(sk) \rangle). \\
&\quad vssk. (\text{let } io_{\text{pke}}^i = io_{\text{pke}}^1 \text{ in } !\mathcal{F}_{\text{enc}} \mid \text{let } io_{\text{pke}}^i = io_{\text{pke}}^2 \text{ in } !\mathcal{F}_{\text{enc}} \mid !\mathcal{F}_{\text{dec}}) \\
\mathcal{F}_{\text{enc}} &:= \text{in}(io_{\text{pke}}^i, \langle = \text{ENC}, m \rangle). vr_1. vr_2. \\
&\quad \text{let } alea = \text{aenc}(m, \text{pk}(ssk), r_1) \text{ in let } menc = \text{aenc}(\langle tag, alea \rangle, \text{pk}(sk), r_2) \text{ in} \\
&\quad \text{out}(io_{\text{pke}}^i, \langle \text{CIPHER}, menc \rangle) \\
\mathcal{F}_{\text{dec}} &:= \text{in}(io_{\text{pke}}^1, \langle = \text{DEC}, m \rangle). \text{let } \langle = tag, m_1 \rangle = \text{adec}(m, sk) \text{ in} \\
&\quad \text{if } \text{testdec}(m_1, ssk) = \text{ok} \text{ then } \text{out}(io_{\text{pke}}^1, \langle \text{PLAIN}, \text{adec}(m_1, ssk) \rangle) \\
&\quad \text{else } \text{out}(io_{\text{pke}}^1, \langle \text{PLAIN}, m_1 \rangle)
\end{aligned}$$

Figure 2: Ideal encryption functionality

During the initialization, the attacker chooses the secret key sk and the tag that will be added in each encryption. Then a secure key ssk is generated and now the process is ready to receive encryption or decryption requests. Each time the process receives an encryption request, it computes the corresponding ciphertext and outputs the corresponding ciphertext. As explained above, the plaintext m is first encrypted using $\text{pk}(ssk)$ before being tagged and encrypted with $\text{pk}(sk)$. When the process receives a decryption request, it tries to decrypt the ciphertext and checks if the tag is the tag provided during the initialization. Then, it checks if the resulting plaintext is encrypted under $\text{pk}(ssk)$. If so, this means that this ciphertext has been produced by the encryption functionality and thus has to be decrypted twice. Otherwise, the ciphertext has been produced by the attacker.

Realization. The real encryption functionality realizes the ideal one, i.e., $\mathcal{P}_{\text{pke}} \leq^{\text{SS}} \mathcal{F}_{\text{pke}}$. This is witnessed by $\mathcal{A}_{\text{pke}} = vnet.(\text{in}(net, = \text{INIT}).vsk. \text{out}(net, (\text{ALGO}, sk, \text{TAG}_0)) \mid -)$.

Composition with joint state. While \leq^{SS} is stable under replication this is not always sufficient to obtain composition guarantees. Indeed replication of a process also replicates all key generation operations. In order to obtain self-composition and inter-protocol composition with common key material we need a *joint state functionality*, i.e. a functionality that realizes $!\mathcal{F}_{\text{pke}}$ while reusing the same key material. We actually consider the functionality $\underline{\mathcal{F}}_{\text{pke}}$, which is a variant of \mathcal{F}_{pke} in which each message is tagged. More precisely, the process $\underline{\mathcal{F}}_{\text{pke}}$ is defined as \mathcal{F}_{pke} , except that: (i) the functionality begins with the instructions $\text{in}(io_{\text{pke}}, io_{\text{pke}}^1). \text{in}(io_{\text{pke}}^1, sid)$ instead of $\text{in}(io_{\text{pke}}, io_{\text{pke}}^1)$, (ii) each input of the form $\text{in}(c, m)$ is replaced by $\text{in}(c, \langle = sid, m \rangle)$, and (iii) each output of the form $\text{out}(c, m)$ is replaced by $\text{out}(c, \langle sid, m \rangle)$.

The joint state functionality $\mathcal{P}_{\text{js}}[\underline{\mathcal{F}}_{\text{pke}}]$ (see Figure 3) uses a single instance of $\underline{\mathcal{F}}_{\text{pke}}$ for all protocol sessions. All the requests to the joint state functionality are received on the public channel io_{pke} in process $\mathcal{P}_{\text{js}}^1$. They are then forwarded using the private IO channel $cont$ to $\mathcal{P}_{\text{js}}^2$. The process $\mathcal{P}_{\text{js}}^2$ shares the private channel io_{pke} with $\underline{\mathcal{F}}_{\text{pke}}$ and forwards all the requests after concatenating the session identifier to the plaintext. Then the response is again forwarded to the process $\mathcal{P}_{\text{js}}^1$ which outputs the result on the public channel io_{pke} .

