

MPRI Lecture Notes

Course 2-30

Cryptographic protocols Formal and Computational Proofs

Authors :

- BRUNO BLANCHET
- HUBERT COMON-LUNDH
- STÉPHANIE DELAUNE
- CÉDRIC FOURNET
- STEVE KREMER
- DAVID POINTCHEVAL

Contents

1	Introduction	5
I	Modelling Protocols and Security Properties	7
2	An Introductory Example	9
2.1	An Informal Description	9
2.2	A More Formal Analysis	10
2.3	An Attack on the Fixed Version of the Protocol	12
2.4	Further Readings	13
2.5	Exercises	14
3	A Small Process Calculus	15
3.1	Preliminaries	15
3.1.1	Messages	15
3.1.2	Assembling Terms into Frames	17
3.1.3	Deduction	17
3.1.4	Static Equivalence	18
3.2	Protocols	19
3.2.1	Protocol Language	19
3.2.2	Operational Semantics	21
3.3	Security Properties	24
3.3.1	Secrecy	24
3.3.2	Correspondence Properties	25
3.3.3	Guessing Attacks	26
3.3.4	Equivalence Properties	27
3.4	Further Readings	30
3.5	Exercises	31
II	Verification in the Symbolic Setting	33
4	Deducibility Constraints	35
4.1	Intruder Deduction problem	35
4.1.1	Preliminaries	35
4.1.2	Decidability via Locality	37
4.2	Deducibility constraints	38

4.3	Decision Procedure	39
4.3.1	Simplification Rules	39
4.3.2	Completeness	40
4.3.3	Complexity	43
4.4	Further Readings	43
4.5	Exercises	43
5	Unbounded process verification	47
5.1	Undecidability	47
5.2	Structure and Main Features of ProVerif	48
5.3	A Formal Model of Security Protocols	48
5.3.1	Syntax and Informal Semantics	48
5.3.2	Example	51
5.3.3	Formal Semantics	52
5.3.4	Security Properties	52
5.3.5	Some Other Models	58
5.4	The Horn Clause Representation of Protocols	58
5.4.1	Definition of this Representation	58
5.4.2	Resolution Algorithm	64
5.5	Translation from the Pi Calculus	71
5.5.1	Clauses for the Attacker	71
5.5.2	Clauses for the Protocol	71
5.5.3	Extension to Equational Theories	73
5.5.4	Extension to Scenarios with Several Phases	75
5.6	Extension to Correspondences	75
5.6.1	From Secrecy to Correspondences	75
5.6.2	Instrumented Processes	77
5.6.3	Generation of Horn Clauses and Resolution	79
5.6.4	Non-injective correspondences	82
5.6.5	Sketch for injective correspondences	83
5.7	Extension to the Proof of Observational Equivalences	84
5.7.1	Weak Secrets	89
5.7.2	Authenticity	90
5.8	Applications	90
5.9	Further Readings	91
5.10	Exercises	91
6	Equivalence based properties (static equivalence)	95
6.1	Definitions and Applications	95
6.1.1	Static equivalence	95
6.1.2	Applications of static equivalence	96
6.1.3	Some properties of static equivalence	98
6.1.4	Further readings	99
6.2	Procedure for subterm convergent equational theories	100
6.2.1	Preliminaries	100
6.2.2	Deciding $\sim_{\mathcal{E}}$ in polynomial time for subterm convergent equational theories	101
6.2.3	Deciding $\sim_{\mathcal{E}}$ vs deciding $\vdash_{\mathcal{E}}$	106

6.2.4	Further readings	106
6.3	Exercises	107
7	Composition Results	109
7.1	Motivation	109
7.2	Parallel Composition under Shared Keys	110
7.2.1	Theorem	110
7.2.2	Main Steps of the Proof	112
7.2.3	Applications	113
7.3	From One Session to Many	114
7.3.1	Protocol Transformation	114
7.3.2	Composition Result	115
7.3.3	Other ways of tagging	116
7.4	Further Readings	116
7.5	Exercises	116
III	Verification in the Computational Setting	119
8	The Computational Model	121
8.1	Introduction	121
8.1.1	Exact Security and Practical Security	122
8.2	Security Proofs and Security Arguments	122
8.2.1	Computational Assumptions	122
8.2.2	Complexity	123
8.2.3	Practical Security	123
8.2.4	The Random-Oracle Model	124
8.2.5	The General Framework	125
9	Provable Security	127
9.1	Security Notions	127
9.1.1	Public-Key Encryption	127
9.1.2	Digital Signature Schemes	129
9.2	The Computational Assumptions	131
9.2.1	Integer Factoring and the RSA Problem	131
9.2.2	The Discrete Logarithm and the Diffie-Hellman Problems	133
9.3	Proof Methodology	134
9.4	Exercises	135
10	Public-Key Encryption Schemes	137
10.1	Introduction	137
10.1.1	The RSA Encryption Scheme	137
10.1.2	The El Gamal Encryption Scheme	137
10.2	The Cramer-Shoup Encryption Scheme	138
10.2.1	Description	138
10.2.2	Security Analysis	138
10.3	A Generic Construction	143

10.3.1	Description	143
10.3.2	Security Analysis	143
10.4	OAEP: the Optimal Asymmetric Encryption Padding	146
10.4.1	Description	146
10.4.2	About the Security	147
10.4.3	The Actual Security of OAEP	148
10.4.4	Intuition behind the Proof of Security	149
11	Digital Signature Schemes	151
11.1	Introduction	151
11.2	Some Schemes	152
11.2.1	The RSA Signature Scheme	152
11.2.2	The Schnorr Signature Scheme	152
11.3	DL-Based Signatures	153
11.3.1	General Tools	154
11.3.2	No-Message Attacks	154
11.3.3	Chosen-Message Attacks	157
11.4	RSA-Based Signatures	158
11.4.1	Basic Proof of the FDH Signature	159
11.4.2	Improved Security Result	161
11.4.3	PSS: The Probabilistic Signature Scheme	162
12	Automating Game-Based Proofs	165
IV	Links between the two Settings	167
13	Soundness of Static Equivalence	169
13.1	Abstract and computational algebras	169
13.1.1	Abstract algebras.	169
13.1.2	Abstract and concrete models.	169
13.2	Soundness for symmetric encryption	170
13.3	Proof of Proposition 13.1	172
13.3.1	Pattern equivalence.	172
13.3.2	Soundness of \equiv_{pat}	174
14	Soundness in presence of an Adaptive Adversary	179
14.1	Adaptive Adversary	179
14.2	Further readings	183
15	Soundness in presence of an Active Adversary	185
V	Implementation	187

Chapter 1

Introduction

We start with the well-known example of the so-called “Needham-Schroeder public-key protocol” [?], that has been designed in 1978 and for which an attack was found in 1996 by G. Lowe [?], using formal methods. This attack is completely independent of the security of cryptographic primitives and of the actual implementation: the attack works even for perfect cryptography and a faithful implementation. The protocol was corrected by G. Lowe himself [?]. The new protocol has been proved secure in a formal model by several authors. However, there is still an attack on this protocol, which we present next. The idea of the attack was suggested by J. Millen and is reported by B. Warinschi in [?].

The reason of the apparent paradox (the protocol is proved secure and there is an attack) is the discrepancy between the models that are used for the proof and for the attack respectively. More precisely, the security proof (that we will also give in these notes) assume a *perfect cryptography*. Such a perfect cryptography does not exist in practice. Based on some hardness assumptions, there are however many “provably secure” encryption schemes. In our example, we are using public-key cryptography and the attack on the corrected protocol is based on a property of the El-Gamal encryption scheme. The El-Gamal public-key encryption scheme is provably secure for *chosen plaintext attacks (IND-CPA)* [?] (assuming the hardness of some discrete logarithm problem). However, it is not secure for some stronger notions of security, for instance the *Chosen Ciphertext Attacks (IND-CCA)*. And it turns out that Lowe’s protocol model for public key cryptography is only sound for IND-CCA secure encryption schemes. In summary, we have two models. In the more abstract one there is a security proof, while in the less abstract one there is an attack. The two models are actually equivalent, provided that the encryption scheme is secure in a strong sense.

Similar problems may occur in actual implementations. A protocol can be proven secure while there is an attack on its actual implementation. Similarly, the soundness of the abstract proof of security requires some properties (typically memory access) of the implementations.

The ambition of this lecture is to cover all these aspects:

- Give a formal symbolic model for security protocols and proof techniques at this level of abstraction.
- Give a computational model of security protocols and proof techniques at this level.
- Provide with soundness results, showing under which assumptions symbolic proofs provide with computational security.

- Provide with correctness results of implementations, that show under which assumptions either the formal or the computational proof yields a security proof of the implementation.

Part I

Modelling Protocols and Security Properties

Chapter 2

An Introductory Example

We start with the well-known example of the so-called “Needham-Schroeder public-key protocol” [?], that has been designed in 1978 and for which an attack was found in 1996 by G. Lowe [?], using formal methods.

2.1 An Informal Description

The protocol is a so-called “mutual authentication protocol”. Two parties A and B wish to agree on some value, *e.g.* they wish to establish a shared secret that they will use later for fast confidential communication. The parties A and B only use a public communication channel (for instance a postal service, Internet or a mobile phone). The transport of the messages on such channels is insecure. Indeed, a malicious agent might intercept the letter (resp. message) look at its content and possibly replace it with another message or even simply destroy it.

In order to secure their communication, the agents use lockers (or encryption). We consider here public-key encryption: the lockers can be reproduced and distributed, but the key to open them is owned by a single person. Encrypting a message m with the public key of A is written $\{m\}_{\text{pk}(A)}$ whereas concatenating two messages m_1 and m_2 is written $\langle m_1, m_2 \rangle$. An informal description of the protocol in the so-called Alice-Bob notation is given in Figure 2.1.

1. $A \rightarrow B : \quad \{\langle A, N_A \rangle\}_{\text{pk}(B)}$
2. $B \rightarrow A : \quad \{\langle N_A, N_B \rangle\}_{\text{pk}(A)}$
3. $A \rightarrow B : \quad \{N_B\}_{\text{pk}(B)}$

Figure 2.1: Informal description of the Needham-Schroeder public key protocol

Description. First the agent A encrypts a nonce N_A , *i.e.* a random number freshly generated, and her identity with the public key of B and sends it on the public channel (message 1). Only the agent B , who owns the corresponding private key can open this message. Upon reception, he gets N_A , generates his own nonce N_B and sends back the pair encrypted with the public key of A (message 2). Only the agent A is able to open this message. Furthermore, since only B was able to get N_A , inserting N_A in the plaintext is a witness that it comes from the agent B . Finally, A , after decrypting, checks that the first component is N_A and retrieves the second component N_B . As an acknowledgement, she sends back N_B encrypted by the

- | | |
|--|---|
| <ol style="list-style-type: none"> 1. $A \rightarrow D : \quad \{\langle A, N_A \rangle\}_{\text{pk}(D)}$ 2. $B \rightarrow A : \quad \{\langle N_A, N_B \rangle\}_{\text{pk}(A)}$ 3. $A \rightarrow D : \quad \{N_B\}_{\text{pk}(D)}$ | <ol style="list-style-type: none"> 1'. $D(A) \rightarrow B : \quad \{\langle A, N_A \rangle\}_{\text{pk}(B)}$ 2'. $B \rightarrow A : \quad \{\langle N_A, N_B \rangle\}_{\text{pk}(A)}$ 3'. $D(A) \rightarrow B : \quad \{N_B\}_{\text{pk}(B)}$ |
|--|---|

Figure 2.2: Attack on the Needham-Schroeder public key protocol

public key of B (message 3). When B receives this message, he checks that the content is N_B . If this succeeds, it is claimed that, if the agents A and B are honest, then both parties agreed on the nonces N_A and N_B (they share these values). Moreover, these values are secret: they are only known by the agents A and B .

Attack. Actually, an attack was found in 1996 by G. Lowe [?] on the Needham-Schroeder public-key protocol. The attack described in Figure 2.2 relies on the fact that the protocol can be used by several parties. Moreover, we have to assume that an honest agent A starts a session of the protocol with a dishonest agent D (message 1). Then D , impersonating A , sends a message to B , starting another instance of the protocol (message 1'). When B receives this message, supposedly coming from A , he answers (messages 2' & 2). The agent A believes this reply comes from C , hence she continues the protocol (message 3). Now, the dishonest agent D decrypts the ciphertext and learn the nonce N_B . Finally, D is able to send the expected reply to B (message 3'). At this stage, two instances of the protocol have been completed with success. In the second instance B believes that he is communicating with A : contrarily to what is expected, A and B do not agree on N_B . Moreover, N_B is not a secret shared only between A and B .

Fixed version of the protocol. It has been proposed to fix the protocol by including the respondent's identity in the response (see Figure 2.3).

1. $A \rightarrow B : \quad \{\langle A, N_A \rangle\}_{\text{pk}(B)}$
2. $B \rightarrow A : \quad \{\langle \langle N_A, N_B \rangle, B \rangle\}_{\text{pk}(A)}$
3. $A \rightarrow B : \quad \{N_B\}_{\text{pk}(B)}$

Figure 2.3: Description of the Needham-Schroeder-Lowe protocol

The attack described above cannot be mounted in the corrected version of the protocol. Actually, it is reported in [?] that the technique that permitted to find the Lowe attack on the Needham-Schroeder public key protocol found no attack on this protocol.

2.2 A More Formal Analysis

The Alice-Bob notation is a semi-formal notation that specifies the conversation between the agents. We have to make more precise the view of each agent. This amounts specifying the concurrent programs that are executed by each party. One has also to be precise when specifying how a message is processed by an agent. In particular, what parts of a received

message are checked by the agent? What are the actions performed by the agent to compute the answer?

A classical way to model protocols is to use a process algebra. However, in order to model the messages that are exchanged, we need a process algebra that allows processes to send first-order terms build over a signature, names and variables. These terms model messages that are exchanged during a protocol.

Example 2.1 Consider for example the signature $\Sigma = \{\{-\}_-, \text{pk}_-, \text{sk}(-), \text{dec}(\langle -, - \rangle), \text{proj}_1, \text{proj}_2\}$ which contains three binary function symbols modelling asymmetric encryption, decryption, and pairing, and four unary function symbols modelling projections, public key and private key. The signature is equipped with an equational theory and we interpret equality up to this theory. For instance the theory

$$\text{dec}(\{x\}_{\text{pk}(y)}, \text{sk}(y)) = x, \text{proj}_1(\langle x_1, x_2 \rangle) = x_1, \text{ and } \text{proj}_2(\langle x_1, x_2 \rangle) = x_2.$$

models that decryption and encryption cancel out whenever suitable keys are used. One can also retrieve the first (resp. second) component of a pair.

Processes P, Q, R, \dots are constructed as follows. The process $\text{new } N.P$ restricts the name N in P and can for instance be used to model that N is a fresh random number. $\text{in}(c, x).P$ models the input of a term on a channel c , which is then substituted for x in process P . $\text{out}(c, t)$ outputs a term t on a channel c . The conditional $\text{if } M = N \text{ then } P \text{ else } Q$ behaves as P when M and N are equal modulo the equational theory and behaves as Q otherwise.

The program (or process) that is executed by an agent, say a , who wants to initiate a session of the Needham-Schroeder protocol with another agent b is as follows:

$A(a, b) \hat{=}$	$\text{new } N_a.$ $\text{out}(c, \{a, N_a\}_{\text{pk}(b)}).$ $\text{in}(c, x).$ $\text{let } x_0 = \text{dec}(x, \text{sk}(a)) \text{ in}$ $\text{if } \text{proj}_1(x_0) = N_a \text{ then}$ $\text{let } x_1 = \text{proj}_2(x_0) \text{ in}$ $\text{out}(c, \{x_1\}_{\text{pk}(b)})$	a generates a fresh message N_a the message is sent on the channel c a is waiting for an input on c a tries to decrypt the message a checks that the first component is N_a a retrieves the second component a sends her answer on c
-------------------	---	--

Note that we use variables for the unknown components of messages. These variables can be (a priori) replaced by any message, provided that the attacker can build it and that it is accepted by the agent. In the program described above, if the decryption fails or if the first component of the message received by a is not equal to N_a , then a will abort the protocol.

Similarly, we have to write the program that is executed by an agent, say b , who has to answer to the messages sent by the initiator of the protocol. This program may look like this:

$B(a, b) \hat{=}$	$\text{in}(c, y).$ $\text{let } (a, y_0) = \text{dec}(y, \text{sk}(b)) \text{ in}$ $\text{new } N_b.$ $\text{out}(c, \{y_0, N_b\}_{\text{pk}(a)}).$ $\text{in}(c, y').$ $\text{if } \text{dec}(y', \text{sk}(b)) = N_b \text{ then Ok.}$	b is waiting for an input on c b tries to decrypt it and then retrieves the second component of the plaintext b generated a fresh random number N_b b sends his answer on the channel c b is waiting for an input on c b tries to decrypt the message and he checks whether its content is N_b or not
-------------------	---	--

The (weak) secrecy property states for instance that, if a, b are honest (their secret keys are unknown to the environment), then, when the process $B(a, b)$ reaches the Ok state, N_b is unknown to the environment. We will also see later how to formalise agreement properties. The “environment knowledge” is actually a component of the description of the global state of the network. Basically, all messages that can be built from the public data and the messages that have been sent are in the knowledge of the environment.

Any number of copies of A and B (with any parameter values) are running concurrently in a hostile environment. Such a hostile environment is modelled by any process that may receive and emit on public channels. We also assume that such an environment owes as many public/private key pairs as it wishes (compromised agents), an agent may also generate new values when needed. The only restrictions on the environment is on the way it may construct new messages: the encryption and decryption functions, as well as public keys are assumed to be known from the environment. However no private key (besides those that it generates) is known. We exhibit now a process that will yield the attack, assuming that the agent d is a dishonest (or compromised) agent who leaked his secret key:

$P \hat{=}$	$\text{in}(c, z_1).$	d receives a message (from a)
	$\text{let } \langle a, z'_1 \rangle = \text{dec}(z_1, \text{sk}(d)) \text{ in}$	d decrypts it
	$\text{out}(c, \{\langle a, z'_1 \rangle\}_{\text{pk}(b)}).$	d sends the plaintext encrypted with $\text{pk}(b)$
	$\text{in}(c, z_2).\text{out}(c, z_2).$	d forwards to a the answer he obtained from b
	$\text{in}(c, z_3).$	d receives the answer from a
	$\text{let } z'_3 = \text{dec}(z_3, \text{sk}(d)) \text{ in}$	d decrypts it and learn N_b
	$\text{out}(c, \{z'_3\}_{\text{pk}(b)}).$	d sends the expected message $\{N_b\}_{\text{pk}(b)}$ to b .

The Needham-Schroeder-Lowe protocol has been proved secure in several formal models close to the one we have sketched in this section [?, ?].

2.3 An Attack on the Fixed Version of the Protocol

Up to now, the encryption is a black-box: nothing can be learnt on a plaintext from a ciphertext and two ciphertexts are unrelated.

Consider however a simple El-Gamal encryption scheme. Roughly (we skip here the group choice for instance), the encryption scheme is given by a cyclic group G of order q and generator g ; these parameters are public. Each agent a may choose randomly a secret key $\text{sk}(a)$ and publish the corresponding public key $\text{pk}(a) = g^{\text{sk}(a)}$. Given a message m (assume for simplicity that it is an element $g^{m'}$ of the group), encrypting m with the public key $\text{pk}(a)$ consists in drawing a random number r and letting $\{m\}_{\text{pk}(a)} = (\text{pk}(a)^r \times g^{m'}, g^r)$. Decrypting the message consists in raising g^r to the power $\text{sk}(a)$ and dividing the first component of the pair by $g^{r \times \text{sk}(a)}$. We have that:

$$[\text{pk}(a)^r \times g^{m'}] / (g^r)^{\text{sk}(a)} = [(g^{\text{sk}(a)})^r \times g^{m'}] / (g^r)^{\text{sk}(a)} = g^{m'} = m.$$

This means that this encryption scheme satisfies the equation $\text{dec}(\{x\}_{\text{pk}(y)}, \text{sk}(y)) = x$. However, as we will see, this encryption scheme also satisfies some other properties that are not taken into account in our previous formal analysis.

1. $a \rightarrow d : \{\langle a, N_a \rangle\}_{\text{pk}(d)}$
 - 1'. $d(a) \rightarrow b : \{\langle a, N_a \rangle\}_{\text{pk}(b)}$
 - 2'. $b \rightarrow a : \{\langle \langle N_a, N_b \rangle, b \rangle\}_{\text{pk}(a)} = (g^{N_a+2^\alpha \times N_b+2^{2\alpha} \times b} \times \text{pk}(a)^r, g^r)$
- d intercepts this message, and computes
$$[g^{N_a+2^\alpha \times N_b+2^{2\alpha} \times b} \times \text{pk}(a)^r] \times g^{-2^{2\alpha} \times b} \times g^{2^{2\alpha} \times d} = g^{N_a+2^\alpha \times N_b+2^{2\alpha} \times d} \times \text{pk}(a)^r$$
2. $d \rightarrow a : \{\langle \langle N_a, N_b \rangle, d \rangle\}_{\text{pk}(a)} = (g^{N_a+2^\alpha \times N_b+2^{2\alpha} \times d} \times \text{pk}(a)^r, g^r)$
3. $a \rightarrow d : \{N_b\}_{\text{pk}(d)}$
 - 3'. $d \rightarrow b : \{N_b\}_{\text{pk}(b)}$

Figure 2.4: Attack on the Needham-Schroeder-Lowe protocol with El-Gamal encryption.

Attack. Assume now that we are using such an encryption scheme in the Needham-Schroeder-Lowe protocol and that pairing two group elements $m_1 = g^{m'_1}$ and $m_2 = g^{m'_2}$ is performed in a naive way: $\langle m_1, m_2 \rangle$ is mapped to $g^{m'_1+2^{|m'_1|} \times m'_2}$ (*i.e.* concatenating the binary representations of the messages m'_1 and m'_2). In such a case, an attack can be mounted on the protocol (see Figure 2.4).

Actually, the attack starts as before. We assume that the honest agent a is starting a session with a dishonest party d . Then d decrypts the message and re-encrypt it with the public key of b . The honest party b replies sending the expected message $\{\langle \langle N_a, N_b \rangle, b \rangle\}_{\text{pk}(a)}$. The attacker intercepts this message. Note that the attacker can not simply forward it to a since it does not have the expected form. The attacker intercepts $\{\langle \langle N_a, N_b \rangle, b \rangle\}_{\text{pk}(a)}$, *i.e.* $(\text{pk}(a)^r \times g^{N_a+2^\alpha \times N_b+2^{2\alpha} \times b}, g^r)$ where α is the length of a nonce. The attacker knows g, α, b , hence he can compute $g^{-2^{2\alpha} \times b} \times g^{2^{2\alpha} \times d}$ and multiply the first component, yielding $\{\langle \langle N_a, N_b \rangle, d \rangle\}_{\text{pk}(a)}$. Then the attack can go on as before: a replies by sending $\{N_b\}_{\text{pk}(d)}$ and the attacker sends $\{N_b\}_{\text{pk}(b)}$ to b , impersonating a .

This example is however a toy example since pairing could be implemented in another way. In [?] there is a real attack that is only based on weaknesses of the El Gamal encryption scheme. In particular, the attack does not depend on how pairing is implemented.

This shows that the formal analysis only proves the security in a formal model, that might not be faithful. Here, the formal analysis assumed a model in which it is not possible to forge a ciphertext from another ciphertext, without decrypting/encrypting. This property is known as *non-malleability*, which is not satisfied by the El Gamal encryption scheme.

2.4 Further Readings

A survey by Clark and Jacob [?] describes several authentication protocols and outlines also the methods that have been used to analyse them. In addition, it provides a summary of the ways in which protocols have been found to fail. The purpose of the SPORE web page [?] is to continue on-line the seminal work of Clark and Jacob, updating their base of security protocols.

As you have seen, some protocols (or some attacks) rely on some algebraic properties of cryptographic primitives. In [?], a list of some relevant algebraic properties of cryptographic

operators is given, and for each of them, some examples of protocols or attacks using these properties are provided. The survey also gives an overview of the existing methods in formal approaches for analysing cryptographic protocols.

2.5 Exercises

Exercise 2.1 (★) *Consider the following protocol:*

$$\begin{aligned} A \rightarrow B &: \langle A, \{K\}_{\text{pk}(B)} \rangle \\ B \rightarrow A &: \langle B, \{K\}_{\text{pk}(A)} \rangle \end{aligned}$$

First, A generates a fresh key K and sends it encrypted with the public key of B. Only B will be able to decrypt this message. In this way, B learns K and B also knows that this message comes from A as indicated in the first part of the message he received. Hence, B answers to A by sending again the key, this time encrypted with the public key of A.

Show that an attacker can learn the key K generated by an honest agent A to another honest agent B.

Exercise 2.2 (★) *The previous protocol is corrected as in the Needham-Schroeder protocol, i.e. we add the identity of the agent inside each encryption.*

$$\begin{aligned} A \rightarrow B &: \{\langle A, K \rangle\}_{\text{pk}(B)} \\ B \rightarrow A &: \{\langle B, K \rangle\}_{\text{pk}(A)} \end{aligned}$$

1. *Check that the previous attack does not exist anymore. Do you think that the secrecy property stated in Exercise 2.1 holds?*
2. *Two agents want to use this protocol to establish a session key. Show that there is an attack.*

Exercise 2.3 (★★) *For double security, all messages in the previous protocol are encrypted twice:*

$$\begin{aligned} A \rightarrow B &: \{\{\langle A, \{K\}_{\text{pk}(B)} \rangle\}_{\text{pk}(B)}\}_{\text{pk}(A)} \\ B \rightarrow A &: \{\{\langle B, \{K\}_{\text{pk}(A)} \rangle\}_{\text{pk}(A)}\}_{\text{pk}(B)} \end{aligned}$$

Show that the protocol then becomes insecure in the sense that an attacker can learn the key K generated by an honest agent A to another honest agent B.

Exercise 2.4 (★★★) *We consider a variant of the Needham-Schroeder-Lowe protocol. This protocol is as follows:*

$$\begin{aligned} 1. A \rightarrow B &: \{\langle A, N_A \rangle\}_{\text{pk}(B)} \\ 2. B \rightarrow A &: \{\langle N_A, \langle N_B, B \rangle \rangle\}_{\text{pk}(A)} \\ 3. A \rightarrow B &: \{N_B\}_{\text{pk}(B)} \end{aligned}$$

1. *Check that the 'man-in-the-middle' attack described in Figure 2.2 does not exist.*
2. *Show that there is an attack on the secrecy of the nonce N_b .*
hint: type confusion
3. *Do you think that this attack is realistic? Why?*

Chapter 3

A Small Process Calculus

We now define our cryptographic process calculus for describing protocols. This calculus is inspired by the applied pi calculus [?] which is the calculus used by the PROVERIF tool [?]. The applied pi calculus is a language for describing concurrent processes and their interactions. It is an extension of the pi calculus [?] with cryptographic primitives. It is designed for describing and analysing a variety of security protocols, such as authentication protocols (*e.g.* [?]), key establishment protocols (*e.g.* [?]), e-voting protocols (*e.g.* [?]), ... These protocols try to achieve various security goals, such as secrecy, authentication, privacy, ...

In this chapter, we present a simplified version that is sufficient for our purpose and we explain how to formalise security properties in such a calculus.

3.1 Preliminaries

The applied pi calculus is similar to the spi calculus [?]. The key difference between the two formalisms concerns the way that cryptographic primitives are handled. The spi calculus has a fixed set of primitives built-in (symmetric and public key encryption), while the applied pi calculus allows one to define less usual primitives by means of an equational theory. This flexibility is particularly useful to model the new protocols that are emerging and which rely on new cryptographic primitives.

3.1.1 Messages

To describe processes, one starts with an infinite set of *names* \mathcal{N} (which are used to represent atomic data, such as keys, nonces, ...), an infinite set of *variables* \mathcal{X} , and a *signature* \mathcal{F} which consists of the *function symbols* which will be used to define *terms*. Each function symbol has an associated integer, its *arity*. In the case of security protocols, typical function symbols will include a binary function symbol `senc` for symmetric encryption, which takes plaintext and a key and returns the corresponding ciphertext, and a binary function symbol `sdec` for decryption, taking ciphertext and a key and returning the plaintext. Variables are used to consider messages containing unknown (unspecified) pieces.

Terms are defined as names, variables, and function symbols applied to other terms. Terms and function symbols may be sorted, and in such a case, function symbol application must respect sorts and arities. We denote by $\mathcal{T}(\Sigma)$ the set of terms built on the symbols in Σ . We denote by $fv(M)$ (resp. $fn(M)$) the set of variables (resp. names) that occur in M . A term

M that does not contain any variable is a *ground term*. The set of positions of a term T is written $pos(T) \subseteq \mathbb{N}^*$, and its set of subterms $st(T)$. The subterm of T at position $p \in pos(T)$ is written $T|_p$. The term obtained by replacing $T|_p$ with a term U in T is denoted $T[U]_p$.

We split the function symbols between *private* and *public* symbols, *i.e.* $\mathcal{F} = \mathcal{F}_{\text{pub}} \uplus \mathcal{F}_{\text{priv}}$. Private function symbols are used to model algorithms or data that are not available to the attacker. Moreover, sometimes, we also split the function symbols into *constructors* and *destructors*, *i.e.* $\mathcal{F} = \mathcal{D} \uplus \mathcal{C}$. Destructors are used to model the fact that some operations fail. A typical destructor symbol could be the symbol `sdec` if we want to model a decryption algorithm that fails when we try to decrypt a ciphertext with a wrong key. A *constructor term* is a term in $\mathcal{T}(\mathcal{C} \cup \mathcal{N} \cup \mathcal{X})$.

By the means of a convergent term rewriting system \mathcal{R} , we describe the equations which hold on terms built from the signature. A *term rewriting system* (TRS) is a set of *rewrite rules* $l \rightarrow r$ where $l \in \mathcal{T}(\mathcal{F} \cup \mathcal{X})$ and $r \in \mathcal{T}(\mathcal{F} \cup \text{fv}(l))$. A term $S \in \mathcal{T}(\mathcal{F} \cup \mathcal{N} \cup \mathcal{X})$ rewrites to T by \mathcal{R} , denoted $S \rightarrow_{\mathcal{R}} T$, if there is $l \rightarrow r$ in \mathcal{R} , $p \in pos(S)$ and a substitution σ such that $S|_p = l\sigma$ and $T = S[r\sigma]_p$. Moreover, we assume that $\{x\sigma \mid x \in \text{Dom}(\sigma)\}$ are constructor terms. We denote by $\rightarrow_{\mathcal{R}}^*$ the reflexive and transitive closure of $\rightarrow_{\mathcal{R}}$, and by $=_{\mathcal{R}}$ the symmetric, reflexive and transitive closure of $\rightarrow_{\mathcal{R}}$. A TRS \mathcal{R} is *convergent* if it is:

- *terminating*, *i.e.* there is no infinite chain $T_1 \rightarrow_{\mathcal{R}} T_2 \rightarrow_{\mathcal{R}} \dots$; and
- *confluent*, *i.e.* for all terms S, T such that $S =_{\mathcal{R}} T$, there exists U such that $S \rightarrow_{\mathcal{R}}^* U$ and $T \rightarrow_{\mathcal{R}}^* U$.

A term T is \mathcal{R} -*reduced* if there is no term S such that $T \rightarrow_{\mathcal{R}} S$. If $T \rightarrow_{\mathcal{R}}^* S$ and S is \mathcal{R} -reduced then S is a \mathcal{R} -*reduced form* of T . When this reduced form is unique (in particular if \mathcal{R} is convergent), we write $S = T \downarrow_{\mathcal{R}}$ (or simply $T \downarrow$ when \mathcal{R} is clear from the context). In the following, we will only consider convergent rewriting system. Hence, we have that $M =_{\mathcal{R}} N$, if and only if, $M \downarrow = N \downarrow$. A ground constructor term in normal form is also called a *message*.

Example 3.1 *In order to model the handshake protocol that we will present later on, we introduce the signature:*

$$\mathcal{F}_{\text{senc}} = \{\text{senc}/2, \text{sdec}/2, \text{f}/1\}$$

together with the term rewriting system $\mathcal{R}_{\text{senc}} = \{\text{sdec}(\text{senc}(x, y), y) \rightarrow x\}$. We will assume that $\mathcal{F}_{\text{senc}}$ only contains constructor symbols. This represents a decryption algorithm that always succeeds. If we decrypt the ciphertext `senc`(n, k) with a key $k' \neq k$, the decryption algorithm will return the message `sdec`(`senc`(n, k), k').

Here, we have that `sdec`(`senc`($n', \text{sdec}(n, n)$), `sdec`(n, n)) $=_{\mathcal{R}}$ n' . Indeed, we have that `sdec`(`senc`($n', \text{sdec}(n, n)$), `sdec`(n, n)) rewrites in one step to n' (with $p = \epsilon$, and $\sigma = \{x \mapsto n', y \mapsto \text{sdec}(n, n)\}$).

Example 3.2 *In order to model the Needham-Schroeder protocol, we will consider the following signature:*

$$\mathcal{F}_{\text{aenc}} = \{\langle -, - \rangle, \text{proj}_1/1, \text{proj}_2/1, \text{aenc}/2, \text{pk}/1, \text{sk}/1, \text{adec}/2\}$$

together with the term rewriting system $\mathcal{R}_{\text{aenc}}$:

$$\text{proj}_1(\langle x, y \rangle) \rightarrow x \quad \text{proj}_2(\langle x, y \rangle) \rightarrow y \quad \text{adec}(\text{aenc}(x, \text{pk}(y)), \text{sk}(y)) \rightarrow x$$

This will allow us to model asymmetric encryption and pairing. We will assume that proj_1 , proj_2 , and adec are destructor symbols. The only private non-constant symbol is the symbol sk . Note that $\text{proj}_1(\langle n, \text{adec}(n, n) \rangle) \neq_{\mathcal{R}} n$. Indeed, the terms $\text{proj}_1(\langle n, \text{adec}(n, n) \rangle)$ and n are both irreducible and not syntactically equal.

3.1.2 Assembling Terms into Frames

At some moment, while engaging in one or more sessions of one or more protocols, an attacker may have observed a sequence of messages M_1, \dots, M_ℓ , *i.e.* a set of ground constructor terms in normal form. We want to represent this knowledge of the attacker. It is not enough for us to say that the attacker knows the *set* of terms $\{M_1, \dots, M_\ell\}$ since he also knows the order that he observed them in. Furthermore, we should distinguish those names that the attacker knows from those that were freshly generated by others and which remain secret from the attacker; both kinds of names may appear in the terms. We use the concept of *frame* from the applied pi calculus [?] to represent the knowledge of the attacker. A *frame* $\phi = \text{new } \bar{n}.\sigma$ consists of a finite set $\bar{n} \subseteq \mathcal{N}$ of *restricted* names (those that the attacker does not know), and a substitution σ of the form:

$$\{M_1/x_1, \dots, M_\ell/x_\ell\}.$$

The variables enable us to refer to each message M_i . We always assume that the terms M_i are ground term in normal form that do not contain destructor symbols. The names \bar{n} are bound and can be renamed. We denote by $=_\alpha$ the α -renaming relation on frames. The *domain* of the frame ϕ , written $\text{Dom}(\phi)$, is defined as $\{x_1, \dots, x_\ell\}$.

3.1.3 Deduction

Given a frame ϕ that represents the information available to an attacker, we may ask whether a given ground constructor term M may be deduced from ϕ . Given a convergent rewriting system \mathcal{R} on \mathcal{F} , this relation is written $\phi \vdash_{\mathcal{R}} M$ and is formally defined below.

Definition 3.1 (Deduction) *Let M be a ground term and $\phi = \text{new } \bar{n}.\sigma$ be a frame. We have that $\text{new } \bar{n}.\sigma \vdash_{\mathcal{R}} M$ if, and only if, there exists a term $N \in \mathcal{T}(\mathcal{F}_{\text{pub}} \cup \mathcal{N} \cup \text{Dom}(\phi))$ such that $\text{fn}(N) \cap \bar{n} = \emptyset$ and $N\sigma =_{\mathcal{R}} M$. Such a term N is a recipe of the term M .*

Intuitively, the deducible messages are the messages of ϕ and the names that are not protected in ϕ , closed by rewriting with \mathcal{R} and closed by application of public function symbols. When $\text{new } \bar{n}.\sigma \vdash_{\mathcal{R}} M$, any occurrence of names from \bar{n} in M is bound by $\text{new } \bar{n}$. So $\text{new } \bar{n}.\sigma \vdash_{\mathcal{R}} M$ could be formally written $\text{new } \bar{n}.\sigma \vdash_{\mathcal{R}} M$.

Example 3.3 *Consider the theory $\mathcal{R}_{\text{senc}}$ given in Example 3.1 and the following frame:*

$$\phi = \text{new } k, s_1. \{^{\text{senc}}(\langle s_1, s_2 \rangle, k) / x_1, k / x_2\}.$$

We have that $\phi \vdash_{\mathcal{R}_{\text{senc}}} k$, $\phi \vdash_{\mathcal{R}_{\text{senc}}} s_1$ and $\phi \vdash_{\mathcal{R}_{\text{senc}}} s_2$. Indeed x_2 , $\text{proj}_1(\text{sdec}(x_1, x_2))$ and s_2 are recipes of the terms k , s_1 and s_2 respectively.

The relation $\text{new } \bar{n}.\sigma \vdash_{\mathcal{R}} M$ can be axiomatized by the following rules:

$$\begin{array}{c} \frac{}{\text{new } \bar{n}.\sigma \vdash_{\mathcal{R}} M} \quad \text{if } \exists x \in \text{dom}(\sigma) \text{ such that } x\sigma = M \\[10pt] \frac{\phi \vdash_{\mathcal{R}} M_1 \quad \dots \quad \phi \vdash_{\mathcal{R}} M_\ell}{\phi \vdash_{\mathcal{R}} f(M_1, \dots, M_\ell)} \quad f \in \mathcal{F}_{\text{pub}} \\[10pt] \frac{}{\text{new } \bar{n}.\sigma \vdash_{\mathcal{R}} s} \quad s \in \mathcal{N} \setminus \bar{n} \\[10pt] \frac{\phi \vdash_{\mathcal{R}} M}{\phi \vdash_{\mathcal{R}} M'} \quad M =_{\mathcal{R}} M' \end{array}$$

Since we only consider convergent rewriting systems, it is easy to prove that the two definitions coincide.

3.1.4 Static Equivalence

The frames we have introduced are too fine-grained as representations of the attacker's knowledge. For example, $\nu k.\{\text{senc}(s_0, k)/x\}$ and $\nu k.\{\text{senc}(s_1, k)/x\}$ represent a situation in which the encryption of the public name s_0 (resp. s_1) by a randomly-chosen key has been observed. Since the attacker cannot detect the difference between these two situations, the frames should be considered equivalent. To formalise this, we note that if two recipes M, N on the frame ϕ produce the same constructor term, we say they are equal in the frame, and write $(M =_{\mathcal{R}} N)\phi$. Thus, the knowledge of the attacker can be thought of as his ability to distinguish such recipes. If two frames have identical distinguishing power, then we say that they are *statically equivalent*.

Definition 3.2 (static equivalence) *We say that two terms M and N in $\mathcal{T}(\mathcal{F}_{\text{pub}} \cup \mathcal{N} \cup \mathcal{X})$ are equal in the frame ϕ , and write $(M =_{\mathcal{R}} N)\phi$, if there exists \bar{n} and a substitution σ such that $\phi =_{\alpha} \nu \bar{n}.\sigma$, $\bar{n} \cap (\text{fn}(M) \cup \text{fn}(N)) = \emptyset$, and $M\sigma \downarrow$ and $N\sigma \downarrow$ are both constructor terms that are equal, i.e. $M\sigma \downarrow = N\sigma \downarrow$.*

We say that two frames $\phi_1 = \bar{n}_1.\sigma_1$ and $\phi_2 = \bar{n}_2.\sigma_2$ are statically equivalent, and write $\phi_1 \sim_{\mathcal{R}} \phi_2$, when:

- $\text{Dom}(\phi_1) = \text{Dom}(\phi_2)$,
- for all term $M \in \mathcal{T}(\mathcal{F}_{\text{pub}} \cup \mathcal{N} \cup \mathcal{X})$ such that $\text{fn}(M) \cap (\bar{n}_1 \cup \bar{n}_2) = \emptyset$, we have that: $M\sigma_1 \downarrow$ is constructor term $\Leftrightarrow M\sigma_2 \downarrow$ is a constructor term.
- for all terms M, N in $\mathcal{T}(\mathcal{F}_{\text{pub}} \cup \mathcal{N} \cup \mathcal{X})$ we have that: $(M =_{\mathcal{R}} N)\phi_1 \Leftrightarrow (M =_{\mathcal{R}} N)\phi_2$.

Note that by definition of \sim , we have that $\phi_1 \sim \phi_2$ when $\phi_1 =_{\alpha} \phi_2$ and we have also that $\text{new } n.\phi \sim \phi$ when n does not occur in ϕ .

Example 3.4 *Consider the rewriting system $\mathcal{R}_{\text{senc}}$ provided in Example 3.1. Consider the frames $\phi = \text{new } k.\{\text{senc}(s_0, k)/x_1, k/x_2\}$, and $\phi' = \text{new } k.\{\text{senc}(s_1, k)/x_1, k/x_2\}$. Intuitively, s_0 and s_1 could be the two possible (public) values of a vote. We have $(\text{sdec}(x_1, x_2) =_{\mathcal{R}_{\text{senc}}} s_0)\phi$ whereas $(\text{sdec}(x_1, x_2) \neq_{\mathcal{R}_{\text{senc}}} s_0)\phi'$. Therefore we have that $\phi \not\sim \phi'$. However, we have that:*

$$\text{new } k.\{\text{senc}(s_0, k)/x_1\} \sim \text{new } k.\{\text{senc}(s_1, k)/x_1\}.$$

Example 3.5 Consider again the rewriting system $\mathcal{R}_{\text{senc}}$ provided in Example 3.1. We have that:

$$\begin{array}{llll}
\text{new } k.\{\text{senc}(0,k)/x\} & \sim & \text{new } k.\{\text{senc}(1,k)/x\} & \\
\{\text{senc}(0,k)/x, \langle 0,k \rangle /y\} & \not\sim & \text{new } k.\{\text{senc}(1,k)/x, \langle 0,k \rangle /y\} & (\text{sdec}(x, \text{proj}_2(y)) \stackrel{?}{=} 0) \\
\text{new } a.\{a/x\} & \sim & \text{new } b.\{b/x\} & \\
\text{new } a.\{a/x\} & \not\sim & \text{new } b.\{b/y\} & (\text{different domains}) \\
\{a/x\} & \not\sim & \{b/x\} & (x \stackrel{?}{=} a)
\end{array}$$

3.2 Protocols

We now described our cryptographic process calculus for describing protocols. For sake of simplicity, we only consider public channels, *i.e.* under the control of the attacker.

3.2.1 Protocol Language

The grammar for *processes* is given below. One has *plain processes* P, Q, R and *extended processes* A, B, C .

Plain processes. Plain processes are formed from the following grammar

$P, Q, R \triangleq$	plain processes
0	null process
$P \parallel Q$	parallel composition
$\text{in}(c, M_i).P$	message input
$\text{out}(c, M_o).P$	message output
$\text{if } M = N \text{ then } P \text{ else } Q$	conditional
$\text{new } n.P$	restriction
$!P$	replication

such that a variable x appears in a term only if the term is in the scope of an input $\text{in}(c, M_i)$ with $x \in \text{fv}(M_i)$. The null process 0 does nothing; $P \parallel Q$ is the parallel composition of P and Q . The replication $!P$ behaves as an infinite number of copies of P running in parallel. The conditional construction $\text{if } M = N \text{ then } P \text{ else } Q$ is standard. We omit $\text{else } Q$ when Q is 0 . The process $\text{in}(c, M_i).P$ is ready to input on the public channel c , then to run P where the variables of M_i are bound by the actual input message. The term M_i is a constructor term with variables. $\text{out}(c, M_o).P$ is ready to output M_o (it may contains some destructors), then to run P . Again, we omit P when P is 0 .

In this definition, we consider both pattern inputs and conditionals, which is redundant in some situations: for any executable process, the patterns can be replaced with conditionals. However, we keep both possibilities, in order to keep some flexibility in writing down the protocols.