We now observe that the following joint state composition result holds. One instance of the encryption functionality can be used to emulate an unbounded number of such instances using the joint state process: $\mathcal{P}_{\text{js}}[\underline{\mathcal{F}}_{\text{pke}}] \leq^{\text{SS}} !\underline{\mathcal{F}}_{\text{pke}}$.

$$\begin{aligned}
\mathcal{P}_{js} &:= vcont.(\mathcal{P}_{js}^1 \mid vio_{pke}, io_{pke}^2.(\mathcal{P}_{js}^2 \mid -)) \\
\mathcal{P}_{js}^1 &:= in(io_{pke}, io_{pke}^1).in(io_1, sid).out(cont, \langle sid, INIT \rangle). \\
&\quad in(cont, (= KEY, pk)).out(io_{pke}^1, \langle sid, KEY, pk \rangle). \\
&\quad (\text{let } io_{pke}^i = io_{pke}^1 \text{ in } !\mathcal{P}_{js\text{-enc}}^1 \mid \text{let } io_{pke}^i = io_{pke}^2 \text{ in } !\mathcal{P}_{js\text{-enc}}^1 \mid !\mathcal{P}_{js\text{-dec}}^1 \mid \\
&\quad \quad !in(io_{pke}, io_{pke}^1).in(io_1, sid).out(io_{pke}^1, \langle sid, KEY, pk \rangle). \\
&\quad \quad (\text{let } io_{pke}^i = io_{pke}^1 \text{ in } !\mathcal{P}_{js\text{-enc}}^1 \mid \text{let } io_{pke}^i = io_{pke}^2 \text{ in } !\mathcal{P}_{js\text{-enc}}^1 \mid !\mathcal{P}_{js\text{-dec}}^1)) \\
\mathcal{P}_{js\text{-enc}}^1 &:= in(io_{pke}^i, \langle = sid, = ENC, m \rangle). out(cont, \langle sid, ENC, m \rangle). \\
&\quad in(cont, \langle = CIPHER, c \rangle). out(io_{pke}^i, \langle sid, CIPHER, c \rangle) \\
\mathcal{P}_{js\text{-dec}}^1 &:= in(io_{pke}^1, \langle = sid, = DEC, c \rangle). out(cont, \langle sid, dec, c \rangle). \\
&\quad in(cont, \langle = PLAIN, m \rangle). out(io_{pke}, \langle sid, PLAIN, m \rangle) \\
\mathcal{P}_{js}^2 &:= in(cont, \langle sid, = INIT \rangle).vio_{pke}'.out(io_{pke}, io_{pke}'). \\
&\quad in(io_{pke}', \langle = KEY, pk \rangle).out(cont, \langle KEY, pk \rangle).(!\mathcal{P}_{js\text{-enc}}^2 \mid !\mathcal{P}_{js\text{-dec}}^2) \\
\mathcal{P}_{js\text{-enc}}^2 &:= in(cont, \langle = sid, = ENC, m \rangle). out(io_{pke}', \langle ENC, \langle sid, m \rangle \rangle). \\
&\quad in(io_{pke}', \langle = CIPHER, c \rangle). out(cont, \langle CIPHER, c \rangle) \\
\mathcal{P}_{js\text{-dec}}^2 &:= in(cont, \langle = sid, = DEC, c \rangle). out(io_{pke}', \langle DEC, c \rangle). \\
&\quad in(io_{pke}', \langle = PLAIN, m \rangle). out(cont, \langle PLAIN, m \rangle)
\end{aligned}$$

Figure 3: Joint state IO-context

$$\begin{aligned}
\mathcal{A}_{js} &:= vcs.(\mathcal{A}_{js}^1 \mid vnet.(\mathcal{A}_{js}^2 \mid -)) \\
\mathcal{A}_{js}^1 &:= in(cs, INIT).out(net, INIT).in(net, \langle = ALGO, sk, tag \rangle).out(cs, \langle ALGO, sk, tag \rangle) \\
\mathcal{A}_{js}^2 &:= in(net, \langle sid, = INIT \rangle). out(cs, INIT). \\
&\quad in(cs, \langle = ALGO, sk, tag \rangle). out(net, \langle sid, ALGO, sk, \langle sid, tag \rangle \rangle). \\
&\quad !in(net, \langle sid', = INIT \rangle).out(net, \langle sid', ALGO, sk, \langle sid', tag \rangle \rangle)
\end{aligned}$$