Example 3.6 We illustrate our syntax with the well-known handshake protocol that can be informally described as follows:

$$\begin{array}{ll}
A & \rightarrow B : \text{senc}(n, w) \\
B & \rightarrow A : \text{senc}(f(n), w)
\end{array}$$

We rely on the signature given in Example 3.1. The goal of this protocol is to authenticate B from A 's point of view, provided that they share an initial secret w . This is done by a simple challenge-response transaction: A sends a random number (a nonce) encrypted with the shared secret key w . Then, B decrypts this message, applies a given function (for instance $f(n) = n+1$) to it, and sends the result back, also encrypted with w . Finally, the agent A checks the validity of the result by decrypting the message and checking the decryption against $f(n)$. In our calculus, we can model the protocol as $\text{new } w.(P_A \parallel P_B)$ where

- $P_A(w) = \text{new } n. \text{out}(c, \text{senc}(n, w)). \text{in}(c, x). \text{if } \text{sdec}(x, w) = f(n) \text{ then } P$
- $P_B(w) = \text{in}(c, y). \text{out}(c, \text{senc}(f(\text{sdec}(y, w)), w)).$

where P models an application that is executed when P_B has been successfully authenticated. Here, we use the formalism with explicit destructors but we could also used pattern inputs.

Example 3.7 Going back to the Needham-Schroeder public key protocol described in Chapter 2 and considering the signature given in Example 3.2, we have that:

$$\begin{aligned} P_A(a, b) &\hat{=} \text{out}(c, \text{aenc}(\langle a, N_a \rangle, \text{pk}(b))). \\ &\quad \text{in}(c, \text{aenc}(\langle N_a, x \rangle, \text{pk}(a))). \\ &\quad \text{out}(c, \text{aenc}(x, \text{pk}(b))) \\ P_B(a, b) &\hat{=} \text{in}(c, \text{aenc}(\langle a, y \rangle, \text{pk}(b))). \\ &\quad \text{out}(c, \text{aenc}(\langle y, N_b \rangle, \text{pk}(a))). \\ &\quad \text{in}(c, \text{aenc}(N_b, \text{pk}(b))) \end{aligned}$$

Here, we have used pattern inputs. We could also have used the alternative formalism of explicit destructors. With pattern inputs, we do not need in general to used destructors to describe the outputs.

Note that all the processes that can be written in this syntax (in particular the one with pattern inputs) are not necessary meaningful. Some of them will not be executable.

Continuing with the Needham-Schroeder protocol, we may define several execution scenarios:

Example 3.8 (Scenario 1) The following specifies a copy of the role of Alice, played by a , with d and a copy of the role of Bob, played by b , with a , as well as the fact that d is dishonest, hence his secret key is leaked.

$$P_1 \hat{=} (\text{new } N_a. P_A(a, d)) \parallel (\text{new } N_b. P_B(a, b)) \parallel \text{out}(c, \text{sk}(d))$$

Example 3.9 (Scenario 2) Assume that we wish a to execute the role of the initiator, however with any other party, which is specified here by letting the environment give the identity of such another party: the process first receives x_b , that might be bound to any value. The other role is specified in the same way.

$$P_2 \hat{=} (\text{new } N_a. \text{in}(c, x_b). P_A(a, x_b)) \parallel (\text{new } N_b. \text{in}(c, x_a). P_B(x_a, b)) \parallel \text{out}(c, \text{sk}(d))$$

Example 3.10 (Scenario 3) In Example 3.8 and Example 3.9, a was only able to engage the protocol once (and b was only able to engage once in a response). We may wish a (resp. b) be able to execute any number of instances of the role of the initiator (resp. responder).

$$P_3 \hat{=} !(\text{new } N_a. \text{in}(c, x_b). P_A(a, x_b)) \parallel !(\text{new } N_b. \text{in}(c, x_a). P_B(x_a, b)) \parallel \text{out}(c, \text{sk}(d))$$

Example 3.11 (Scenario 4) Finally, in general, the role of the initiator could be executed by any agent, including b and the role of the responder could be executed by any number of agents as well. We specify an unbounded number of parties, engaging in an unbounded number of sessions by:

$$P_4 \triangleq \left\{ \begin{array}{l} !(\text{new } N_a. \text{in}(c, x_a). \text{in}(c, x_b). P_A(x_a, x_b)) \parallel \\ !(\text{new } N_b. \text{in}(c, x_a). \text{in}(c, x_b). P_B(x_a, x_b)) \parallel \text{out}(c, \text{sk}(d)) \end{array} \right.$$

We can imagine other scenarios as well. Verifying security will only be relative to a given scenario.

Extended Processes. Further, we extend processes with active substitutions and restrictions:

$$A, B, C := P \mid A \parallel B \mid \text{new } n.A \mid \{^M/x\}$$

where M is a ground constructor term in normal form. As usual, names and variables have scopes, which are delimited by restrictions and by inputs. We write $fv(A)$, $bv(A)$, $fn(A)$, $bn(A)$ for the sets of free and bound variables (resp. names). Moreover, we require processes to be *name and variable distinct*, meaning that $bn(A) \cap fn(A) = \emptyset$, $bv(A) \cap fv(A) = \emptyset$, and also that any name and variable is bound at most once in A . Note that the only free variables are introduced by active substitutions (the x in $\{^M/x\}$). Lastly, in an extended process, we require that there is at most one substitution for each variable. An *evaluation context* is an extended process with a hole instead of an extended process.

Extended processes built up from the null process, active substitutions using parallel composition and restriction are called *frames* (extending the notion of frame introduced in Section 3.1.2). Given an extended process A we denote by $\phi(A)$ the frame obtained by replacing any embedded plain processes in it with 0.

Example 3.12 Consider the following process:

$$A = \text{new } s, k_1. (\text{out}(c, a) \parallel \{^{\text{senc}(s, k_1)}/x\} \parallel \text{new } k_2. \text{out}(c, \text{senc}(s, k_2))).$$

We have that $\phi(A) = \text{new } s, k_1. (0 \parallel \{^{\text{senc}(s, k_1)}/x\} \parallel \text{new } k_2. 0)$.

3.2.2 Operational Semantics

To formally define the operational semantics of our calculus, we have to introduce three relations, namely *structural equivalence*, *internal reduction*, and *labelled transition*.

Structural Equivalence. Informally, two processes are *structurally equivalent* if they model the same thing, even if the grammar permits different encodings. For example, to describe a pair of processes P_A and P_B running in parallel, we have to write either $P_A \parallel P_B$, or $P_B \parallel P_A$. These two processes are said to be structurally equivalent. More formally, structural equivalence is the smallest equivalence relation closed by application of evaluation contexts and

such that:

$$\begin{array}{lll}
\text{PAR-0} & A \parallel 0 & \equiv A \\
\text{PAR-C} & A \parallel B & \equiv B \parallel A \\
\text{PAR-A} & (A \parallel B) \parallel C & \equiv A \parallel (B \parallel C) \\
\\
\text{NEW-PAR} & A \parallel \text{new } n.B & \equiv \text{new } n.(A \parallel B) \quad n \notin \text{fn}(A) \\
\text{NEW-C} & \text{new } n_1.\text{new } n_2.A & \equiv \text{new } n_2.\text{new } n_1.A
\end{array}$$

Note that the side condition of the rule NEW-PAR is always true on processes that are name and variable distinct. Using structural equivalence, every extended process A can be rewritten to consist of a substitution and a plain process with some restricted names, *i.e.*

$$A \equiv \text{new } \bar{n}.(\{M_1/x_1\} \parallel \dots \parallel \{M_k/x_k\} \parallel P).$$

In particular, any frame can be rewritten as $\text{new } \bar{n}.\sigma$ matching the notion of frame introduced in Section 3.1.2. We note that unlike in the original applied pi calculus, active substitutions cannot “interact” with the extended processes. As we will see in the following, active substitutions record the outputs of a process to the environment. The notion of frames will be particularly useful to define equivalence based security properties such as resistance against guessing attacks and privacy type properties.

Internal Reduction. A process can be executed without contact with its environment, *e.g.* execution of conditionals, or internal communications between processes in parallel. Formally, *internal reduction* is the smallest relation on processes closed under structural equivalence and application of evaluation contexts such that:

$$\begin{array}{ll}
\text{REPL} & !P \xrightarrow{\tau} P' \parallel !P \quad \text{where } P' \text{ is a fresh renaming of } P \\
\\
\text{THEN} & \text{if } M = N \text{ then } P \text{ else } Q \xrightarrow{\tau} P \quad \text{where } M\downarrow = N\downarrow \text{ and } M\downarrow \text{ is a message} \\
\\
\text{ELSE} & \text{if } M = N \text{ then } P \text{ else } Q \xrightarrow{\tau} Q \quad \text{where } M\downarrow \neq N\downarrow \text{ and } M\downarrow, N\downarrow \text{ are messages} \\
\\
\text{COMM} & \text{out}(c, M_1).P_1 \parallel \text{in}(c, M_2).P_2 \xrightarrow{\tau} P_1 \parallel P_2\theta \quad \text{where } \theta \text{ is such that} \\
& \quad \text{Dom}(\theta) = \text{fv}(M_2), M_2\theta\downarrow = M_1\downarrow, \text{ and } M_1\downarrow \text{ is a message.}
\end{array}$$

We write \rightarrow^* for the reflexive and transitive closure of $\xrightarrow{\tau}$. Note that, in some situations, a process of the form $\text{if } M = N \text{ then } P \text{ else } Q$ may block. This happens when $M\downarrow$ (resp. $N\downarrow$) contains some destructors.

Labelled Transition. Communications are synchronous, but (as long as there is no private channel) we can assume that they occur with the environment. We sketch here a labelled transition semantics. The semantics given previously allow us to reason about protocols with an adversary represented by a context. In order to prove that security properties hold for all adversaries, quantification over all contexts is typically required, which can be difficult in practise. The *labelled semantics* aim to eliminate universal quantification of the context. We have two main rules:

$$\begin{array}{ll}
\text{IN} & \text{in}(c, x).P \xrightarrow{\text{in}(c, M)}_{\ell} P\{M/x\} \quad \text{where } M \text{ is a message} \\
\\
\text{OUT} & \text{out}(c, M).P \xrightarrow{\text{out}(c, M\downarrow)}_{\ell} P \parallel \{M\downarrow/x\} \quad \text{where } x \text{ is a fresh variable and } M\downarrow \text{ is a message}
\end{array}$$

The labelled operational semantics is closed by structural equivalence and under some evaluation contexts. Actually, we have that:

$$\frac{A \equiv A' \quad A' \xrightarrow{\alpha}_{\ell} B' \quad B' \equiv B}{A \xrightarrow{\alpha}_{\ell} B} \quad \frac{A \xrightarrow{\alpha}_{\ell} B}{C[A] \xrightarrow{\alpha}_{\ell} C[B]}$$

where C is an evaluation context, and in case of an input, i.e. $\alpha = \text{in}(c, M)$, we have that $\phi(C[A]) \vdash_{\mathcal{R}} M$.

We write \rightarrow_{ℓ} to denote $\xrightarrow{\tau} \cup \xrightarrow{\alpha}_{\ell}$ and \rightarrow_{ℓ}^* to denote the reflexive and transitive closure of \rightarrow_{ℓ} .

Example 3.13 *Going back to the handshake protocol described in Example 3.6, the derivation described below represents a normal execution of the protocol. For simplicity of this example we suppose that $x \notin \text{fv}(P)$.*

$$\begin{array}{lcl} & \text{new } w.(P_A(w) \parallel P_B(w)) & \\ \xrightarrow{\text{out}(c, \text{senc}(n, w))}_{\ell} & \text{new } w, n.(P_B(w) \parallel \{\text{senc}(n, w)/x_1\} \parallel \text{in}(c, x). \text{ if } \text{sdec}(x, w) = f(n) \text{ then } P) & \\ \xrightarrow{\text{in}(c, \text{senc}(n, w))}_{\ell} & \text{new } w, n.(\text{out}(c, M) \parallel \{\text{senc}(n, w)/x_1\} \parallel \text{in}(c, x). \text{ if } \text{sdec}(x, w) = f(n) \text{ then } P) & \\ \xrightarrow{\text{out}(c, M\downarrow)}_{\ell} & \text{new } w, n.(\{\text{senc}(n, w)/x_1\} \parallel \{M\downarrow/x_2\} \parallel \text{in}(c, x). \text{ if } \text{sdec}(x, w) = f(n) \text{ then } P) & \\ \xrightarrow{\text{in}(c, \text{senc}(f(n), w))}_{\ell} & \text{new } w, n.(\{\text{senc}(n, w)/x_1\} \parallel \{M\downarrow/x_2\} \parallel \text{if } \text{sdec}(\text{senc}(f(n), w), w) = f(n) \text{ then } P) & \\ \xrightarrow{\tau} & \text{new } w, n.(\{\text{senc}(n, w)/x_1\} \parallel \{M\downarrow/x_2\} \parallel P) & \end{array}$$

where $M = \text{senc}(f(\text{sdec}(\text{senc}(n, w), w)), w) \rightarrow_{\mathcal{R}_{\text{senc}}} \text{senc}(f(n), w)$.

Example 3.14 *Continuing Example 3.7 we develop some transitions from*

$$P_1 = (\text{new } N_a. P_A(a, d)) \parallel (\text{new } N_b. P_B(a, b)) \parallel \text{out}(c, \text{sk}(d))$$

For convenience, the names N_a and N_b are pushed out. We obtain another process that is structurally equivalent.

Case 1: The process P_A may move first, yielding

$$\begin{aligned} P_1 \xrightarrow{\text{out}(c, \text{aenc}(\langle a, N_a \rangle, \text{pk}(d)))}_{\ell} & \text{new } N_a. \text{new } N_b. (\quad \{\text{aenc}(\langle a, N_a \rangle, \text{pk}(d))/x_1\} \\ & \parallel (\text{in}(c, \text{aenc}(\langle N_a, x \rangle, \text{pk}(a))). \text{out}(c, \text{aenc}(x, \text{pk}(b)))) \\ & \parallel P_B(a, b) \\ & \parallel \text{out}(c, \text{sk}(d)) \quad) \end{aligned}$$

Case 2: The process P_B may also move first, and the resulting process depends on an input M_1 such that $\text{new } N_a, N_b.(\sigma \vdash \text{aenc}(\langle a, M_1 \rangle, \text{pk}(b)))$ where $\text{Dom}(\sigma) = \emptyset$.

$$\begin{aligned} P_1 \xrightarrow{\text{in}(c, M_1)}_{\ell} & \text{new } N_a, \text{new } N_b. (\quad P_A(a, d) \\ & \parallel \text{out}(c, \text{aenc}(\langle M_1, N_b \rangle, \text{pk}(a))). \text{in}(c, \text{aenc}(N_b, \text{pk}(b))) \\ & \parallel \text{out}(c, \text{sk}(d)) \quad) \end{aligned}$$

Case 3: The last process may also move first, yielding

$$P_1 \xrightarrow{\text{out}(c, \text{sk}(d))}_{\ell} \text{new } N_a, \text{new } N_b. (\quad \{\text{sk}(d)/x_1\} \parallel P_A(a, d) \parallel P_B(a, b) \quad)$$

From the resulting processes, there are again several possible transitions. We do not continue here the full transition sequence, which is too large to be displayed.

In the above example, we see that the transition system might actually be infinite. Indeed, the term M_1 is an arbitrary message that satisfies some deducibility conditions. Such deducibility conditions can be simplified (and decided). This will be the subject of Chapter 4 on bounded process verification.

3.3 Security Properties

This section presents mainly through examples how to formalise definitions of the most standard security properties. To prove that security properties hold for all adversaries, quantification over all contexts is required. However, in order to consider realistic adversary, we have to consider processes that are built using public function symbols only and we have to ensure that these processes are executable.

In practise, it may be difficult to reason with the quantification over all contexts. The labelled transition semantics aim to eliminate universal quantification of the context and is easier to manipulate. In this section, we rely on this semantics. Since our small process calculus does not allow us to model private channels, we do not have to consider the rule COMM. The attacker controls the entire network and can eavesdrop, block, intercept, and inject messages.

3.3.1 Secrecy

Intuitively, a protocol preserves the secrecy of some message M if an adversary cannot obtain M by constructing it from the outputs of the protocol. We can formalise the adversary as a process running in parallel with the protocol, that after constructing M outputs it on a public channel. The adversary process does not have any of the secrets of the protocol. As explained in introduction of this section, another possibility is to rely on the labelled semantics and to simply ask that in any reachable extended process, M can not be deduced from the frame. Below, you illustrate this property through several examples based on the Needham-Schroeder protocol.

Example 3.15 (Scenario 1) *Consider again the following process defined in Example 3.7 :*

$$P_1 = (\text{new } N_a. P_A(a, d)) \parallel (\text{new } N_b. P_B(a, b)) \parallel \text{out}(c, \text{sk}(d)).$$

We typically wish to ensure the secrecy of the nonce N_b . For this, we have to show that, for any extended process B such that $P_1 \rightarrow_\ell^ B$, we have that $\phi(B) \not\vdash N_b$. Actually, this secrecy property does not hold because of the attack described in Chapter 2. Note that, in this scenario, it is not reasonable to require the secrecy of N_a since N_a is generated by an honest agent for a dishonest one.*

We may also want to express the secrecy of the nonce N_a received by $P_B(a, b)$. This means that we want that the value of y (this is the variable that represents the nonce N_a in the process $P_B(a, b)$) is not known by the attacker in any possible executions. For this, we have to show that for each process B such that $P_1 \rightarrow_\ell^ B$ and in which the variable y has been instantiated by some message M , we have that $\phi(B) \not\vdash M$.*

Example 3.16 (Scenario 2) *Consider now*

$$P_2 = (\text{new } N_a. \text{in}(c, x_b). P_A(a, x_b)) \parallel (\text{new } N_b. \text{in}(c, x_a). P_B(x_a, b)) \parallel \text{out}(c, \text{sk}(d))$$

In such a situation, neither N_a nor N_b can be required to remain secret: this depends on the inputs x_a and x_b . In this case, to express the secrecy of N_b , we can ask that for each process B such that $P_1 \rightarrow_\ell^* B$ and in which the variable x_a has been instantiated by an honest agent (i.e. not d), we have that $\phi(B) \not\models N_b$.

To express secrecy of a nonce in the scope of a replication, we need extra material. Consider the following scenario

$$P_3 = !(\text{new } N_a. \text{in}(c, x_b). P_A(a, x_b)) \parallel !(\text{new } N_b. \text{in}(c, x_a). P_B(x_a, b)) \parallel \text{out}(c, \text{sk}(d)).$$

Intuitively, we wish that, in any copy of the process, in which $x_b \neq d$, then N_a is secret. Be careful that x_b is actually a local variable of the process and should actually be renamed in each copy. Similarly, N_a and N_b are renamed in each instance.

There are again several ways of specifying the desired properties. For instance, we may split the processes in those for which x_b is bound to a honest party and those in which $x_b = d$ and then forget about the different copies in the specification. We may also enrich the calculus with status events. These status events are also very useful to express correspondence properties explained in the following section.

3.3.2 Correspondence Properties

Correspondence properties are used to capture relationships between events that can be expressed in the form "if an event e has been executed then event e has been previously executed." Moreover, these events may contain arguments. This will allow one to express agreement properties. To reason with correspondence properties, we have to annotate processes with events. These events will mark the different control points reached by the protocol.

We say that an extended process A has reached an event $\text{event}(M_1, \dots, M_n)$ if, and only if, there exist an evaluation context C , a plain process P and an extended process B such that $A \equiv C[\text{event}(M_1, \dots, M_n).P \parallel B]$.

Aliveness. This property is the weakest form of authentication in Lowe's hierarchy [?].

A protocol satisfies *aliveness* if, whenever an honest agent completes a run of the protocol, apparently with another honest agent B , then B has previously run the protocol.

Note that B may not necessarily believe that he was running the protocol with A . Also, the agent B may not have run the protocol *recently*. The aliveness of principal B to initiator A can be specified in our formalism. First, we have to consider two status events **start** and **end**. We insert them at the beginning and at the end of each role respectively. For instance, in $P_A(a, d)$, we insert **start**(a) at the beginning and **end**(a, d) at the end. This expresses the fact that the role is executed by a with d . We insert **start**(b) and **end**(b, a) in $P_B(a, b)$. Now, the aliveness property from the point of view of b can be specified as follows:

For any trace execution such that $P_1 \rightarrow_\ell A_1 \rightarrow_\ell \dots \rightarrow_\ell A_n$ such that A_n has reached **end**(M_1, M_2) with $M_1 \neq d$ and $M_2 \neq d$, there exists i such that A_i has reached **start**(M_2). This corresponds to the fact that the property "if **end**(x, y) has been executed then **start**(y) has been previously executed when x and y are both honest agents." For the Needham-Schroeder public key protocol (e.g. Scenario 1) the aliveness property is satisfied.

Weak agreement. Weak agreement is slightly stronger than aliveness.

A protocol guarantees *weak agreement* if, whenever an honest agent completes a run of the protocol, apparently with another honest agent B , then B has previously been running the protocol, apparently with A .

The weak agreement property can also be expressed in our formalism. We have again to add status events **start** and **end** in our specification. However, the predicate **start** will have also two parameters: **start**(a, d) expresses the fact that a has started a session with d . Now, the weak agreement property can be specified as follows:

For any trace execution such that $P_1 \rightarrow_\ell A_1 \rightarrow_\ell \dots \rightarrow_\ell A_n$ such that A_n has reached $\text{end}(M_1, M_2)$ with $M_1 \neq d$ and $M_2 \neq d$, there exists i such that A_i has reached **start**(M_2, M_1). For the Needham-Schroeder public key protocol, it is well-known that this property is not satisfied: b can complete a session apparently with a whereas a has never started a session with b . The property is already falsified on Scenario 1.

We can also express some refinements of these properties by distinguishing the case where an agent starts a session as an initiator or as a responder. Moreover, we can also express the fact that the two agents agreed on some message M , *e.g.* the value of a nonce or a key. This allows us to express the non-injective agreement security property. There are also stronger agreement properties, that would require the mapping from **end** to **start** to be injective.

3.3.3 Guessing Attacks

Guessing attacks are a kind of dictionary attack in which the password is supposed to be weak, *i.e.* part of a dictionary for which a brute force attack is feasible. A guessing attack works in two phases. In a first phase the attacker eavesdrops and interacts with one or several protocol sessions. In a second *offline* phase, the attacker tries each of the possible passwords on the data collected during the first phase. To resist against a guessing attack, the protocol must be designed such that the attacker cannot discover on the basis of the data collected whether his current guess of the password is the actual password or not.

The idea behind the definition is the following. Suppose the frame ϕ represents the information gained by the attacker by eavesdropping one or more sessions and let w be the weak password. Then, we can represent resistance against guessing attacks by checking whether the attacker can distinguish a situation in which he guesses the correct password w and a situation in which he guesses an incorrect one, say w' . We model these two situations by adding $\{w/x\}$ (resp. $\{w'/x\}$) to the frame. We use static equivalence to capture the notion of indistinguishability. This definition is due to M. Baudet [?], inspired from the one of [?]. In our definition, we allow multiple shared secrets, and write \bar{w} for a sequence of such secrets.

Definition 3.3 Let $\phi \equiv \text{new } \bar{w}.\phi'$ be a frame. We say that the frame ϕ is resistant to guessing attacks against \bar{w} if

$$\text{new } \bar{w}.\phi' \parallel \{\bar{w}/\bar{x}\} \sim \text{new } \bar{w}'.\text{new } \bar{w}.\phi' \parallel \{\bar{w}'/\bar{x}\}$$

where \bar{w}' is a sequence of fresh names and \bar{x} a sequence of variables such that $\bar{x} \cap \text{Dom}(\phi) = \emptyset$.

Note that this definition is general w.r.t. to the equational theory and the number of guessable data items. Now, we can define what it means for a protocol to be resistant against guessing attacks.

Definition 3.4 *Let A be a process and $\bar{w} \subseteq \text{bn}(A)$. We say that A is resistant to guessing attacks against \bar{w} if, for every process B such that $A \rightarrow_\ell^* B$, we have that the frame $\phi(B)$ is resistant to guessing attacks against \bar{w} .*

Example 3.17 *Consider the handshake protocol described in Example 3.6. An interesting problem arises if the shared key w is a weak secret, i.e. vulnerable to brute-force off-line testing. In such a case, the protocol has a guessing attack against w . Indeed, we have that*

$$\text{new } w.(P_A(w) \parallel P_B(w)) \rightarrow_\ell^* D$$

with $\phi(D) = \text{new } w.\text{new } n.(\{\text{senc}(n,w)/x_1\} \parallel \{\text{senc}(f(n),w)/x_2\})$. The frame $\phi(D)$ is not resistant to guessing attacks against w . The test $f(\text{sdec}(x_1, x)) \stackrel{?}{=} \text{sdec}(x_2, x)$ allows us to distinguish the two associated frames:

- $\text{new } w.\text{new } n.(\{\text{senc}(n,w)/x_1\} \parallel \{\text{senc}(f(n),w)/x_2\} \parallel \{w/x\})$, and
- $\text{new } w'.\text{new } w.\text{new } n.(\{\text{senc}(n,w)/x_1\} \parallel \{\text{senc}(f(n),w)/x_2\} \parallel \{w'/x\})$.

Hence, the process $\text{new } w.(P_A \parallel P_B)$ is not resistant to guessing attacks against w .

3.3.4 Equivalence Properties

The notion of indistinguishability is a powerful concept which allows us to reason about complex properties that cannot be expressed as secrecy or correspondence properties. Intuitively, two processes are said to be equivalent if an observer has no way to tell them apart. While static equivalence models indistinguishability of sequences of terms, it is also possible to lift it to an observational equivalence, i.e. indistinguishability of processes that interact with an arbitrary adversary. We define this observational equivalence by the means of a labelled bisimulation. The processes may perform different computations, but they have to look the same to an external observer. This notion allows us to define strong notions of secrecy and also privacy properties.

Before we formalise this notion of equivalence, we have to adapt the labelled semantics provided in Section 3.2.2. Indeed, we will now assume that the attacker can observe the interactions with the environment and we have to capture the fact that the attacker performs the same experiment on both processes. Intuitively, we want that, for any experiment s (sequence of labels) such that $A \xrightarrow_\ell^s A'$, there exists B' such that $B \xrightarrow_\ell^s B'$ and $\phi(A') \sim \phi(B')$. However, our labels are too fine grained.

Let $A = \text{new } n.\text{out}(c, n)$ and $B = \text{new } n, k.\text{out}(c, \text{senc}(n, k))$. The only transitions that can be performed by A and B are as follows:

- $A \xrightarrow_\ell \text{out}(c, n) \text{new } n.\{n/x\}$, and
- $B \xrightarrow_\ell \text{out}(c, \text{senc}(n, k)) \text{new } n, k.\{\text{senc}(n, k)/x\}$.

However, in reality an attacker has no way to distinguish these two processes since he will not see any difference between a fresh nonce and an encryption (he does not know the key). The same situation also occurs with the two processes $A = \text{new } n.\text{in}(c, y)$ and $B = \text{new } n'.\text{in}(c, y)$. We have that $A \xrightarrow{\text{in}(c, n')} 0$ and B can not mimic this step. B is not allowed to use the name n' since it is restricted. Our labels contains too much information. We modify the IN and OUT rules as follows:

$$\begin{array}{ll} \text{IN} & \text{in}(c, x).P \xrightarrow{\text{in}(c, M)}_{\ell} P\{M/x\} \quad \text{where } M \text{ is a message} \\ \text{OUT} & \text{out}(c, M).P \xrightarrow{\text{out}(c, x)}_{\ell} P \parallel \{M\downarrow/x\} \quad \text{where } x \text{ is a fresh variable and } M\downarrow \text{ is a message} \end{array}$$

The labelled operational semantics is closed by structural equivalence and under some evaluation contexts. Actually, we have that:

$$\frac{A \equiv A' \quad A' \xrightarrow{\alpha}_{\ell} B' \quad B' \equiv B}{A \xrightarrow{\alpha}_{\ell} B} \quad \frac{A \xrightarrow{\alpha}_{\ell} B}{C[A] \xrightarrow{\alpha'}_{\ell} C[B]}$$

where C is an evaluation context, and in case of an input, *i.e.* $\alpha = \text{in}(c, M)$, we have that $\phi(C[A]) \vdash_{\mathcal{R}} M$ and $\alpha' = \text{in}(c, M')$ where M' is a recipe witnessing the fact that $\phi(C[A]) \vdash_{\mathcal{R}} M$.

Moreover, we now consider that structural equivalence is closed under α -renaming.

Example 3.18 *Going back to the handshake protocol described in Example 3.6, the derivation described below represents a normal execution of the protocol in the new labelled semantics.*

$$\begin{array}{l} \text{new } w.(P_A(w) \parallel P_B(w)) \\ \xrightarrow{\text{out}(c, x_1)}_{\ell} \text{new } w, n.(P_B(w) \parallel \{\text{senc}(n, w)/x_1\} \parallel \text{in}(c, x). \text{ if sdec}(x, w) = f(n) \text{ then } P) \\ \xrightarrow{\text{in}(c, x_1)}_{\ell} \text{new } w, n.(\text{out}(c, M) \parallel \{\text{senc}(n, w)/x_1\} \parallel \text{in}(c, x). \text{ if sdec}(x, w) = f(n) \text{ then } P) \\ \xrightarrow{\text{out}(c, x_2)}_{\ell} \text{new } w, n.(\{\text{senc}(n, w)/x_1\} \parallel \{M\downarrow/x_2\} \parallel \text{in}(c, x). \text{ if sdec}(x, w) = f(n) \text{ then } P) \\ \xrightarrow{\text{in}(c, x_2)}_{\ell} \text{new } w, n.(\{\text{senc}(n, w)/x_1\} \parallel \{M\downarrow/x_2\} \parallel \text{if sdec}(\text{senc}(f(n), w), w) = f(n) \text{ then } P) \\ \xrightarrow{\tau} \text{new } w, n.(\{\text{senc}(n, w)/x_1\} \parallel \{M\downarrow/x_2\} \parallel P) \end{array}$$

where $M = \text{senc}(f(\text{sdec}(\text{senc}(n, w), w)), w) \rightarrow_{\mathcal{R}_{\text{senc}}} \text{senc}(f(n), w)$.

For every closed extended process A we define its set of traces, each trace consisting in a sequence of visible actions (*i.e.* different from τ) together with the sequence of sent messages:

$$\text{trace}(A) = \{(s, \phi(B)) \mid A \xrightarrow{s}_{\ell} B \text{ for some } B\}.$$

Note that, in the new versions of our labelled semantics, the sent messages are exclusively stored in the frame and not in the sequence s (the outputs are made by “reference”).

Definition 3.5 (trace equivalence \approx_t) *Let A and B be two closed extended processes, $A \sqsubseteq_t B$ if for every $(s, \varphi) \in \text{trace}(A)$, there exists $(s', \varphi') \in \text{trace}(B)$ such that $s = s'$ and $\varphi \sim \varphi'$. The extended processes A and B are trace equivalent, denoted by $A \approx_t B$, if $A \sqsubseteq_t B$ and $B \sqsubseteq_t A$.*

Example 3.19 Consider the equational theory described in Example 3.1. We have that:

$$\begin{aligned}
& \text{new } s, k. \text{out}(c, \text{senc}(s, k)).\text{in}(c, x). \text{ if } x = s \text{ then } \text{out}(c, a) \approx_t \text{new } s, k. \text{out}(c, \text{senc}(s, k)).\text{in}(c, x) \\
& \text{new } s. \text{out}(c, \text{senc}(s, k)).\text{in}(c, x). \text{ if } x = s \text{ then } \text{out}(c, a) \not\approx_t \text{new } s. \text{out}(c, \text{senc}(s, k)).\text{in}(c, x) \\
& \text{out}(c, a).\text{out}(c, a) \approx_t \text{out}(c, a) \parallel \text{out}(c, a) \\
& \text{out}(c_1, a).\text{out}(c_2, a) \not\approx_t \text{out}(c_1, a) \parallel \text{out}(c_2, a) \\
& \text{out}(c, a); \text{out}(c, b) \not\approx_t \text{out}(c, a) \parallel \text{out}(c, b)
\end{aligned}$$

Now, we develop an example to illustrate how this notion of equivalence can be used to formalise anonymity.

Example 3.20 We consider a slightly simplified version of a protocol given in [?] designed for transmitting a secret without revealing its identity to other participants. In this protocol, A is willing to engage in communication with B and wants to reveal its identity to B . However, A does not want to compromise its privacy by revealing its identity or the identity of B more broadly. The participants A and B proceed as follows:

$$\begin{aligned}
A \rightarrow B & : \text{aenc}(\langle N_a, \text{pk}(A) \rangle, \text{pk}(B)) \\
B \rightarrow A & : \text{aenc}(\langle N_a, \langle N_b, \text{pk}(B) \rangle \rangle, \text{pk}(A))
\end{aligned}$$

First A sends to B a nonce N_a and its public key encrypted with the public key of B . If the message is of the expected form then B sends to A the nonce N_a , a freshly generated nonce N_b and its public key, all of this being encrypted with the public key of A . Otherwise, B sends out a “decoy” message: $\text{aenc}(N_b, \text{pk}(B))$. This message should basically look like B ’s other message from the point of view of an outsider. This is important since the protocol is supposed to protect the identity of the participants.

A session of role A played by agent a with b can be modelled by the following basic process where $M = \text{dec}(x, \text{sk}(a))$.

$$\begin{aligned}
A(a, b) & \triangleq \\
& \text{out}(c, \text{aenc}(\langle N_a, \text{pk}(a) \rangle, \text{pk}(b))). \\
& \text{in}(c, x). \\
& \text{if } \langle \text{proj}_1(M), \text{proj}_2(\text{proj}_2(M)) \rangle = \langle N_a, \text{pk}(b) \rangle \text{ then } 0
\end{aligned}$$

Similarly, a session of role B played by agent b with a can be modelled by the basic process $B(b, a)$ where $N = \text{dec}(y, \text{sk}(b))$.

$$\begin{aligned}
B(b, a) & \triangleq \text{in}(c, y). \\
& \text{if } \text{proj}_2(N) = \text{pk}(a) \text{ then } \text{out}(c, \text{aenc}(\langle \text{proj}_1(N), \langle N_b, \text{pk}(b) \rangle \rangle, \text{pk}(a))) \\
& \text{else } \text{out}(c, \text{aenc}(N_b, \text{pk}(b))).
\end{aligned}$$

Intuitively, this protocol preserves anonymity if an attacker cannot distinguish whether b is willing to talk to a (represented by the process $B(b, a)$) or willing to talk to a' (represented by the process $B(b, a')$), provided a , a' and b are honest participants. For illustration purposes, we also consider the process $B'(b, a)$ obtained from $B(b, a)$ by replacing the else branch by else 0. We will see that the “decoy” message plays a crucial role to ensure privacy.

We can ask whether the two following processes P_{ex} and P'_{ex} are in equivalence:

- $P_{\text{ex}} = \text{new } N_a.\text{new } N_b.[A(a, b) \parallel B(b, a) \parallel K(a, a', b)]$, and
- $P'_{\text{ex}} = \text{new } N_a.\text{new } N_b.[A(a', b) \parallel B(b, a') \parallel K(a, a', b)]$.

where $K(a, a', b) = \text{out}(c, \text{pk}(a)).\text{out}(c, \text{pk}(a')).\text{out}(c, \text{pk}(b))$.

Actually, the 'decoy' message is crucial to have this equivalence, and thus anonymity. We have that $P_{\text{ex}} \approx_t P'_{\text{ex}}$ whereas:

$$\text{new } N_a, N_b.[A(a, b) \parallel B'(b, a) \parallel K(a, a', b)] \not\approx_t \text{new } N_a, N_b.[A(a', b) \parallel B'(b, a') \parallel K(a, a', b)].$$

Another notion of equivalence that has been quite well-studied is the notion of *observationally equivalent*. However, proofs of observational equivalences are difficult because of the universal quantification over all contexts. In the context of the applied pi calculus, it has been shown that observational equivalence coincides with labelled bisimilarity [?]. This should also hold in the calculus presented here.

Definition 3.6 (labelled bisimilarity \approx_ℓ) Labeled bisimilarity is the largest symmetric relation \mathcal{R} on closed extended processes such that $A \mathcal{R} B$ implies

1. $\phi(A) \sim \phi(B)$,
2. if $A \xrightarrow{\tau} A'$, then $B \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' ,
3. if $A \xrightarrow{\alpha}_\ell A'$ then $B \rightarrow^* \xrightarrow{\alpha}_\ell \rightarrow^* B'$ and $A' \mathcal{R} B'$ for some B' .

It is easy to see that observational equivalence (or labelled bisimilarity) implies trace equivalence while the converse is false in general (see Example 3.21).

Lemma 3.1 Let A and B be two extended processes: $A \approx_\ell B$ implies $A \approx_t B$.

Example 3.21 For convenience we introduce for this example a non-deterministic choice operator $+$ and extend the internal reduction by the rule $A+B \rightarrow A$ and structural equivalence by associativity and commutativity of $+$. Consider the two following processes:

$$A = (\text{out}(c, a).\text{out}(c, b_1)) + (\text{out}(c, a).\text{out}(c, b_2))$$

$$B = \text{out}(c, a).(\text{out}(c, b_1) + \text{out}(c, b_2))$$

We have that $A \approx_t B$ whereas $A \not\approx_\ell B$. Intuitively, after B 's first move, B still has the choice of emitting b_1 or b_2 , while A , trying to follow B 's first move, is forced to choose between two states from which she can only emit one of the two.

The notion of labelled bisimilarity is quite strong but is also used to express privacy type properties. It is more or less a matter of taste to define anonymity w.r.t. trace equivalence or w.r.t. the stronger version with labelled bisimilarity.

3.4 Further Readings

The calculus presented in this chapter is very close to the applied pi calculus [?]. A presentation of this calculus in a tutorial style is also available [?]. Another calculus that is very close to the applied pi calculus and for which there exists a tutorial presentation is the spi-calculus [?].

3.5 Exercises

Exercise 3.1 (★) Consider the signature described in Example 3.1. Let

$$\phi = \text{new } s, k_1. \{ \text{senc}(s, \langle k_1, k_2 \rangle) /_{x_1}, \text{senc}(k_1, k_2) /_{x_2} \}.$$

1. Is s deducible from ϕ ?
2. Could you enumerate the subterms of ϕ ? Among these subterms, give those that are deducible.
3. Give a term that is deducible from ϕ and that is not a subterm.

Exercise 3.2 (★) Give a reasonable formalisation of the following protocol:

$$\begin{aligned} A \rightarrow B : & \langle A, \{K\}_{\text{pk}(B)} \rangle \\ B \rightarrow A : & \langle B, \{K\}_{\text{pk}(A)} \rangle \end{aligned}$$

First, A generates a fresh key K and sends it encrypted with the public key of B . Only B will be able to decrypt this message. In this way, B learns K and B also knows that this message comes from A as indicated in the first part of the message he received. Hence, B answers to A by sending again the key, this time encrypted with the public key of A .

Exercise 3.3 (★) Consider the formalisation of the Needham-Schroeder protocol as described in Example 3.7, and the following scenario (see Example 3.15).

$$(\text{new } N_a. P_A(a, d)) \parallel (\text{new } N_b. P_B(a, b)) \parallel \text{out}(c, \text{sk}(d)).$$

Give the complete transition sequence that yields the attack on the secrecy of the nonce N_b .

Exercise 3.4 (★★) Give a reasonable formalisation of the handshake protocol without using the conditional (`if then else`). Give a trace that exists in the model presented in 3.6 and that does not exist in this new formalisation.

Part II

Verification in the Symbolic Setting

Chapter 4

Deducibility Constraints

In this chapter, we present the NP-complete decision procedure for a bounded number of sessions by H. Comon-Lundh *et al.* [?]. In this setting (*i.e.* finite number of sessions), deducibility constraint systems have become the standard model for verifying security properties, with a special focus on secrecy. Starting with a paper by J. Millen and V. Shmatikov [?], many results (*e.g.* [?, ?]) have been obtained within this framework.

Here, we consider only symmetric/asymmetric encryptions, and pairing. We show that any deducibility constraint system can be transformed in (possibly several) much simpler deducibility constraint systems that are called solved forms, preserving all solutions of the original system, and not only its satisfiability. In other words, the deducibility constraint system represents in a symbolic way all the possible sequences of messages that are produced, following the protocol rules, whatever are the intruder's actions. This set of symbolic traces is infinite in general. Solved forms are a simple (and finite) representation of such traces. The procedure preserves all solutions. Hence, we can represent for instance, all attacks on the secrecy and not only decide if there exists one. Moreover, presenting the decision procedure using a small set of simplification rules yields more flexibility for further extensions and modifications.

4.1 Intruder Deduction problem

4.1.1 Preliminaries

An *inference rule* is a rule of the form
$$\frac{u_1 \dots u_n}{u_0}$$
 where u_0, u_1, \dots, u_n are terms (with variables). An *inference system* is a set of inference rules.

Example 4.1 *The following inference system \mathcal{I}_{DY} represents the deduction capabilities of an attacker. We consider the signature $\mathcal{F} = \{\text{senc}, \text{aenc}, \langle -, - \rangle, \text{sk}\}$ and the underlying rewriting system \mathcal{R} is empty. There are several possible ways of defining the intruder capabilities, we choose here the “implicit destructors” formulation, in which the destructors do not appear. This leads to an inference system that is slightly different from the one proposed in Section 3.1.3. For sake of simplicity, we make a confusion between the identity of an agent*

and his public key.

$$\begin{array}{c}
\frac{x \quad y}{\langle x, y \rangle} \text{ P} \qquad \frac{x \quad y}{\text{aenc}(x, y)} \text{ PKE} \qquad \frac{x \quad y}{\text{senc}(x, y)} \text{ SE} \\
\\
\frac{\langle x, y \rangle}{x} \text{ Left} \quad \frac{\langle x, y \rangle}{y} \text{ Right} \quad \frac{\text{aenc}(x, y) \quad \text{sk}(y)}{x} \text{ PKD} \quad \frac{\text{senc}(x, y) \quad y}{x} \text{ SD}
\end{array}$$

The rules P, SE, and PKE are composition rules whereas the rules Left, Right, SD, and PKD are decomposition rules.

Definition 4.1 (proof) Let \mathcal{I} be an inference system. A proof Π of $T \vdash u$ in \mathcal{I} is a tree such that:

- every leaf of Π is labelled with a term v such that $v \in T$,
- for every node labelled with v_0 having n sons labelled with v_1, \dots, v_n , there is an instance of an inference rule with conclusion v_0 and hypotheses v_1, \dots, v_n . We say that Π ends with this instance if the node is the root of Π ,
- the root is labelled with u .

We denote by $\text{Hyp}(\Pi)$ the set of labels of the leaves of a proof Π and $\text{Conc}(\Pi)$ is the label of the root of Π . $\text{Steps}(\Pi)$ is the set of labels of all nodes of Π . The *size* of a proof Π is the number of nodes in it. A proof Π of $T \vdash u$ is *minimal* if it does not exist any proof Π' of $T \vdash u$ having a size strictly smaller than the size of Π .

Example 4.2 Let $\phi = \text{new } a, b, s. \{ \langle \text{senc}(s, \langle a, b \rangle), a \rangle /_{x_1}, \text{senc}(b, a) /_{x_2} \}$. We may ask whether s is deducible from ϕ , i.e. does there exist a proof of $\langle \text{senc}(s, \langle a, b \rangle), a \rangle, \text{senc}(b, a) \vdash s$. Such a proof is given below:

$$\frac{\frac{\langle \text{senc}(s, \langle a, b \rangle), a \rangle}{\text{senc}(s, \langle a, b \rangle)} \quad \frac{\frac{\langle \text{senc}(s, \langle a, b \rangle), a \rangle \quad \text{senc}(b, a)}{a} \quad \frac{\text{senc}(b, a) \quad a}{b}}{\langle a, b \rangle}}{s}$$

The problem whether an intruder can gain certain information s from a set of knowledge T , i.e. whether there is a proof of $T \vdash s$ is called the *intruder deduction problem*.

Intruder deduction problem (for a fixed inference system \mathcal{I})

INPUT: a finite set of terms T , and a term s (the secret).

OUTPUT: Does there exist a proof of $T \vdash s$?

This definition is in-line with the concept of deduction introduced in Section 3.1.3. Here, we do not explicitly rely on the concept of frame. Note that for deduction, the ordering in which the messages have been sent is not relevant. Moreover, restriction on names are not necessary. It is assumed that each name is restricted.

4.1.2 Decidability via Locality

To show that the intruder deduction problem is decidable (in PTIME) for an inference system \mathcal{I} , we use the notion of *locality* introduced by D. McAllester [?].

Definition 4.2 (locality) *Let \mathcal{I} be an inference system. The system \mathcal{I} is local if whenever $T \vdash u$ in \mathcal{I} , there exists a proof Π of $T \vdash u$ such that $\text{Steps}(\Pi) \subseteq st(T \cup \{u\})$.*

Given an inference system \mathcal{I} , to establish that the intruder deduction problem is decidable, it is actually sufficient to prove that:

1. a *locality result* for the inference system \mathcal{I} : checking the existence of a proof of $T \vdash u$ amounts to checking the existence of a local proof that only contains subterms of u and T (there is a polynomial number of subterms),
2. a *one-step-deducibility result* to ensure that we can test (in PTIME) whether a term is deducible in one step from a set of terms by using an instance of one of the inference rules. This result trivially holds for the inference system presented in Example 4.1.