Figure 4: Joint state adversary

This relation is witnessed by the adversary \mathcal{A}_{js} described in Figure 4 as we have that: $\mathcal{P}_{js}[\mathcal{F}_{pke}] \preceq \mathcal{A}_{js}[\mathcal{F}_{pke}]$. This adversary launches several functionalities with the same key sk . However, note that the session identifier sid used to tag each encryption associated could be different. The value of these session identifiers is selected by the attacker.

Note that it is crucial to introduce an ideal encryption functionality. We indeed have that $\mathcal{P}_{js}[\mathcal{P}_{pke}] \leq^{SS} \mathcal{P}_{js}[\mathcal{F}_{pke}] \leq^{SS} !\mathcal{F}_{pke}$ as well as $!\mathcal{P}_{pke} \leq^{SS} !\mathcal{F}_{pke}$ (where \mathcal{P}_{pke} is defined from \mathcal{P}_{pke} in the same way as \mathcal{F}_{pke} from \mathcal{F}_{pke}). However, $\mathcal{P}_{js}[\mathcal{P}_{pke}] \not\leq^{SS} !\mathcal{P}_{pke}$. In particular $!\mathcal{P}_{pke}$ will provide multiple public keys while $\mathcal{P}_{js}[\mathcal{P}_{pke}]$ only provides a single one. Taking the more abstract ideal functionality allows this to be avoided by a simulator that chooses the same secret key for each instance of the functionality.

4.2 Mutual authentication

Because of lack of space we only briefly sketch the mutual authentication functionality. The details of the functionalities and the simulator are given in [10].

Ideal functionality for mutual authentication. The functionality $\mathcal{F}_{\text{auth}}$ roughly works as follows. Both the initiator ($\mathcal{F}_{\text{init}}$) and the responder ($\mathcal{F}_{\text{resp}}$) receive a request for mutual authentication on their *io* channel. They forward this request to the adversary and, if both parties are honest, to a trusted host \mathcal{F}_{th} which compares these requests and authorizes going further if they match. Eventually, when the adversary asks to finish the protocol, then both participants complete the protocol session.

Realization of mutual authentication. The functionality $\mathcal{F}_{\text{auth}}$ can be realized by a functionality \mathcal{P}_{nsi} implementing the well-known Needham-Schroeder-Lowe protocol [17]. We have that $\mathcal{P}_{\text{nsi}} \leq^{\text{SS}} \mathcal{F}_{\text{auth}}$ by showing that $\mathcal{P}_{\text{nsi}} \preceq \mathcal{S}[\mathcal{F}_{\text{auth}}]$ for some \mathcal{S} .

4.3 From one to many sessions

We have that $\mathcal{P}_{\text{nsi}} \leq^{\text{SS}} \mathcal{F}_{\text{auth}}$. This result only shows that \mathcal{P}_{nsi} is as secure as $\mathcal{F}_{\text{auth}}$ for a single session of the protocol. By Proposition 14 we have that $!\mathcal{P}_{\text{nsi}} \leq^{\text{SS}} !\mathcal{F}_{\text{auth}}$ but this does not correspond to the expected security for an unbounded number of sessions, as each session uses a different key. To show that $!\mathcal{F}_{\text{auth}}$ can be realized with shared key material we use our joint state result. To apply this result we need the following technical lemma which allows pushing the replication under the restricted channel c .

LEMMA 17. *Let n be a name and c be a channel name such that $c \notin \text{fn}(P) \cup \text{fn}(Q)$.*

$$vc.![vn.(\text{out}(c, n) \mid P) \mid \text{in}(c, x).Q] \preceq_\ell !vc.![vn.(\text{out}(c, n) \mid P) \mid \text{in}(c, x).Q].$$

This lemma allows us to apply the joint state result and obtain a result for an unbounded number of sessions sharing keys. Note that the joint state context uses a tagging mechanism.