Then, the existence of a local proof of $T \vdash u$ can be checked in polynomial time by saturation of T with terms deducible in one-step. Thanks to locality, the number of iteration to obtain a saturated set is bounded by the number of terms that can be involved in a local proof. This yields a PTIME algorithm.

Lemma 4.1 (locality) *Let T be a set of terms and u be a term. A minimal proof Π of $T \vdash u$ only contains terms in $st(T \cup \{u\})$, i.e. $\text{Steps}(\Pi) \subseteq st(T \cup \{u\})$. Moreover, if Π is reduced to a leaf or ends with a decomposition rule, then we have that $\text{Steps}(\Pi) \subseteq st(T)$.*

Proof : Let Π be a minimal proof of $T \vdash u$. We prove the result by induction on the size of the proof Π .

Base case: In such a case, the proof Π is reduced to a leaf and we easily conclude.

Induction step: We have that:

$$\Pi = \left\{ \begin{array}{ccc} \Pi_1 & & \Pi_n \\ u_1 & \cdots & u_n \\ \hline & u & \end{array} \right. R$$

We distinguish several cases depending on the last inference rule of Π .

- If R is a composition rule, then u_1, \dots, u_n are subterms of u and we easily conclude by relying on our induction hypothesis.
- If R is a projection rule (say proj_1), then $u_1 = \langle u, v \rangle$ for some v . In such a case, by minimality of Π , we know that Π_1 does not end with a composition rule. Hence, by relying on our induction hypothesis, we have that $\text{Steps}(\Pi_1) \subseteq st(T)$, and thus $u_1 \in st(T)$. Moreover, we have that $u \in st(u_1)$, and thus $u \in st(T)$. This allows us to conclude that $\text{Steps}(\Pi) \subseteq st(T)$.

The cases where Π ends with a decryption rule (symmetric and asymmetric) can be done in a similar way. \square

Proposition 4.1 *The intruder deduction problem is decidable in PTIME for \mathcal{I}_{DY} . Actually, this problem is PTIME complete.*

The PTIME-hardness can be proved by a reduction from HORNSAT.

The concept of locality has been used to establish decidability of several inference systems. For instance, we may want to model digital signature, exclusive or operator, commutative encryption, ...

4.2 Deducibility constraints

Assume processes without replication. Then the transition system is finite in depth but might be infinitely branching, as we saw in Example 3.14. The idea then is to represent in a simple symbolic way the set of terms that satisfy the required conditions. This is what we formalise now.

Definition 4.3 *A Deducibility constraint system is either \perp or a conjunction of deducibility constraints of the form:*

$$T_1 \overset{?}{\vdash} u_1 \wedge \dots \wedge T_n \overset{?}{\vdash} u_n$$

in which T_1, \dots, T_n are finite sets of terms, u_1, \dots, u_n are terms. Moreover, we assume that the constraints can be ordered in such a way that:

- *monotonicity:* $\emptyset \subsetneq T_1 \subseteq T_2 \cdots \subseteq T_n$
- *origination:* for every i , we have that $fv(T_i) \subseteq fv(u_1, \dots, u_{i-1})$

Intuitively, the sets T_i correspond to messages that have been sent on the network, while u_1, \dots, u_n are the messages that are expected by the processes, hence have to be constructed by the environment. The first condition, called *monotonicity* reflects the fact that the set of messages that have been sent on the network can only increase. In other words, the ordering on the atomic deducibility constraints is a temporal ordering of actions. The second condition (called *origination*) reflects the properties of our processes: variables that occur in a message sent on the network must appear before in messages received from the network.

Definition 4.4 (T_x) *Let $\mathcal{C} = T_1 \overset{?}{\vdash} u_1 \wedge \dots \wedge T_n \overset{?}{\vdash} u_n$ be a deducibility constraint system and x be a variable that occurs in \mathcal{C} . T_x is the minimal set (w.r.t. inclusion) among the sets T_1, \dots, T_n such that $T \overset{?}{\vdash} u \in \mathcal{C}$ and $x \in fv(u)$.*

Thanks to the monotonicity and the origination properties, for any $x \in fv(\mathcal{C})$, the set T_x exists and is uniquely defined.

Such constraint systems may be enriched with equations/disequations between terms or other constraints, that correspond to the conditions in the process calculus. We consider (for now) only these simple constraints.

Definition 4.5 (solution) *Let \mathcal{I} be an inference system. A substitution σ is a solution of a deducibility constraint system $\mathcal{C} = T_1 \overset{?}{\vdash} u_1 \wedge \dots \wedge T_n \overset{?}{\vdash} u_n$ if there exists a proof of $T_i \sigma \vdash u_i \sigma$ in \mathcal{I} for every $i \in \{1, \dots, n\}$.*

Example 4.3 Consider the constraints corresponding to one of the possible Needham-Schroeder symbolic trace. We give explicitly the free names to the attacker and assume that all names that are not explicitly given are (supposedly) secret:

$$\mathcal{C} \triangleq \left\{ \begin{array}{l} a, b, d, \text{sk}(d), \text{aenc}(\langle a, N_a \rangle, d) \stackrel{?}{\vdash} \text{aenc}(\langle a, x \rangle, b) \\ a, b, d, \text{sk}(d), \text{aenc}(\langle a, N_a \rangle, d), \text{aenc}(\langle x, N_b \rangle, a) \stackrel{?}{\vdash} \text{aenc}(\langle N_a, y \rangle, a) \end{array} \right.$$

The failure of the secrecy of N_b (for this scenario) is given by the additional constraint:

$$a, b, d, \text{sk}(d), \text{aenc}(\langle a, N_a \rangle, d), \text{aenc}(\langle x, N_b \rangle, a), \text{aenc}(y, d) \stackrel{?}{\vdash} N_b$$

A solution of \mathcal{C} in \mathcal{I}_{DY} is $\sigma = \{x \mapsto N_a, y \mapsto N_b\}$.

4.3 Decision Procedure

We describe here a non-deterministic simplification procedure. It can be simplified in many respects, but we will see that the problem of deciding whether a constraint system has at least one solution is NP-complete anyway (for the \mathcal{I}_{DY} inference system given in Example 4.1). Many parts of this section, including the set of simplification rules, are borrowed from [?].

4.3.1 Simplification Rules

We prove that any deducibility constraint system can be transformed into simpler ones, called *solved*. Such simplified constraints are then used to decide the security properties.

$$\begin{array}{ll} \text{R}_1 & \mathcal{C} \wedge T \stackrel{?}{\vdash} u \rightsquigarrow \mathcal{C} \quad \text{if } T \cup \{x \mid (T' \stackrel{?}{\vdash} x) \in \mathcal{C}, T' \subsetneq T\} \vdash u \\ \text{R}_2 & \mathcal{C} \wedge T \stackrel{?}{\vdash} u \rightsquigarrow_{\sigma} \mathcal{C}\sigma \wedge T\sigma \stackrel{?}{\vdash} u\sigma \quad \text{if } t \in \text{st}(T), \sigma = \text{mgu}(t, u), t \neq u \\ & \quad \quad \quad t, u \text{ not variables} \\ \text{R}_3 & \mathcal{C} \wedge T \stackrel{?}{\vdash} u \rightsquigarrow_{\sigma} \mathcal{C}\sigma \wedge T\sigma \stackrel{?}{\vdash} u\sigma \quad \text{if } t_1, t_2 \in \text{st}(T), \sigma = \text{mgu}(t_1, t_2), \text{ and } t_1 \neq t_2 \\ \text{R}_4 & \mathcal{C} \wedge T \stackrel{?}{\vdash} u \rightsquigarrow \perp \quad \text{if } \text{fv}(T \cup \{u\}) = \emptyset \text{ and } T \not\vdash u \\ \text{R}_f & \mathcal{C} \wedge T \stackrel{?}{\vdash} f(u, v) \rightsquigarrow \mathcal{C} \wedge T \stackrel{?}{\vdash} u \wedge T \stackrel{?}{\vdash} v \quad \text{for } f \in \{\langle, \rangle, \text{senc}, \text{aenc}\} \end{array}$$

Figure 4.1: Simplification rules.

All the rules are indexed by a substitution (when there is no index then the identity substitution is assumed). We write $\mathcal{C} \rightsquigarrow_{\sigma}^* \mathcal{C}'$ if there are constraint systems $\mathcal{C}_1, \dots, \mathcal{C}_n$ such that $\mathcal{C} \rightsquigarrow_{\sigma_0} \mathcal{C}_1 \rightsquigarrow_{\sigma_1} \dots \rightsquigarrow_{\sigma_n} \mathcal{C}'$ and $\sigma = \sigma_0 \sigma_1 \dots \sigma_n$. We denote by $\sigma = \text{mgu}(u, v)$ a most general unifier of u and v , such that $\text{fv}(v\sigma, u\sigma) \subseteq \text{fv}(v, u)$.

A constraint system is called *solved* if it is different from \perp and if each of its constraints is of the form $T \stackrel{?}{\vdash} x$, where x is a variable. Note that the empty constraint system is solved. Solved constraint systems are particularly simple since they always have a solution. Indeed, let T_1 be the smallest (w.r.t. inclusion) left-hand side of a constraint. From the definition

of a constraint system we have that $T_1 \neq \emptyset$ and has no variable. Then the substitution τ defined by $x\tau = t_1$ where $t_1 \in T_1$ for every variable x , is a solution since $T \vdash x\theta$ for any constraint $T \vdash x$ of the solved constraint system.

Given a constraint system \mathcal{C} , we say that T_i is a *minimal unsolved left-hand side* of \mathcal{C} if T_i is a left-hand side of \mathcal{C} and for all $T \vdash u \in \mathcal{C}$ such that $T \subsetneq T_i$, we have that u is a variable.

Lemma 4.2 *The simplification rules transform a deducibility constraint system into a deducibility constraint system.*

Theorem 4.1 *Let \mathcal{C} be an unsolved constraint system.*

1. (Termination) *There is no infinite chain $\mathcal{C} \rightsquigarrow_{\sigma_1} \mathcal{C}_1 \dots \rightsquigarrow_{\sigma_n} \mathcal{C}_n$.*
2. (Correctness) *If $\mathcal{C} \rightsquigarrow_{\sigma}^* \mathcal{C}'$ for some constraint system \mathcal{C}' and some substitution σ and if θ is a solution of \mathcal{C}' then $\sigma\theta$ is a solution of \mathcal{C} .*
3. (Completeness) *If θ is a solution of \mathcal{C} , then there exist a solved constraint system \mathcal{C}' and substitutions σ, θ' such that $\theta = \sigma\theta'$, $\mathcal{C} \rightsquigarrow_{\sigma}^* \mathcal{C}'$ and θ' is a solution of \mathcal{C}' .*

Termination and correctness are quite easy to show. For termination, it is easy to see that the number of variables is non-increasing. Furthermore, this number strictly decreases by the rules R_2 and R_3 . Any other rule strictly reduces the total size of the right hand sides of the constraint (here, the “size” is the number of symbols in the term). Completeness is more involved and its proof is detailed in Section 4.3.2. Getting a polynomial bound on the length of simplification sequences requires to consider a particular strategy.

4.3.2 Completeness

First, we show that proofs considered in solutions of constraints can be narrowed to so-called *simple proofs*. Let $T_1 \subseteq T_2 \subseteq \dots \subseteq T_n$. We say that a proof Π of $T_i \vdash u$ is *left minimal* if, whenever there is a proof of $T_j \vdash u$ for some $j < i$, then Π is also a proof of $T_j \vdash u$. In other words, the left-minimal proofs are those that can be performed in a minimal T_j . We say that a proof is *simple* if all its subproofs are left minimal and there is no repeated label on any branch. Note that a subproof of a simple proof is simple.

Lemma 4.3 *Let $T_1 \subseteq T_2 \subseteq \dots \subseteq T_n$ be a sequence of sets of terms and u be a term such that $T_i \vdash u$. There exists a simple proof Π of $T_i \vdash u$.*

Proof : Let i be a minimal index for which there is a proof of $T_i \vdash u$. Thanks to Lemma 4.1, there is a local proof Π_0 of $T_i \vdash u$. We prove the lemma by induction on the size of Π_0 .

Base case: Π_0 is reduced to a leaf. In such a case, Π_0 is a simple proof.

Induction step: Consider the last rule in the proof of u :

$$\Pi_0 = \left\{ \begin{array}{c} \Pi_1 \quad \dots \quad \Pi_n \\ \dots \\ \frac{u_1 \quad \dots \quad u_n}{u} R \end{array} \right.$$

For every $j = 1, \dots, n$, we have that Π_j is a proof of $T_i \vdash u_j$. By induction hypothesis, there are simple proofs Π'_j of u_j . If u appears as a node in some of these proofs, let Π be the corresponding subproof and we get the desired result. Otherwise, let

$$\Pi = \left\{ \frac{\begin{array}{ccc} \Pi'_1 & \dots & \Pi'_n \\ u_1 & & u_n \end{array}}{u} R \right.$$

The proof Π is a simple proof of u . □

Lemma 4.4 *Let \mathcal{C} be an unsolved constraint system, θ be a solution of \mathcal{C} and $T_i \vdash^? u_i$ be a minimal unsolved constraint of \mathcal{C} . Let u be a term. If there is a simple proof of $T_i \theta \vdash u$ having the last rule an axiom or a decomposition then there is $t \in st(T_i) \setminus \mathcal{X}$ such that $t\theta = u$.*

Proof : Let Π be a simple proof of $T_i \theta \vdash u$ such that its last rule is an axiom or a decomposition. Let j be the minimal indice such that $T_j \theta \vdash u$. Note that $j \leq i$ and by definition of a simple proof, we have that Π is also a simple proof of $T_j \theta \vdash u$.

- The last rule is an axiom. Then $u \in T_j \theta$. There is $t \in T_j$ (thus $t \in st(T_j)$) such that $t\theta = u$. If t is a variable then $T_t \vdash^? t$ is a constraint in \mathcal{C} with $T_t \subsetneq T_j$ (see the definition of a constraint system). Hence $T_t \theta \vdash t\theta$, that is $T_t \theta \vdash u$, which contradicts the minimality of j . Thus, as required, t is not a variable.
- The last rule is a decomposition. Suppose that it is a symmetric decryption. That is, there is w such that $T_j \theta \vdash \text{senc}(u, w)$, and $T_j \theta \vdash w$. By simplicity of the proof, the last rule applied when obtaining $\text{senc}(u, w)$ is an axiom or a decomposition, otherwise the same node would appear twice. Then, applying the induction hypothesis we have that there is $t \in st(T_j) \setminus \mathcal{X}$ such that $t\theta = \text{senc}(u, w)$. It follows that $t = \text{senc}(t', t'')$ with $t'\theta = u$. If t' is a variable then $T_{t'} \theta \vdash t'\theta$. That is $T_{t'} \theta \vdash u$, which again contradicts the minimality of j . Hence t' is not variable, as required.

For the other decomposition rules the same reasoning holds. □

Lemma 4.5 *Let $\mathcal{C} = T_0 \vdash^? x_0, \dots, T_{i-1} \vdash^? x_{i-1}, T_i \vdash^? u, \dots$ be a constraint system and σ be a solution of \mathcal{C} such that*

1. T_i does not contain two distinct subterms t_1, t_2 with $t_1\sigma = t_2\sigma$,
2. u is a non-variable subterm of T_i .

Then $T'_i \vdash u$, where $T'_i = T_i \cup \{x \mid (T \vdash^? x) \in \mathcal{C}, T \subsetneq T_i\}$.

Proof : Let j be minimal such that $T_j \sigma \vdash u\sigma$. Thus $j \leq i$ and $T_j \subseteq T_i$. Consider a simple proof Π of $T_j \sigma \vdash u\sigma$. We reason by induction on the depth of Π .

Base case: Π is reduced to a leaf. Then there is $t \in T_j$ such that $t\sigma = u\sigma$. By hypothesis 1, we deduce that $t = u$. Hence, we have that $u \in T_j$ and thus $T'_i \vdash u$, as required.

Induction step: We analyse the different cases, depending on the last rule R of Π :

- *Case R is a composition rule.* Assume for example that $R = SE$. In such a case, we have that:

$$\Pi = \left\{ \frac{\frac{\Pi_1 \quad \Pi_2}{v_1 \quad v_2}}{\text{senc}(v_1, v_2)} \right.$$

with $u\sigma = \text{senc}(v_1, v_2)$. Since u is not a variable, $u = \text{senc}(u_1, u_2)$, $u_1\sigma = v_1$, and $u_2\sigma = v_2$. If u_1 (resp. u_2) is a variable then u_1 (resp. u_2) belongs to $fv(T_i)$ since $u \in st(T_i)$. Again, this implies $u_1 \in T'_i$ (resp. $u_2 \in T'_i$). Otherwise, u_1 (resp. u_2) is not a variable. Then, by induction hypothesis, $T'_i \vdash u_1$ (resp. $T'_i \vdash u_2$). Hence in both cases we have that $T'_i \vdash u_1$ and $T'_i \vdash u_2$. This allows us to conclude that $T'_i \vdash u$.

- *Case $R = SD$.* In such a case, there is w such that $T_j\sigma \vdash \text{senc}(u\sigma, w)$, and $T_j\sigma \vdash w$:

$$\Pi = \left\{ \frac{\frac{\Pi_1 \quad \Pi_2}{\text{senc}(u\sigma, w) \quad w}}{u\sigma} \right.$$

By simplicity, the last rule of the proof Π_1 is a decomposition or an axiom. By Lemma 4.4, there is $t \in st(T_j) \setminus \mathcal{X}$ such that $t\sigma = \text{senc}(u\sigma, w)$. Let $t = \text{senc}(t_1, t_2)$ with $t_1\sigma = u\sigma$, and $t_2\sigma = w$. By induction hypothesis, $T'_i \vdash t$. Since $t_1\sigma = u\sigma$, by hypothesis 1, we have that $t_1 = u$.

Now, if t_2 is a variable, and since $t_2 \in fv(T_i)$, we have that $T_{t_2} \subsetneq T_i$ and thus $t_2 \in T'_i$. If t_2 is not a variable, then, from $T_j\sigma \vdash t_2\sigma$ and by induction hypothesis, $T'_i \vdash t_2$. So, in any case, $T'_i \vdash t_2$.

Hence, we have both that $T'_i \vdash \text{senc}(u, t_2)$ and $T'_i \vdash t_2$, from which we conclude that $T'_i \vdash u$, by symmetric decryption.

- *Case $R = PKD$.* In such a case, there is w such that $T_j\sigma \vdash \text{sk}(w)$ and $T_j\sigma \vdash \text{aenc}(u\sigma, w)$. As in the previous case, there is $t \in st(T_j) \setminus \mathcal{X}$ such that $t\sigma = \text{aenc}(u\sigma, w)$. By induction hypothesis, $T'_i \vdash t$. Let $t = \text{aenc}(t_1, t_2)$. As in the previous case, we have that $t_1\sigma = u\sigma$, and thus $t_1 = u$ (thanks to hypothesis 1).

The last rule in the proof of $T_j\sigma \vdash \text{sk}(w)$ is a decomposition (no composition rule can yield a term headed with $\text{sk}(-)$). Then, by Lemma 4.4 (T_j satisfies the hypotheses of the lemma since $T_j \subseteq T_i$), there is a non-variable subterm $w_1 \in st(T_j)$ such that $w_1\sigma = \text{sk}(w)$. Let $w_1 = \text{sk}(w_2)$. By induction hypothesis, $T'_j \vdash \text{sk}(w_2)$. Moreover, since $w_2\sigma = t_2\sigma$, by hypothesis 2, we have that $w_2 = t_2$.

Finally, from $T'_i \vdash \text{aenc}(u, w_2)$ and $T'_i \vdash \text{sk}(w_2)$, we conclude that $T'_i \vdash u$.

The proof is similar for the other decomposition rules. □

Proposition 4.2 (Completeness for one step) *If \mathcal{C} is an unsolved deducibility constraint system and σ is a solution of \mathcal{C} , then there is a deducibility constraint system \mathcal{C}' , a substitution τ , and a solution σ' of \mathcal{C}' such that $\mathcal{C} \rightsquigarrow_\tau \mathcal{C}'$ and $\sigma = \tau\sigma'$.*

Proof : Let \mathcal{C} be an unsolved constraint system and σ be a solution of \mathcal{C} . We show that there is a constraint system \mathcal{C}' and a solution σ' of \mathcal{C}' such that $\mathcal{C} \rightsquigarrow_\tau \mathcal{C}'$ and $\sigma = \tau\sigma'$.

Consider a minimal unsolved constraint $T_i \vdash^? u_i$ such that u_i is not a variable. We have that $T_i\sigma \vdash u_i\sigma$. Consider a simple proof Π of $T_i\sigma \vdash u_i\sigma$. We analyse the different cases depending on the last rule of Π .

1. *The last rule is a composition.* Suppose that it is the pairing rule. That is, there are w_1, w_2 such that $T_i\theta \vdash w_1$, $T_i\theta \vdash w_2$ and $\langle w_1, w_2 \rangle = u_i\theta$. Since u_i is not a variable there exists u', u'' such that $u_i = \langle u', u'' \rangle$. Hence we can apply the simplification rule R_f in order to obtain C' . Since $u'\theta = w_1$ and $u''\theta = w_2$, the substitution θ is also a solution to C' . For the other composition rules the same reasoning holds.
2. *The last rule is an axiom or a decomposition.* Applying Lemma 4.4 we obtain that there is $t \in st(T_i) \setminus \mathcal{X}$ such that $t\theta = u_i\theta$. We can have the following two possibilities:
 - (a) If $t \neq u_i$ then we apply the simplification rule R_2 .
 - (b) Otherwise, if $t = u_i$, then $u_i \in st(T_i)$ and we already know that u_i is not a variable. We consider two cases:
 - i. There are two distinct terms $t_1, t_2 \in st(T)$ such that $t_1\theta = t_2\theta$. Then we apply the simplification rule R_3 .
 - ii. Otherwise, the simplification rule R_1 can be applied (Lemma 4.5). \square

4.3.3 Complexity

The termination stated in Theorem 4.1 does not provide with tight complexity bounds. In fact, applying the simplification rules may lead to branches of exponential length in the size of the constraint system [?]. Inspecting the completeness proof, there is still some room for choosing a strategy to ensure that the length of each branch is polynomially bounded in \mathcal{C} (while keeping completeness). Note that correctness is independent of the order of the rules application.

Moreover, for any suitable representation of terms, we have that $|u\sigma, v\sigma| < |u, v|$ where $\sigma = \text{mgu}(u, v)$. Hence, if we use a DAG representation of terms, when $\mathcal{C} \rightsquigarrow_\sigma^* \mathcal{C}'$, we have that the size of \mathcal{C}' is polynomially bounded in the size of \mathcal{C} . As a consequence, the security problem is in co-NP and it is actually co-NP-complete [?]. The NP-hardness can be established with a reduction from 3-SAT.

4.4 Further Readings

Many parts of this section are borrowed from [?]. Hence, more details can be found in this paper. Another decision procedure based on constraint simplification rules has been proposed by J. Millen and V. Shmatikov [?]. Many results (e.g. [?, ?]) have been obtained within this framework. In particular, this framework has been extended by several authors to deal with algebraic properties of cryptographic primitives.

4.5 Exercises

Exercise 4.1 *Say whether each couple of terms are unifiable or not. If so, give a most general unifier (mgu).*

1. $\langle x, b \rangle$ and $\langle a, y \rangle$,
2. $\text{aenc}(x, a)$ and $\text{aenc}(b, x)$,
3. $\langle x, y \rangle$ and $\langle \langle y, y \rangle, a \rangle$,
4. z and $\langle x, y \rangle$.

Exercise 4.2 (\star) Consider the following inference system:

$$\frac{x \quad y}{\langle x, y \rangle} \quad \frac{\langle x, y \rangle}{x} \quad \frac{\langle x, y \rangle}{y} \quad \frac{x \quad y}{\text{senc}(x, y)} \quad \frac{\text{senc}(x, y) \quad y}{x}$$

Let $T = \{\text{senc}(s, \langle k_1, k_2 \rangle), \text{senc}(k_1, k_3), k_3, k_2\}$.

1. Enumerate all the subterms of T .
2. The term s is deducible from T . Give a derivation witnessing this fact.
3. Among the subterms of T , give those that are deducible.
4. Give a term u that is not a subterm of T and such that $T \vdash u$.

Exercise 4.3 ($\star\star\star$) Consider the following inference system:

$$\frac{x \quad y}{\langle x, y \rangle} \quad \frac{\langle x, y \rangle}{x} \quad \frac{\langle x, y \rangle}{y} \quad \frac{x \quad y}{\text{senc}(x, y)} \quad \frac{\text{senc}(x, y) \quad y}{x}$$

In order to decide whether a term s is deducible from a set of terms T in the inference system described above, we propose the following algorithm:

Algorithm:

1. Apply as much as possible the decryption and the projection rules. This leads to a set of terms called $\text{analz}(T)$.
2. Check whether s can be obtained by applying the encryption and the pairing rules. The (infinite) set of terms obtained by applying the composition rules is denoted $\text{synth}(\text{analz}(T))$.

If $s \in \text{synth}(\text{analz}(T))$ then the algorithm return yes. Otherwise, it returns no.

1. Show that this algorithm terminates.
2. Show that this algorithm is sound, i.e. if the algorithm returns yes then $T \vdash s$.
3. The algorithm is not complete, i.e. there exist T and s such that $T \vdash s$, and for which the algorithm returns no. Find an example illustrating this fact.
4. Give an hypothesis on T that allows one to restore completeness.
5. Show that the algorithm is complete when this hypothesis is fulfilled.

Exercise 4.4 (★) We consider the following inference system allowing us to model asymmetric encryption.

$$\frac{x \quad y}{\text{aenc}(x, y)} \quad \frac{\text{aenc}(x, \text{pk}(z)) \quad \text{sk}(z)}{x} \quad \frac{z}{\text{pk}(z)}$$

Is this inference system local, or not? If so, give a proof. If not, give a derivation witnessing this fact.

Exercise 4.5 (★★) Consider the following inference system allowing us to model digital signature.

$$\frac{x \quad \text{sk}(z)}{\text{sign}(x, \text{sk}(z))} \quad \frac{\text{sign}(x, \text{sk}(z)) \quad \text{vk}(z)}{x} \quad \frac{z}{\text{vk}(z)}$$

1. This inference system is not local according to Definition 4.2. Give an example witnessing this fact.

2. Show that the intruder deduction problem is decidable.

You can use the technique described in this chapter and extend the notion of subterm to restore the locality property.

Exercise 4.6 (★) We consider the signature and the inference system given in Example 4.1. Let $T_0 = \{a, b, c, \text{sk}(c), \text{aenc}(\langle a, \text{aenc}(s, b) \rangle, b)\}$ and $\mathcal{C} = \{T_0 \vdash^? \text{aenc}(\langle a, \text{aenc}(x_1, b) \rangle, b)\}$. What are the solutions of \mathcal{C} ?

Exercise 4.7 (★★) Consider the following protocol (defined informally):

$$\begin{aligned} A \rightarrow B &: \langle \text{aenc}(k_1, \text{pk}(b)), \text{aenc}(k_2, \text{pk}(b)) \rangle \\ B \rightarrow A &: \text{senc}(k_1, k_2) \end{aligned}$$

Here k_1 and k_2 represent two keys that are freshly generated by A at the beginning of each session.

1. Write formally the processes corresponding to an instance of the role A played by two honest agents a, b and an instance of the role B with the same two honest agents
2. Give a deduction constraint system corresponding to the only relevant symbolic trace for the processes of the previous question.
3. Apply the simplification rules to this constraint system and derive all possible attacks on the secrecy of k_1 (resp. k_2) for this scenario.

Exercise 4.8 (★) Give an example showing that the rule R_3 is necessary for the completeness of the procedure. More precisely, this example has to show that Proposition 4.2 will be wrong without this rule.

Exercise 4.9 (★★) We consider the following variant of the rule R_3

$$R'_3: \quad \mathcal{C} \wedge T \vdash^? u \rightsquigarrow_\sigma \mathcal{C}\sigma \wedge T\sigma \vdash^? u\sigma \text{ if } t_1, t_2 \in st(T) \setminus \mathcal{X}, \sigma = \text{mgu}(t_1, t_2), \text{ and } t_1 \neq t_2$$

1. Show that R'_3 is not sufficient to restore completeness, i.e. give an example witnessing the fact that Proposition 4.2 is wrong if we use the rule R'_3 instead of R_3 .
2. Consider the set of simplification rules R_1, R_2, R'_3, R_f . Show that this set of rules is complete if we consider symmetric encryption/decryption and pairing/projection (no asymmetric encryption).

Chapter 5

Unbounded process verification

In this chapter, we first show that the problem of the verification of security protocols with an unbounded number of sessions is undecidable. Next, we present a technique for solving this problem, based on an abstract representation of the protocol by Horn clauses and implemented in the tool ProVerif. Obviously, because of the undecidability, this technique does not always solve the problem. (It may answer “I don’t know” or not terminate.)

5.1 Undecidability

First observe that the existence of an attack is undecidable:

Theorem 5.1 *Given a process P (in our algebra) and a term t , the question of whether there are σ, Q such that $P \xrightarrow{*} \nu \bar{n}(\sigma \| Q)$ and $\nu \bar{n}.\sigma \vdash t$ is undecidable.*

This is even true when the attacker is passive, when the keys are atomic and/or the messages are of bounded sizes ([?] for instance).

To give an idea of the proof, we encode a (modified) PCP:

The following problem is undecidable:

Input: two finite sequences $v_0, v_1, \dots, v_n, w_0, w_1, \dots, w_n \in \{a, b\}^*$

Question: are there a k and a sequence of indices $i_1, \dots, i_k \in [1..n]$ such that $v_0.v_{i_1} \dots .v_{i_k} = w_0.w_{i_1} \dots .w_{i_k}$?

For any word $u \in \{a, b\}^*$, we define inductively \tilde{u} as a function from terms to terms as follows:

- if u is the empty word, then \tilde{u} is the identity
- $\widetilde{a \cdot u}(t) = \langle a, \tilde{u}(t) \rangle$.

$$P_A = \text{out}(c, \{\langle \tilde{v}_0(0), \tilde{w}_0(0) \rangle\}_{k_{AB}}^s) . \text{in}(c, \{\langle x, x \rangle\}_{k_{AB}}^s) . \text{out}(c, k_{AB})$$

$$P_B = \text{in}(c, \{\langle x, y \rangle\}_{k_{AB}}^s) . \text{out}(c, \{\langle \tilde{v}_1(x), \tilde{w}_1(x) \rangle\}_{k_{AB}}^s) \dots \text{out}(c, \{\langle \tilde{v}_n(x), \tilde{w}_n(x) \rangle\}_{k_{AB}}^s) \dots$$

Let $P = \nu k_{AB} . (P_A \| !P_B)$. We can deduce k_{AB} iff the modified PCP has a solution. Note that the attacker has not much to do: only select one of the outputs of P_B and send it to the next instance of P_B .

5.2 Structure and Main Features of ProVerif

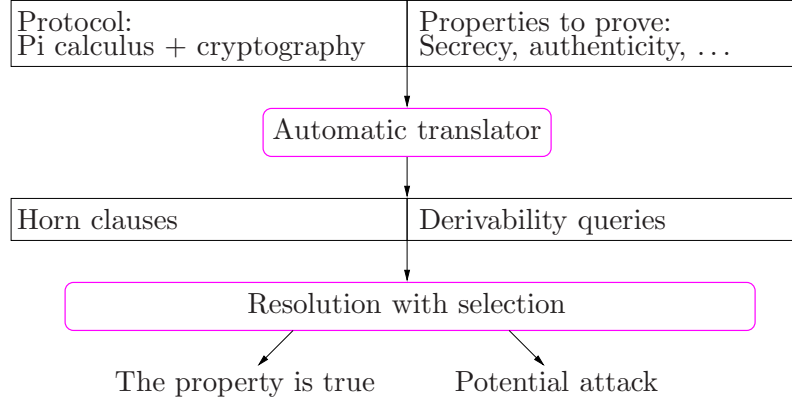


Figure 5.1: Structure of ProVerif

The structure of ProVerif is represented in Figure 5.1. ProVerif takes as input a model of the protocol in an extension of the pi calculus with cryptography, similar to the calculus of Chapter 3. It supports a wide variety of cryptographic primitives, modeled by rewrite rules or by equations. ProVerif also takes as input the security properties that we want to prove. It can verify various security properties, including secrecy, authentication (correspondences), and some observational equivalence properties. It automatically translates this information into an internal representation by Horn clauses: the protocol is translated into a set of Horn clauses, the security properties to prove are translated into derivability queries on these clauses. ProVerif uses an algorithm based on resolution with free selection to determine whether a fact is derivable from the clauses. If the fact is *not* derivable, then the desired security property is proved. If the fact is derivable, then there may be an attack against the considered property: the derivation may correspond to an attack, but it may also correspond to a “false attack”, because the Horn clause representation makes some abstractions. These abstractions are key to the key of an unbounded number of sessions of protocols.

Section 5.3 presents the model of protocols. Section 5.4 presents the Horn clause representation of protocols and the resolution algorithm. Section 5.5 gives the translation from the pi calculus model to Horn clauses for secrecy properties, with extensions to correspondences and equivalences in Sections 5.6 and 5.7 respectively.

5.3 A Formal Model of Security Protocols

This section details the model of protocols used by ProVerif. This calculus was presented in [?]; we adapt that presentation.

5.3.1 Syntax and Informal Semantics

Figure 5.2 gives the syntax of terms (data) and processes (programs) of ProVerif’s input language. The identifiers a, b, c, k , and similar ones range over names, and x, y , and z range over variables. The syntax also assumes a set of symbols for constructors and destructors; we often use f for a constructor and g for a destructor.

$M, N ::=$	terms
x, y, z	variable
a, b, c, k, s	name
$f(M_1, \dots, M_n)$	constructor application
$P, Q ::=$	processes
$\text{out}(M, N).P$	output
$\text{in}(M, x).P$	input
0	nil
$P \parallel Q$	parallel composition
$!P$	replication
$\text{new } a.P$	restriction
$\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q$	destructor application
$\text{let } x = M \text{ in } P$	local definition
$\text{if } M = N \text{ then } P \text{ else } Q$	conditional

Figure 5.2: Syntax of the process calculus

Constructors are used to build terms. Therefore, the terms are variables, names, and constructor applications of the form $f(M_1, \dots, M_n)$; the terms are untyped. On the other hand, destructors do not appear in terms, but only manipulate terms in processes. They are partial functions on terms that processes can apply. The process $\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q$ tries to evaluate $g(M_1, \dots, M_n)$; if this succeeds, then x is bound to the result and P is executed, else Q is executed. More precisely, the semantics of a destructor g of arity n is given by a set $\text{def}(g)$ of rewrite rules of the form $g(M_1, \dots, M_n) \rightarrow M$ where M_1, \dots, M_n, M are terms without names, and the variables of M also occur in M_1, \dots, M_n . We extend these rules by $g(M'_1, \dots, M'_n) \rightarrow M'$ if and only if there exist a substitution σ and a rewrite rule $g(M_1, \dots, M_n) \rightarrow M$ in $\text{def}(g)$ such that $M'_i = \sigma M_i$ for all $i \in \{1, \dots, n\}$, and $M' = \sigma M$. We assume that the set $\text{def}(g)$ is finite. (It usually contains one or two rules in examples.)

Using these constructors and destructors, we can represent data structures, such as tuples, and cryptographic operations, for instance as follows:

- $\text{ntuple}(M_1, \dots, M_n)$ is the tuple of the terms M_1, \dots, M_n , where ntuple is a constructor. (We sometimes abbreviate $\text{ntuple}(M_1, \dots, M_n)$ to (M_1, \dots, M_n) .) The n projections are destructors ith_n for $i \in \{1, \dots, n\}$, defined by

$$\text{ith}_n(\text{ntuple}(x_1, \dots, x_n)) \rightarrow x_i$$

- $\text{senc}(M, N)$ is the symmetric (shared-key) encryption of the message M under the key N , where senc is a constructor. The corresponding destructor sdec is defined by

$$\text{sdec}(\text{senc}(x, y), y) \rightarrow x$$

Thus, $\text{sdec}(M', N)$ returns the decryption of M' if M' is a message encrypted under N .

- In order to represent asymmetric (public-key) encryption, we may use two constructors pk and aenc : $\text{pk}(M)$ builds a public key from a secret M and $\text{aenc}(M, N)$ encrypts M under N . The corresponding destructor adec is defined by

$$\text{adec}(\text{aenc}(x, \text{pk}(y)), y) \rightarrow x$$

- As for digital signatures, we may use a constructor `sign`, and write `sign(M, N)` for M signed with the signature key N , and the two destructors `check` and `getmess` with the rewrite rules:

$$\begin{aligned}\text{check}(\text{sign}(x, y), \text{pk}(y)) &\rightarrow x \\ \text{getmess}(\text{sign}(x, y)) &\rightarrow x\end{aligned}$$

- We may represent a one-way hash function by the constructor h . There is no corresponding destructor; so we model that the term M cannot be retrieved from its hash $h(M)$.

Thus, the process calculus supports many of the operations common in security protocols. It has limitations, though: for example, modular exponentiation or XOR cannot be directly represented by a constructor or by a destructor. We explain how we can treat some of these primitives in Section 5.5.3.

Exercise 5.1 *Model signatures that do not reveal the signed message.*

The other constructs in the syntax of Figure 5.2 are standard; most of them come from the pi calculus, and were also present in the calculus of Chapter 3.

- The input process $\text{in}(M, x).P$ inputs a message on channel M , and executes P with x bound to the input message. The output process $\text{out}(M, N).P$ outputs the message N on the channel M and then executes P . Here, we use an arbitrary term M to represent a channel: M can be a name, a variable, or a constructor application. The calculus is monadic (in that the messages are terms rather than tuples of terms), but a polyadic calculus can be simulated since tuples are terms. It is also synchronous (in that a process P is executed after the output of a message). As usual, we may omit P when it is 0.
- The nil process 0 does nothing.
- The process $P \parallel Q$ is the parallel composition of P and Q .
- The replication $!P$ represents an unbounded number of copies of P in parallel.
- The restriction $\text{new } a.P$ creates a new name a , and then executes P .
- The local definition $\text{let } x = M \text{ in } P$ executes P with x bound to the term M .
- The conditional $\text{if } M = N \text{ then } P \text{ else } Q$ executes P if M and N reduce to the same term at runtime; otherwise, it executes Q . As usual, we may omit an `else` clause when it consists of 0.

The name a is bound in the process $\text{new } a.P$. The variable x is bound in P in the processes $\text{in}(M, x).P$, $\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q$, and $\text{let } x = M \text{ in } P$. We write $fn(P)$ and $fv(P)$ for the sets of names and variables free in P , respectively. A process is closed if it has no free variables; it may have free names. We write $\{M_1/x_1, \dots, M_n/x_n\}$ for the substitution that replaces x_1, \dots, x_n with M_1, \dots, M_n , respectively. When σ is such a substitution and D is some expression, we may write σD or $D\sigma$ for the result of applying σ to D ; the distinction is

one of emphasis at most. Except when stated otherwise, substitutions always map variables (not names) to expressions.

ProVerif's calculus resembles the applied pi calculus [?]. Both calculi are extensions of the pi calculus with (fairly arbitrary) functions on terms. However, there are also important differences between these calculi. The first one is that ProVerif uses destructors instead of the equational theories of the applied pi calculus. (Section 5.5.3 contains further material on equational theories.) The second difference is that ProVerif has a built-in error-handling construct (the `else` clause of the destructor application), whereas in the applied pi calculus the error-handling must be done “by hand”.

5.3.2 Example

We use as a running example a simplified version of the Denning-Sacco key distribution protocol [?], omitting certificates and timestamps:

Message 1. $A \rightarrow B : \{\{k\}_{sk_A}\}_{pk_B}$
 Message 2. $B \rightarrow A : \{s\}_k$

This protocol involves two principals A and B . The key sk_A is the secret key of A , pk_A its public key. Similarly, sk_B and pk_B are the secret and public keys of B , respectively. The key k is a fresh key created by A . A sends this key signed with its private key sk_A and encrypted under the public key of B , pk_B . When B receives this message, B decrypts it and assumes, seeing the signature, that the key k has been generated by A . Then B sends a secret s encrypted under k . Only A should be able to decrypt the message and get the secret s . (The second message is not really part of the protocol, we use it to check if the key k can really be used to exchange secrets between A and B . In fact, there is an attack against this protocol [?], so s will not remain secret.)

This protocol can be encoded by the following process:

$$\begin{aligned} P_0 &= \text{new } sk_A. \text{new } sk_B. \text{let } pk_A = \text{pk}(sk_A) \text{ in let } pk_B = \text{pk}(sk_B) \text{ in out}(c, pk_A). \text{out}(c, pk_B). \\ &\quad (P_A(pk_A, sk_A) \parallel P_B(pk_B, sk_B, pk_A)) \\ P_A(pk_A, sk_A) &= ! \text{in}(c, x_pk_B). \text{new } k. \text{out}(c, \text{aenc}(\text{sign}(k, sk_A), x_pk_B)). \\ &\quad \text{in}(c, x). \text{let } z = \text{sdec}(x, k) \text{ in } 0 \\ P_B(pk_B, sk_B, pk_A) &= ! \text{in}(c, y). \text{let } y' = \text{adec}(y, sk_B) \text{ in} \\ &\quad \text{let } x_k = \text{check}(y', pk_A) \text{ in out}(c, \text{senc}(s, x_k)) \end{aligned}$$

Such a process can be given as input to ProVerif, in an ASCII syntax. This process first creates the secret keys sk_A and sk_B , computes the corresponding public keys pk_A and pk_B , and sends these keys on the public channel c , so that the adversary has these public keys. Then, it runs the processes P_A and P_B in parallel. These processes correspond respectively to the roles of A and B in the protocol. They both start with a replication, which makes it possible to model an unbounded number of sessions of the protocol.

The process P_A first receives on the public channel c the key x_pk_B , which is the public key of A 's interlocutor in the protocol. This message is not strictly speaking part of the protocol; it makes it possible for the adversary to choose with whom A is going to execute a session. In a standard session of the protocol, this key is pk_B , but the adversary can also choose another key, for instance one of his own keys. Then, P_A executes the role of A : it creates a fresh

key k , signs it with its secret key sk_A , then encrypts this message under x_{pk_B} , and sends the obtained message on channel c . P_A then expects the second message of the protocol on channel c , stores it in x and decrypts it. If decryption succeeds, the result (normally the secret s) is stored in z .

The process P_B receives the first message of the protocol on channel c , stores it in y , decrypts it with sk_B , and verifies the signature with pk_A . (The signature is verified with the key pk_A of A and not with an arbitrary key chosen by the adversary since B sends the second message $\{s\}_k$ only if its interlocutor is the honest participant A .) If these verifications succeed, B believes that x_k is a key shared between A and B , and it sends the secret s encrypted under x_k . If the protocol is correct, s should remain secret.

In the above model, we have assumed for simplicity that A and B each play only one role of the protocol. One could easily write a more general model in which they play both roles, or one could even provide the adversary with an interface that allows it to dynamically create new protocol participants.

Exercise 5.2 *Model the Needham-Schroeder public key protocol of Figure 2.1 in this calculus.*

5.3.3 Formal Semantics

The formal semantics of this calculus can be defined in two ways. We can use a structural congruence and a reduction relation (Figure 5.3), as in Chapter 3. We identify processes up to renaming of bound names and variables. We just add reduction rules for the destructor application, (Red Destr 1) and (Red Destr 2), corresponding respectively to the success and failure of the destructor application. In rule (Red I/O), we allow communications on channels that can be any term. We do not consider frames, extended processes, or labeled transitions. (We define the adversary as a process in Definition 5.1 below, so we do not need to study the interaction of the protocol with an external environment, which is the purpose of the labeled semantics.)

We can also define the semantics by a reduction relation on semantic configurations, as in Figure 5.4. A semantic configuration is a pair E, \mathcal{P} where the environment E is a finite set of names and \mathcal{P} is a finite multiset of closed processes. The environment E must contain at least all free names of processes in \mathcal{P} . The configuration $\{a_1, \dots, a_n\}, \{P_1, \dots, P_n\}$ corresponds intuitively to the process $\text{new } a_1 \dots \text{new } a_n. (P_1 \parallel \dots \parallel P_n)$. The semantics of the calculus is defined by a reduction relation \rightarrow on semantic configurations, shown in Figure 5.4. The rule (Red Res) is the only one that uses renaming. The new semantics (in particular the fact that only (Red Res) uses renaming) provides simplifications in the definitions of correspondences (see Section 5.3.4.2) and in the proofs that correspondences hold. Except when mentioned otherwise, we will focus on this new semantics.