5 Conclusions

This paper proposes a symbolic framework for the analysis of security protocols along the lines of the simulation based security approach, while adopting the applied pi calculus as its basic layer. We state central definitions and security notions, show general composition theorems and specific joint-state composition results for asymmetric encryption, and illustrate their use in the analysis of a mutual authentication protocol.

This framework brings the benefits of the secure composition theorems associated to simulation based security into the symbolic world, and opens the path to the analysis of more sophisticated protocols that can naturally be specified by the behavior of an ideal functionality. At a more fundamental level, we use preorder notions, which can be established by labeled simulation. While the use of labeled bisimulations is quite common in the applied pi calculus and has been integrated in automatic provers, the automation of proofs relying on labeled simulation appears as an interesting challenge for future works. Another direction for future work is to give a precise characterization of what properties are preserved by strong simulatability.

References

- [1] M. Abadi and C. Fournet. Mobile values, new names, and secure communication. In *Proc. 28th ACM Symp. on Principles of Programming Languages (POPL'01)*. ACM, 2001.

- [2] M. Abadi and A. D. Gordon. A calculus for cryptographic protocols: The spi calculus. Technical Report 149, SRC, 1998.
- [3] A. Armando et al. The AVISPA Tool for the automated validation of internet security protocols and applications. In *Proc. 17th Int. Conference on Computer Aided Verification (CAV'05)*, LNCS. Springer, 2005.
- [4] M. Backes, B. Pfitzmann, and M. Waidner. The reactive simulatability (RSIM) framework for asynchronous systems. *Information and Computation*, 205(12):1685–1720, 2007.
- [5] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proc. 42nd IEEE Symp. on Foundations of Computer Science (FOCS'01)*, 2001.
- [6] R. Canetti, L. Cheung, D. Kaynar, N. Lynch, and O. Pereira. Compositional security for Task-PIOAs. In *Proc. 20th Computer Security Foundations Symposium (CSF'07)*, 2007.
- [7] R. Canetti and J. Herzog. Universally composable symbolic analysis of mutual authentication and key exchange protocols. In *Proc. Theory of Cryptography Conference (TCC'06)*, LNCS. Springer, 2006.
- [8] I. Cervesato, A. Jaggard, A. Scedrov, J.-K. Tsay, and C. Walstad. Breaking and fixing public-key kerberos. *Information and Computation*, 206(2-4):402–424, 2008.
- [9] A. Datta, A. Derek, J. C. Mitchell, and D. Pavlovic. Abstraction and refinement in protocol derivation. In *Proc. 17th Comp. Security Foundations Workshop (CSFW'04)*, 2004.
- [10] S. Delaune, S. Kremer, and O. Pereira. Simulation based security in the applied pi calculus. Cryptology ePrint Archive, Report 2009/267. <http://eprint.iacr.org/>.
- [11] S. Delaune, S. Kremer, and M. D. Ryan. Symbolic bisimulation for the applied pi-calculus. In *Proc. 27th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'07)*, LNCS. Springer, 2007.
- [12] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game: A completeness theorem for protocols with honest majority. In *Proc. 19th ACM Symposium on the Theory of Computing (STOC'87)*. ACM Press, 1987.
- [13] J. D. Guttman and F. J. Thayer. Protocol independence through disjoint encryption. In *Proc. 13th IEEE Computer Security Foundations Workshop (CSFW'00)*, 2000.
- [14] R. Küsters. Simulation-Based Security with Inexhaustible Interactive Turing Machines. In *Proc. 19th IEEE Computer Security Foundations Workshop (CSFW'06)*, 2006.
- [15] R. Küsters, A. Datta, J. C. Mitchell, and A. Ramanathan. On the relationships between notions of simulation-based security. *Journal of Cryptology*, 21(4):492–546, 2008.
- [16] R. Küsters and M. Tuengerthal. Joint State Theorems for Public-Key Encryption and Digital Signature Functionalities with Local Computation. In *Proc. 21st IEEE Computer Security Foundations Symposium (CSF'08)*, 2008.
- [17] G. Lowe. An attack on the Needham-Schroeder public key authentication protocol. *Information Processing Letters*, 56(3):131–133, 1995.
- [18] P. Mateus, J. Mitchell, and A. Scedrov. Composition of cryptographic protocols in a probabilistic polynomial-time calculus. In *Proc. 14th Conference on Concurrency Theory (CONCUR'03)*, LNCS. Springer, 2003.