5.3.4 Security Properties

5.3.4.1 Secrecy

We assume that the protocol is executed in the presence of an adversary that can listen to all messages, compute, and send all messages it has, following the so-called Dolev-Yao model [?]. Thus, an adversary can be represented by any process that has a set of public names S in its initial knowledge. (Although the initial knowledge of the adversary contains only names in S , one can give any terms to the adversary by sending them on a channel in S .)

$$\begin{array}{ll}
P \parallel 0 \equiv P & P \equiv Q \Rightarrow P \parallel R \equiv Q \parallel R \\
P \parallel Q \equiv Q \parallel P & P \equiv Q \Rightarrow !P \equiv !Q \\
(P \parallel Q) \parallel R \equiv P \parallel (Q \parallel R) & P \equiv Q \Rightarrow \text{new } a.P \equiv \text{new } a.Q \\
\text{new } a_1.\text{new } a_2.P \equiv \text{new } a_2.\text{new } a_1.P & P \equiv P \\
\text{new } a.(P \parallel Q) \equiv P \parallel \text{new } a.Q \text{ if } a \notin \text{fn}(P) & Q \equiv P \Rightarrow P \equiv Q \\
& P \equiv Q, Q \equiv R \Rightarrow P \equiv R \\
\\
\text{out}(N, M).Q \parallel \text{in}(N, x).P \rightarrow Q \parallel P\{M/x\} & (\text{Red I/O}) \\
\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q \rightarrow P\{M'/x\} & (\text{Red Destr 1}) \\
\quad \text{if } g(M_1, \dots, M_n) \rightarrow M' & \\
\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q \rightarrow Q & (\text{Red Destr 2}) \\
\quad \text{if there exists no } M' \text{ such that } g(M_1, \dots, M_n) \rightarrow M' & \\
\text{let } x = M \text{ in } P \rightarrow P\{M/x\} & (\text{Red Let}) \\
\text{if } M = M \text{ then } P \text{ else } Q \rightarrow P & (\text{Red Cond 1}) \\
\text{if } M = N \text{ then } P \text{ else } Q \rightarrow Q \text{ if } M \neq N & (\text{Red Cond 2}) \\
!P \rightarrow P \parallel !P & (\text{Red Repl}) \\
P \rightarrow Q \Rightarrow P \parallel R \rightarrow Q \parallel R & (\text{Red Par}) \\
P \rightarrow Q \Rightarrow \text{new } a.P \rightarrow \text{new } a.Q & (\text{Red Res}) \\
P' \equiv P, P \rightarrow Q, Q \equiv Q' \Rightarrow P' \rightarrow Q' & (\text{Red } \equiv)
\end{array}$$

Figure 5.3: Structural congruence and reduction

Definition 5.1 *Let S be a finite set of names. The closed process Q is an S -adversary if and only if $\text{fn}(Q) \subseteq S$.*

Intuitively, a process P preserves the secrecy of M when M cannot be output on a public channel, in a run of P with any adversary. Formally, we define that a trace outputs M as follows:

Definition 5.2 *We say that a trace $\mathcal{T} = E_0, \mathcal{P}_0 \rightarrow^* E', \mathcal{P}'$ outputs M if and only if \mathcal{T} contains a reduction $E, \mathcal{P} \cup \{\text{out}(c, M).Q, \text{in}(c, x).P\} \rightarrow E, \mathcal{P} \cup \{Q, P\{M/x\}\}$ for some E, \mathcal{P}, x, P, Q , and $c \in S$.*

We can finally define secrecy:

Definition 5.3 *The closed process P preserves the secrecy of M from S if and only if for any S -adversary Q , for any E_0 containing $\text{fn}(P_0) \cup S \cup \text{fn}(M)$, for any trace $\mathcal{T} = E_0, \{P_0, Q\} \rightarrow^* E', \mathcal{P}'$, the trace \mathcal{T} does not output M .*

This notion of secrecy is similar to that of [?, ?, ?]: a term M is secret if the adversary cannot get it by listening and sending messages, and performing computations.

5.3.4.2 Correspondences

As already mentioned in Section 3.3.2, correspondence properties are used to capture relationships between events that can be expressed in the form "if an event e has been executed,

$E, \mathcal{P} \cup \{0\} \rightarrow E, \mathcal{P}$	(Red Nil)
$E, \mathcal{P} \cup \{!P\} \rightarrow E, \mathcal{P} \cup \{P, !P\}$	(Red Repl)
$E, \mathcal{P} \cup \{P \parallel Q\} \rightarrow E, \mathcal{P} \cup \{P, Q\}$	(Red Par)
$E, \mathcal{P} \cup \{\text{new } a.P\} \rightarrow E \cup \{a'\}, \mathcal{P} \cup \{P\{a'/a\}\}$	(Red Res)
where $a' \notin E$.	
$E, \mathcal{P} \cup \{\text{out}(N, M).Q, \text{in}(N, x).P\} \rightarrow E, \mathcal{P} \cup \{Q, P\{M/x\}\}$	(Red I/O)
$E, \mathcal{P} \cup \{\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q\} \rightarrow E, \mathcal{P} \cup \{P\{M'/x\}\}$	(Red Destr 1)
if $g(M_1, \dots, M_n) \rightarrow M'$	
$E, \mathcal{P} \cup \{\text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q\} \rightarrow E, \mathcal{P} \cup \{Q\}$	(Red Destr 2)
if there exists no M' such that $g(M_1, \dots, M_n) \rightarrow M'$	
$E, \mathcal{P} \cup \{\text{let } x = M \text{ in } P\} \rightarrow E, \mathcal{P} \cup \{P\{M/x\}\}$	(Red Let)
$E, \mathcal{P} \cup \{\text{if } M = M \text{ then } P \text{ else } Q\} \rightarrow E, \mathcal{P} \cup \{P\}$	(Red Cond 1)
$E, \mathcal{P} \cup \{\text{if } M = N \text{ then } P \text{ else } Q\} \rightarrow E, \mathcal{P} \cup \{Q\}$	(Red Cond 2)
if $M \neq N$	

Figure 5.4: Operational semantics

then events e'_1, \dots, e'_n have been previously executed.” Moreover, these events may contain arguments. Correspondences can be used to model authentication [?, ?].

In order to define correspondences, we add an instruction for executing events in the syntax of processes:

$P, Q ::=$	processes
\dots	(see Figure 5.2)
$\text{event}(M).P$	event

The semantics of events is defined by the transition:

$$E, \mathcal{P} \cup \{\text{event}(M).P\} \rightarrow E, \mathcal{P} \cup \{P\} \quad (\text{Red Event})$$

We recall that, in the semantics of Figure 5.4, the rule (Red Res) is the only one that uses renaming. This is important so that the parameters of events are not renamed after the execution of the event, to be able to compare them with the parameters of events executed later. This is why this semantics simplifies the definition of correspondences.

We adapt the definition of an S -adversary, so that it does not contain events.

Non-injective correspondences

Definition 5.4 We say that a trace $\mathcal{T} = E_0, \mathcal{P}_0 \rightarrow^* E', \mathcal{P}'$ executes $\text{event}(M)$ if and only if \mathcal{T} contains a reduction $E, \mathcal{P} \cup \{\text{event}(M).P\} \rightarrow E, \mathcal{P} \cup \{P\}$ for some E, \mathcal{P}, P .

The correspondence $\text{event}(M) \rightsquigarrow \bigwedge_{k=1}^l \text{event}(M_k)$, formally defined below, means intuitively that, if an instance of $\text{event}(M)$ is executed, then a corresponding instance of each of the events $\text{event}(M_1), \dots, \text{event}(M_l)$ has been executed.

Definition 5.5 *The closed process P_0 satisfies the correspondence*

$$\text{event}(M) \rightsquigarrow \bigwedge_{k=1}^l \text{event}(M_k)$$

against S -adversaries if and only if, for any S -adversary Q , for any E_0 containing $\text{fn}(P_0) \cup S \cup \text{fn}(M) \cup \bigcup_k \text{fn}(M_k)$, for any substitution σ , for any trace $\mathcal{T} = E_0, \{P_0, Q\} \rightarrow^ E', \mathcal{P}'$, if \mathcal{T} executes $\text{event}(\sigma M)$, then there exists σ' such that $\sigma' M = \sigma M$ and, for all $k \in \{1, \dots, l\}$, \mathcal{T} executes $\text{event}(\sigma' M_k)$ as well.*

The variables in M are universally quantified (because, in Definition 5.5, σ is universally quantified). The variables in M_k that do not occur in M are existentially quantified (because σ' is existentially quantified).

This definition is very general; a particular case is non-injective agreement:

Definition 5.6 *Non-injective agreement is a correspondence of the form $\text{event}(e(x_1, \dots, x_n)) \rightsquigarrow \text{event}(e'(x_1, \dots, x_n))$.*

Intuitively, the correspondence $\text{event}(e(x_1, \dots, x_n)) \rightsquigarrow \text{event}(e'(x_1, \dots, x_n))$ means that, if an event $e(M_1, \dots, M_n)$ is executed, then the event $e'(M_1, \dots, M_n)$ has also been executed. This definition can be used to represent Lowe's notion of non-injective agreement [?].

Injective correspondences

Definition 5.7 *We say that the event $\text{event}(M)$ is executed at step τ in a trace $\mathcal{T} = E_0, P_0 \rightarrow^* E', \mathcal{P}'$ if and only if the τ -th reduction of \mathcal{T} is of the form $E, \mathcal{P} \cup \{\text{event}(M).P\} \rightarrow E, \mathcal{P} \cup \{P\}$ for some E, \mathcal{P}, P .*

Intuitively, an injective correspondence $\text{event}(M) \rightsquigarrow \text{inj event}(M')$ requires that each event $\text{event}(\sigma M)$ is enabled by distinct events $\text{event}(\sigma M')$, while a non-injective correspondence $\text{event}(M) \rightsquigarrow \text{event}(M')$ allows several events $\text{event}(\sigma M)$ to be enabled by the same event $\text{event}(\sigma M')$.

Definition 5.8 *The closed process P_0 satisfies the correspondence*

$$\text{event}(M) \rightsquigarrow \bigwedge_{k=1}^l \text{inj event}(M_k)$$

against S -adversaries if and only if, for any S -adversary Q , for any E_0 containing $\text{fn}(P_0) \cup S \cup \text{fn}(M) \cup \bigcup_k \text{fn}(M_k)$, for any trace $\mathcal{T} = E_0, \{P_0, Q\} \rightarrow^ E', \mathcal{P}'$, there exist injective functions ϕ_{jk} from a subset of steps in \mathcal{T} to steps in \mathcal{T} such that for all τ , if the event $\text{event}(\sigma M)$ is executed at step τ in \mathcal{T} for some σ , then there exists σ' such that $\sigma' M = \sigma M$ and, for all $k \in \{1, \dots, l\}$, $\phi_k(\tau)$ is defined and $\text{event}(\sigma' M_k)$ is executed at step $\phi_k(\tau)$ in \mathcal{T} .*

The functions ϕ_k map execution steps of events $\text{event}(\sigma M)$ to the execution steps of the events $\text{event}(\sigma' M_k)$ that enable $\text{event}(\sigma M)$. The injectivity of ϕ_k guarantees that distinct executions of $\text{event}(\sigma M)$ correspond to distinct executions of $\text{event}(\sigma' M_k)$.

Definition 5.9 Injective agreement is a correspondence of the form $\text{event}(e(x_1, \dots, x_n)) \rightsquigarrow \text{inj event}(e'(x_1, \dots, x_n))$.

Injective agreement requires that the number of executions of $\text{event}(e(M_1, \dots, M_n))$ is smaller than the number of executions of $\text{event}(e'(M_1, \dots, M_n))$: each execution of $\text{event}(e(M_1, \dots, M_n))$ corresponds to a distinct execution of $\text{event}(e'(M_1, \dots, M_n))$. This corresponds to Lowe's agreement specification [?].

More general definitions of correspondences can be found in [?].

Example As an example, we consider a simplified version of the Woo and Lam one-way public-key authentication protocol, version of [?], in which host names are replaced by public keys, which makes interaction with a server useless. The protocol is:

Message 1. $A \rightarrow B : pk_A$
 Message 2. $B \rightarrow A : b$
 Message 3. $A \rightarrow B : \{pk_A, pk_B, b\}_{sk_A}$

A first sends to B its public key, B replies with a nonce (fresh name), and A sends its public key, the public key of B , and the nonce all signed with its private key sk_A . This protocol can be represented by the process P :

$$\begin{aligned} P_A(sk_A, pk_A) &= ! \text{in}(c, x_pk_B). \text{event}(e_A(x_pk_B)). \text{out}(c, pk_A). \text{in}(c, x_b). \\ &\quad \text{out}(c, \text{sign}((pk_A, x_pk_B, x_b), sk_A)) \\ P_B(sk_B, pk_B, pk_A) &= ! \text{in}(c, x_pk_A). \text{new } b. \text{out}(c, b). \text{in}(c, m). \\ &\quad \text{if } (x_pk_A, pk_B, b) = \text{check}(m, x_pk_A) \text{ then} \\ &\quad \text{if } x_pk_A = pk_A \text{ then event}(e_B(pk_B)) \\ P &= \text{new } sk_A. \text{new } sk_B. \text{let } pk_A = \text{pk}(sk_A) \text{ in let } pk_B = \text{pk}(sk_B) \text{ in} \\ &\quad \text{out}(c, pk_A). \text{out}(c, pk_B). (P_A(sk_A, pk_A) \parallel P_B(sk_B, pk_B, pk_A)) \end{aligned}$$

The channel c is public: The adversary can send and listen on it. The process P begins with the creation of the secret and public keys of A and B . The public keys are output on channel c to model that the adversary has them in its initial knowledge. Then the protocol itself starts: P_A represents A , P_B represents B . Both principals can run an unbounded number of sessions, hence the replications. We consider that A and B are both willing to talk to any principal. So, to determine to whom A will talk, we consider that A first inputs a message containing the public key of its interlocutor (this interlocutor is therefore chosen by the adversary). Then A starts the protocol by executing an event $\text{event}(e_A(x_pk_B))$, whose intuitive meaning is “ A has started a session with the host of public key x_pk_B ”. At the end of P_B , when B thinks he talks with A (that is, when $x_pk_A = pk_A$), he executes an event $\text{event}(e_B(pk_B))$, whose meaning is “ B thinks he has completed a session with A ”.¹ If the protocol is correct, B should only execute $\text{event}(e_B(pk_B))$ when A has first executed $\text{event}(e_A(pk_B))$. This can be formalized by the correspondence:

$$\text{event}(e_B(x)) \rightsquigarrow \text{event}(e_A(x))$$

¹Note that the event e_B must not be executed when B thinks he talks to the adversary. Indeed, in this case, it is correct that no event e_A has been executed by the interlocutor of B , since the adversary never executes events.

We may also require that each execution of e_B corresponds to a distinct event e_A , so that different runs of B correspond to different runs of A . This is formalized by the injective correspondence:

$$\text{event}(e_B(x)) \rightsquigarrow \text{inj event}(e_A(x))$$

This is what we are going to check next.

5.3.4.3 Observational Equivalence

The notion of indistinguishability is a powerful concept which allows us to reason about complex properties that cannot be expressed as reachability or correspondence properties. The notion of indistinguishability is generally named *observational equivalence* in the formal model. Intuitively, two processes P and Q are observationally equivalent, written $P \approx Q$, when an active adversary cannot distinguish P from Q . Using this notion, one can for instance specify that a process P follows its specification Q by saying that $P \approx Q$. ProVerif can prove some observational equivalences, but not all of them because their proof is complex. Formally, observational equivalence can be defined as follows, using the first semantics of ProVerif's calculus. An evaluation context C is a closed context built from $[]$, $C \parallel P$, $P \parallel C$, and $\text{new } a.C$.

Definition 5.10 *The process P emits on M ($P \downarrow_M$) if and only if $P \equiv C[\text{out}(M, N).R]$ for some evaluation context C that does not bind $\text{fn}(M)$.*

Observational equivalence \approx is the largest symmetric relation \mathcal{R} on closed processes such that $P \mathcal{R} Q$ implies

1. if $P \rightarrow^* \downarrow_M$ then $Q \rightarrow^* \downarrow_M$;
2. if $P \rightarrow^* P'$ then $Q \rightarrow^* Q'$ and $P' \mathcal{R} Q'$ for some Q' ;
3. $C[P] \mathcal{R} C[Q]$ for all evaluation contexts C .

In the applied pi calculus, this notion coincides with the notion of labeled bisimilarity presented in Section 3.3.4 [?]. It can be defined without using labeled transitions; on the other hand, it uses a universal quantification over all contexts, which complicates its proof. (Typically, the goal is to prove observational equivalence; a proof method is to design a labeled transition system such that observational equivalence is equivalent to labeled bisimilarity, and to prove labeled bisimilarity instead.)

Observational equivalence can also be used to formalize a notion of secrecy, strong secrecy, which means that the attacker is unable to distinguish when the secret changes. In other words, the value of the secret should not affect the observable behavior of the protocol. Such a notion is useful to capture the adversary's ability to learn partial information about the secret: when the adversary learns the first component of a pair, for instance, the whole pair is secret in the sense of reachability (the adversary cannot reconstruct the whole pair because it does not have the second component), but it is not secret in the sense of strong secrecy (the adversary can notice changes in the value of the pair, since it has its first component). Strong secrecy also detects implicit flows, which happen for instance when the secret is compared with another value. The concept of strong secrecy is particularly important when the secret consists of known values. Consider for instance a process P that uses a boolean b . The variable b can take two values, `true` or `false`, which are both known to the adversary, so it is

$p ::=$	patterns
x	variable
$a[p_1, \dots, p_n]$	name
$f(p_1, \dots, p_n)$	constructor application
$F ::= \text{pred}(p_1, \dots, p_n)$	fact
$R ::= F_1 \wedge \dots \wedge F_n \Rightarrow F$	Horn clause

Figure 5.5: Syntax of ProVerif's internal protocol representation

not secret in the sense of reachability. However, one may express that b is strongly secret by saying that $P\{\text{true}/b\} \approx P\{\text{false}/b\}$: the adversary cannot determine whether b is `true` or `false`. Strong secrecy can be formally defined as follows:

Definition 5.11 *P preserves the secrecy of its free variables if and only if for all closed substitutions θ and θ' such that $\text{dom}(\theta) = \text{dom}(\theta') = \text{fv}(P)$, $P\theta \approx P\theta'$. (The substitution avoids name captures by first alpha renaming P if necessary.)*

5.3.5 Some Other Models

In addition to variants of the pi calculus, such as the spi calculus [?] and the applied pi calculus [?], there also exist other models of security protocols, for instance strand spaces [?] and multiset rewriting [?].

5.4 The Horn Clause Representation of Protocols

5.4.1 Definition of this Representation

Internally, ProVerif translates the protocol into a representation by a set of Horn clauses; the syntax of these clauses is given in Figure 5.5. In this figure, x ranges over variables, a over names, f over function symbols, and p over predicate symbols. The patterns p represent messages that are exchanged between participants of the protocol. (Patterns are terms; we use the word patterns to distinguish them from terms of the process calculus.) A variable can represent any pattern. Names represent atomic values, such as keys and nonces (random numbers). Each principal has the ability of creating new names: fresh names are created at each run of the protocol. Here, the created names are considered as functions of the messages previously received by the principal that creates the name. Thus, names are distinguished only when the preceding messages are different. As noticed by Martín Abadi (personal communication), this approximation is in fact similar to the approximation done in some type systems (such as [?]): the type of the new name depends on the types in the environment. It is enough to handle many protocols, and can be enriched by adding other parameters to the name. In particular, we shall see in Section 5.6 that, by adding as parameter a session identifier that takes a different value in each run of the protocol, one can distinguish all names. This is necessary for proving authentication but not for secrecy, so we omit session identifiers here for simplicity. The constructor applications $f(M_1, \dots, M_n)$ build patterns. A fact $F = \text{pred}(p_1, \dots, p_n)$ expresses a property of the messages p_1, \dots, p_n .

Several predicates *pred* can be used but, for a first example, we are going to use a single predicate **attacker**, such that the fact **attacker**(*p*) means “the attacker may have the message *p*”. A clause $R = F_1 \wedge \dots \wedge F_n \Rightarrow F$ means that, if all facts F_1, \dots, F_n are true, then F is also true. A clause with no hypothesis $\Rightarrow F$ is written simply F .

We use illustrate the encoding of a protocol on the example of Section 5.3.2:

$$\begin{aligned} \text{Message 1. } & A \rightarrow B : \{\{k\}_{sk_A}\}_{pk_B} \\ \text{Message 2. } & B \rightarrow A : \{s\}_k \end{aligned}$$

5.4.1.1 Representation of the Abilities of the Attacker

We first present the encoding of the computation abilities of the attacker. The encoding of the protocol itself will be detailed in Section 5.4.1.2.

During its computations, the attacker can apply all constructors and destructors. If f is a constructor of arity n , this leads to the clause:

$$\mathbf{attacker}(x_1) \wedge \dots \wedge \mathbf{attacker}(x_n) \Rightarrow \mathbf{attacker}(f(x_1, \dots, x_n)).$$

If g is a destructor, for each rewrite rule $g(M_1, \dots, M_n) \rightarrow M$ in $\text{def}(g)$, we have the clause:

$$\mathbf{attacker}(M_1) \wedge \dots \wedge \mathbf{attacker}(M_n) \Rightarrow \mathbf{attacker}(M).$$

The destructors never appear in the clauses, they are coded by pattern-matching on their parameters (here M_1, \dots, M_n) in the hypothesis of the clause and generating their result in the conclusion. In the particular case of public-key encryption, this yields:

$$\begin{aligned} \mathbf{attacker}(m) \wedge \mathbf{attacker}(pk) &\Rightarrow \mathbf{attacker}(\mathbf{aenc}(m, pk)), \\ \mathbf{attacker}(sk) &\Rightarrow \mathbf{attacker}(\mathbf{pk}(sk)), \\ \mathbf{attacker}(\mathbf{aenc}(m, \mathbf{pk}(sk))) \wedge \mathbf{attacker}(sk) &\Rightarrow \mathbf{attacker}(m), \end{aligned} \tag{5.1}$$

where the first two clauses correspond to the constructors **aenc** and **pk**, and the last clause corresponds to the destructor **pdec**. When the attacker has an encrypted message **aenc**(m, pk) and the decryption key sk , then it also has the cleartext m . (We assume that the cryptography is perfect, hence the attacker can obtain the cleartext from the encrypted message only if it has the key.)

Clauses for signatures (**sign**, **getmess**, **check**) and for shared-key encryption (**senc**, **sdec**) are given in Figure 5.6.

The clauses above describe the computation abilities of the attacker. Moreover, the attacker initially has the public keys of the protocol participants. Therefore, we add the clauses **attacker**(**pk**($sk_A[]$)) and **attacker**(**pk**($sk_B[]$)). We also give a name a to the attacker, that will represent all names it can generate: **attacker**($a[]$). In particular, $a[]$ can represent the secret key of any dishonest participant, his public key being **pk**($a[]$), which the attacker can compute by the clause for constructor **pk**.

5.4.1.2 Representation of the Protocol Itself

Now, we describe how the protocol itself is represented. We consider that A and B are willing to talk to any principal, A , B but also malicious principals that are represented by the attacker. Therefore, the first message sent by A can be **aenc**(**sign**($k, sk_A[]$), **pk**(x))

for any x . We leave to the attacker the task of starting the protocol with the principal it wants, that is, the attacker will send a preliminary message to A , mentioning the public key of the principal with which A should talk. This principal can be B , or another principal represented by the attacker. Hence, if the attacker has some key $\text{pk}(x)$, it can send $\text{pk}(x)$ to A ; A replies with his first message, which the attacker can intercept, so the attacker obtains $\text{aenc}(\text{sign}(k, sk_A[]), \text{pk}(x))$. Therefore, we have a clause of the form

$$\text{attacker}(\text{pk}(x)) \Rightarrow \text{attacker}(\text{aenc}(\text{sign}(k, sk_A[]), \text{pk}(x))).$$

Moreover, a new key k is created each time the protocol is run. Hence, if two different keys $\text{pk}(x)$ are received by A , the generated keys k are certainly different: k depends on $\text{pk}(x)$. The clause becomes:

$$\text{attacker}(\text{pk}(x)) \Rightarrow \text{attacker}(\text{aenc}(\text{sign}(k[\text{pk}(x)], sk_A[]), \text{pk}(x))). \quad (5.2)$$

When B receives a message, he decrypts it with his secret key sk_B , so B expects a message of the form $\text{aenc}(x', \text{pk}(sk_B[]))$. Next, B tests whether A has signed x' , that is, B evaluates $\text{check}(x', pk_A)$, and this succeeds only when $x' = \text{sign}(y, sk_A[])$. If so, he assumes that the key y is only known by A , and sends a secret s (a constant that the attacker does not have a priori) encrypted under y . We assume that the attacker relays the message coming from A , and intercepts the message sent by B . Hence the clause:

$$\text{attacker}(\text{aenc}(\text{sign}(y, sk_A[]), \text{pk}(sk_B[]))) \Rightarrow \text{attacker}(\text{senc}(s, y)).$$

Remark 5.1 *With these clauses, A cannot play the role of B and vice-versa. In order to model a situation in which all principals play both roles, we can replace all occurrences of $sk_B[]$ with $sk_A[]$ in the clauses above. Then A plays both roles, and is the only honest principal. A single honest principal is sufficient for proving secrecy properties by [?].*

More generally, a protocol that contains n messages is encoded by n sets of clauses. If a principal X sends the i th message, the i th set of clauses contains clauses that have as hypotheses the patterns of the messages previously received by X in the protocol, and as conclusion the pattern of the i th message. There may be several possible patterns for the previous messages as well as for the sent message, in particular when the principal X uses a function defined by several rewrite rules, such as the function exp of Section 5.5.3. In this case, a clause must be generated for each combination of possible patterns. Moreover, the hypotheses of the clauses describe all messages previously received, not only the last one. This is important since in some protocols the fifth message for instance can contain elements received in the first message. The hypotheses summarize the history of the exchanged messages.

5.4.1.3 Summary

To sum up, a protocol can be represented by three sets of Horn clauses, as detailed in Figure 5.6 for the protocol of Section 5.3.2:

- Clauses representing the computation abilities of the attacker: constructors, destructors, and name generation.

Computation abilities of the attacker:

For each constructor f of arity n :

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

For each destructor g , for each rewrite rule $g(M_1, \dots, M_n) \rightarrow M$ in $\text{def}(g)$:

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

that is

aenc	$\text{attacker}(m) \wedge \text{attacker}(pk) \Rightarrow \text{attacker}(\text{aenc}(m, pk))$
pk	$\text{attacker}(sk) \Rightarrow \text{attacker}(\text{pk}(sk))$
pdec	$\text{attacker}(\text{aenc}(m, \text{pk}(sk))) \wedge \text{attacker}(sk) \Rightarrow \text{attacker}(m)$
sign	$\text{attacker}(m) \wedge \text{attacker}(sk) \Rightarrow \text{attacker}(\text{sign}(m, sk))$
getmess	$\text{attacker}(\text{sign}(m, sk)) \Rightarrow \text{attacker}(m)$
check	$\text{attacker}(\text{sign}(m, sk)) \wedge \text{attacker}(\text{pk}(sk)) \Rightarrow \text{attacker}(m)$
senc	$\text{attacker}(m) \wedge \text{attacker}(k) \Rightarrow \text{attacker}(\text{senc}(m, k))$
sdec	$\text{attacker}(\text{senc}(m, k)) \wedge \text{attacker}(k) \Rightarrow \text{attacker}(m)$

Name generation: $\text{attacker}(a[])$

Initial knowledge: $\text{attacker}(\text{pk}(sk_A[])), \text{attacker}(\text{pk}(sk_B[]))$

The protocol:

First message: $\text{attacker}(\text{pk}(x)) \Rightarrow \text{attacker}(\text{aenc}(\text{sign}(k[\text{pk}(x)], sk_A[]), \text{pk}(x)))$

Second message: $\text{attacker}(\text{aenc}(\text{sign}(y, sk_A[]), \text{pk}(sk_B[]))) \Rightarrow \text{attacker}(\text{senc}(s, y))$

Figure 5.6: Summary the Horn clause representation of the protocol of Section 5.3.2

- Facts corresponding to the initial knowledge of the attacker. In general, there are facts giving the public keys of the participants and/or their names to the attacker.
- Clauses representing the messages of the protocol itself. There is one set of clauses for each message in the protocol. In the set corresponding to the i th message, sent by principal X , the clauses are of the form $\text{attacker}(p_{j_1}) \wedge \dots \wedge \text{attacker}(p_{j_n}) \Rightarrow \text{attacker}(p_i)$ where p_{j_1}, \dots, p_{j_n} are the patterns of the messages received by X before sending the i th message, and p_i is the pattern of the i th message.

Exercise 5.3 Give the representation by Horn clauses of the Needham-Schroeder public-key protocol of Figure 2.1.

5.4.1.4 Approximations

The reader can notice that the Horn clause representation of protocols is approximate. Specifically, the number of repetitions of each action is ignored, since Horn clauses can be applied any number of times. So a step of the protocol can be completed several times, as long as the previous steps have been completed at least once between the same principals (even when future steps have already been completed). For instance, consider the following protocol (communicated by Véronique Cortier)

First step: A sends $\{(N_1, M)\}_k, \{(N_2, M)\}_k$
 Second step: If A receives $\{(x, M)\}_k$, he replies with x
 Third step: If A receives N_1, N_2 , he replies with s

where N_1 , N_2 , and M are nonces. In an exact model, A never sends s , since $\{(N_1, M)\}_k$ or $\{(N_2, M)\}_k$ can be decrypted, but not both. In the Horn clause model, even though the first step is executed once, the second step may be executed twice for the same M (that is, the corresponding clause can be applied twice), so that both $\{(N_1, M)\}_k$ and $\{(N_2, M)\}_k$ can be decrypted, and A may send s . We have a false attack against the secrecy of s .


However, the important point is that the approximations are sound: if an attack exists in a more precise model, such as the applied pi calculus [?] or multiset rewriting [?], then it also exists in the Horn clause representation. This is shown for the applied pi calculus in [?] and for multiset rewriting in [?]. In particular, [?] shows formally that the only approximation with respect to the multiset rewriting model is that the number of repetitions of actions is ignored. Performing approximations enables us to build a much more efficient verifier, which will be able to handle larger and more complex protocols. Another advantage is that the verifier does not have to limit the number of runs of the protocol. The price to pay is that false attacks may be found by the verifier: sequences of clause applications that do not correspond to a protocol run, as illustrated above. False attacks appear in particular for protocols with temporary secrets: when some value first needs to be kept secret and is revealed later in the protocol, the Horn clause model considers that this value can be reused in the beginning of the protocol, thus breaking the protocol. When a false attack is found, we cannot know whether the protocol is secure or not: a real attack may also exist. A more precise analysis is required in this case. Fortunately, the Horn clause representation is precise enough so that false attacks are rare. (This is demonstrated by the experiments, see Section 5.8.)

5.4.1.5 Secrecy Criterion

A basic goal is to determine secrecy properties: for instance, can the attacker get the secret s ? That is, can the fact $\text{attacker}(s)$ be derived from the clauses? If $\text{attacker}(s)$ can be derived, the sequence of clauses applied to derive $\text{attacker}(s)$ will lead to the description of an attack. This is the notion of secrecy of Section 5.3.4.1.

In our running example, $\text{attacker}(s)$ is derivable from the clauses. The derivation is as follows. The attacker generates a fresh name $a[]$ (considered as a secret key), it computes $\text{pk}(a[])$ by the clause for pk , obtains $\text{aenc}(\text{sign}(k[\text{pk}(a[])], \text{sk}_A[]), \text{pk}(a[]))$ by the clause for the first message. It decrypts this message using the clause for pdec and its knowledge of $a[]$, thus obtaining $\text{sign}(k[\text{pk}(a[])], \text{sk}_A[])$. It reencrypts the signature under $\text{pk}(\text{sk}_B[])$ by the clause for aenc (using its initial knowledge of $\text{pk}(\text{sk}_B[])$), thus obtaining $\text{aenc}(\text{sign}(k[\text{pk}(a[])], \text{sk}_A[]), \text{pk}(\text{sk}_B[]))$. By the clause for the second message, it obtains $\text{senc}(s, k[\text{pk}(a[])])$. On the other hand, from $\text{sign}(k[\text{pk}(a[])], \text{sk}_A[])$, it obtains $k[\text{pk}(a[])]$ by the clause for getmess , so it can decrypt $\text{senc}(s, k[\text{pk}(a[])])$ by the clause for sdec , thus obtaining s . In other words, the attacker starts a session between A and a dishonest participant of secret key $a[]$. It gets the first message $\text{aenc}(\text{sign}(k, \text{sk}_A[]), \text{pk}(a[]))$, decrypts it, reencrypts it under $\text{pk}(\text{sk}_B[])$, and sends it to B . For B , this message looks like the first message of a session between A and B , so B replies with $\text{senc}(s, k)$, which the attacker can decrypt since it obtains k from the first message. The obtained derivation corresponds to the known attack against this protocol. In contrast, if we fix the protocol by adding the public key of B in the first message $\{\{(pk_B, k)\}_{sk_A}\}_{pk_B}$, $\text{attacker}(s)$ is not derivable from the clauses, so the fixed protocol preserves the secrecy of s .

Next, we formally define when a given fact can be derived from a given set of clauses. We shall see in the next section how we determine that. Technically, the hypotheses F_1, \dots, F_n of a clause are considered as a multiset. This means that the order of the hypotheses is

Figure 5.7: Derivation of F

irrelevant, but the number of times a hypothesis is repeated is important. (This is not related to multiset rewriting models of protocols: the semantics of a clause does not depend on the number of repetitions of its hypotheses, but considering multisets is necessary in the proof of the resolution algorithm.) We use R for clauses (logic programming *rules*), H for hypothesis, and C for conclusion.

Definition 5.12 (Subsumption) We say that $H_1 \Rightarrow C_1$ subsumes $H_2 \Rightarrow C_2$, and we write $(H_1 \Rightarrow C_1) \sqsupseteq (H_2 \Rightarrow C_2)$, if and only if there exists a substitution σ such that $\sigma C_1 = C_2$ and $\sigma H_1 \subseteq H_2$ (multiset inclusion).

We write $R_1 \sqsupseteq R_2$ when R_2 can be obtained by adding hypotheses to a particular instance of R_1 . In this case, all facts that can be derived by R_2 can also be derived by R_1 .

A derivation is defined as follows, as illustrated in Figure 5.7.

Definition 5.13 (Derivability) Let F be a closed fact, that is, a fact without variable. Let \mathcal{R} be a set of clauses. F is derivable from \mathcal{R} if and only if there exists a derivation of F from \mathcal{R} , that is, a finite tree defined as follows:

1. Its nodes (except the root) are labeled by clauses $R \in \mathcal{R}$;
2. Its edges are labeled by closed facts;
3. If the tree contains a node labeled by R with one incoming edge labeled by F_0 and n outgoing edges labeled by F_1, \dots, F_n , then $R \sqsupseteq F_1 \wedge \dots \wedge F_n \Rightarrow F_0$.
4. The root has one outgoing edge, labeled by F . The unique son of the root is named the subroot.

In a derivation, if there is a node labeled by R with one incoming edge labeled by F_0 and n outgoing edges labeled by F_1, \dots, F_n , then F_0 can be derived from F_1, \dots, F_n by the clause R . Therefore, there exists a derivation of F from \mathcal{R} if and only if F can be derived from clauses in \mathcal{R} (in classical logic).

5.4.2 Resolution Algorithm

The internal protocol representation is a set of Horn clauses, and our goal is to determine whether a given fact can be derived from these clauses or not. This is exactly the problem solved by usual Prolog systems. However, we cannot use such systems here, because they would not terminate. For instance, the clause:

$$\text{attacker}(\text{aenc}(m, \text{pk}(sk))) \wedge \text{attacker}(sk) \Rightarrow \text{attacker}(m)$$

leads to considering more and more complex terms, with an unbounded number of encryptions. We could of course limit arbitrarily the depth of terms to solve the problem, but we can do much better than that.

As detailed below, the main idea is to combine pairs of clauses by resolution, and to guide this resolution process by a selection function: ProVerif's resolution algorithm is resolution with free selection [?, ?, ?]. This algorithm is similar to ordered resolution with selection, used by [?], but without the ordering constraints.

Notice that, since a term is secret when a fact is *not* derivable from the clauses, soundness in terms of security (if the verifier claims that there is no attack, then there is no attack) corresponds to the completeness of the resolution algorithm in terms of logic programming (if the algorithm claims that a fact is not derivable, then it is not). The resolution algorithm that we use must therefore be complete.

5.4.2.1 The Basic Algorithm

Let us first define resolution: when the conclusion of a clause R unifies with a hypothesis of another (or the same) clause R' , resolution infers a new clause that corresponds to applying R and R' one after the other. Formally, resolution is defined as follows:

Definition 5.14 *Let R and R' be two clauses, $R = H \Rightarrow C$, and $R' = H' \Rightarrow C'$. Assume that there exists $F_0 \in H'$ such that C and F_0 are unifiable and σ is the most general unifier of C and F_0 . In this case, we define $R \circ_{F_0} R' = \sigma(H \cup (H' \setminus \{F_0\})) \Rightarrow \sigma C'$. The clause $R \circ_{F_0} R'$ is the result of resolving R' with R upon F_0 ; it can be inferred from R and R' :*

$$\frac{R = H \Rightarrow C \quad R' = H' \Rightarrow C'}{R \circ_{F_0} R' = \sigma(H \cup (H' \setminus \{F_0\})) \Rightarrow \sigma C'}$$

For example, if R is the clause (5.2), R' is the clause (5.1), and the fact F_0 is $F_0 = \text{attacker}(\text{aenc}(m, \text{pk}(sk)))$, then $R \circ_{F_0} R'$ is

$$\text{attacker}(\text{pk}(x)) \wedge \text{attacker}(x) \Rightarrow \text{attacker}(\text{sign}(k[\text{pk}(x)], sk_A[]))$$

with the substitution $\sigma = \{sk \mapsto x, m \mapsto \text{sign}(k[\text{pk}(x)], sk_A[])\}$.

We guide the resolution by a selection function:

Definition 5.15 *A selection function sel is a function from clauses to sets of facts, such that $sel(H \Rightarrow C) \subseteq H$. If $F \in sel(R)$, we say that F is selected in R . If $sel(R) = \emptyset$, we say that no hypothesis is selected in R , or that the conclusion of R is selected.*

The resolution algorithm is correct (sound and complete) with any selection function, as we show below. However, the choice of the selection function can change dramatically the behavior of the algorithm. The essential idea of the algorithm is to combine clauses by resolution

$\text{saturate}(\mathcal{R}_0) =$

1. $\mathcal{R} \leftarrow \emptyset$.
For each $R \in \mathcal{R}_0$, $\mathcal{R} \leftarrow \text{elim}(\{R\} \cup \mathcal{R})$.
2. Repeat until a fixpoint is reached
for each $R \in \mathcal{R}$ such that $\text{sel}(R) = \emptyset$,
for each $R' \in \mathcal{R}$, for each $F_0 \in \text{sel}(R')$ such that $R \circ_{F_0} R'$ is defined,
 $\mathcal{R} \leftarrow \text{elim}(\{R \circ_{F_0} R'\} \cup \mathcal{R})$.
3. Return $\{R \in \mathcal{R} \mid \text{sel}(R) = \emptyset\}$.

Figure 5.8: Resolution algorithm

only when the facts unified in the resolution are selected. We will therefore choose the selection function to reduce the number of possible unifications between selected facts. Having several selected facts slows down the algorithm, because it has more choices of resolutions to perform, therefore we will select at most one fact in each clause. In the case of protocols, facts of the form $\text{attacker}(x)$, with x variable, can be unified with all facts of the form $\text{attacker}(p)$. Therefore, we should avoid selecting them. So a basic selection function is a function sel_0 that satisfies the constraint

$$\text{sel}_0(H \Rightarrow C) = \begin{cases} \emptyset & \text{if } \forall F \in H, \exists x \text{ variable, } F = \text{attacker}(x) \\ \{F_0\} & \text{where } F_0 \in H \text{ and } \forall x \text{ variable, } F_0 \neq \text{attacker}(x) \end{cases} \quad (5.3)$$

The resolution algorithm is described in Figure 5.8. It transforms the initial set of clauses into a new one that derives the same facts.

The resolution algorithm, $\text{saturate}(\mathcal{R}_0)$, contains 3 steps.

- The first step inserts in \mathcal{R} the initial clauses representing the protocol and the attacker (clauses that are in \mathcal{R}_0), after elimination of subsumed clauses by elim : if R' subsumes R , and R and R' are in \mathcal{R} , then R is removed by $\text{elim}(\mathcal{R})$.
- The second step is a fixpoint iteration that adds clauses created by resolution. The resolution of clauses R and R' is added only if no hypothesis is selected in R and the hypothesis F_0 of R' that we unify is selected. When a clause is created by resolution, it is added to the set of clauses \mathcal{R} . Subsumed clauses are eliminated from \mathcal{R} .
- At last, the third step returns the set of clauses of \mathcal{R} with no selected hypothesis.

Basically, saturate preserves derivability (it is both sound and complete):

Theorem 5.2 (Correctness of saturate) *Let F be a closed fact. F is derivable from \mathcal{R}_0 if and only if it is derivable from $\text{saturate}(\mathcal{R}_0)$.*

This result is proved by transforming a derivation of F from \mathcal{R}_0 into a derivation of F from $\text{saturate}(\mathcal{R}_0)$. Basically, when the derivation contains a clause R' with $\text{sel}(R') \neq \emptyset$, we replace in this derivation two clauses R , with $\text{sel}(R) = \emptyset$, and R' that have been combined by resolution during the execution of saturate with a single clause $R \circ_{F_0} R'$. This replacement decreases the number of clauses in the derivation, so it terminates, and, upon termination, all

clauses of the obtained derivation satisfy $\text{sel}(R') = \emptyset$ so they are in $\text{saturate}(\mathcal{R}_0)$. A detailed proof is given in Section 5.4.2.2.

Usually, resolution with selection is used for proofs by refutation. That is, the negation of the goal F is added to the clauses, under the form of a clause without conclusion: $F \Rightarrow$. The goal F is derivable if and only if the empty clause “ \Rightarrow ” can be derived. Here, for non-closed goals, we also want to be able to know which instances of the goal can be derived. That is why we prove that the clauses in $\text{saturate}(\mathcal{R}_0)$ derive the same facts as the clauses in \mathcal{R}_0 .

We can determine which instances of $\text{pred}(p_1, \dots, p_n)$ are derivable, as follows:

Corollary 5.1 *Let $\text{solve}_{\mathcal{R}_0}(\text{pred}(p_1, \dots, p_n)) = \{H \Rightarrow \text{pred}(p'_1, \dots, p'_n) \mid H \Rightarrow \text{pred}'(p'_1, \dots, p'_n) \in \text{saturate}(\mathcal{R}'_0)\}$, where pred' is a new predicate and $\mathcal{R}'_0 = \mathcal{R}_0 \cup \{\text{pred}(p_1, \dots, p_n) \Rightarrow \text{pred}'(p'_1, \dots, p'_n)\}$.*

The fact $\sigma \text{pred}(p_1, \dots, p_n)$ is derivable from \mathcal{R}_0 if and only if there exists a clause $H \Rightarrow \text{pred}(p'_1, \dots, p'_n)$ in $\text{solve}_{\mathcal{R}_0}(\text{pred}(p_1, \dots, p_n))$ and a substitution σ' such that $\sigma' \text{pred}(p'_1, \dots, p'_n) = \sigma \text{pred}(p_1, \dots, p_n)$ and $\sigma' H$ is derivable from \mathcal{R}'_0 .

Proof : The fact $\sigma \text{pred}(p_1, \dots, p_n)$ is derivable from \mathcal{R}_0 if and only if $\sigma \text{pred}'(p'_1, \dots, p'_n)$ is derivable from \mathcal{R}'_0 , so by Theorem 5.2, if and only if $\sigma \text{pred}'(p'_1, \dots, p'_n)$ is derivable from $\text{saturate}(\mathcal{R}'_0)$, so if and only if there exists a clause $H \Rightarrow \text{pred}(p'_1, \dots, p'_n)$ in $\text{solve}_{\mathcal{R}_0}(\text{pred}(p_1, \dots, p_n))$ and a substitution σ' such that $\sigma' \text{pred}(p'_1, \dots, p'_n) = \sigma \text{pred}(p_1, \dots, p_n)$ and $\sigma' H$ is derivable from $\text{saturate}(\mathcal{R}'_0)$, that is, from \mathcal{R}'_0 . \square

In particular, if $\text{solve}_{\mathcal{R}_0}(\text{attacker}(p)) = \emptyset$, then $\text{attacker}(p)$ is not derivable from \mathcal{R}_0 (and if $\text{solve}_{\mathcal{R}_0}(\text{attacker}(p))$ is not empty for the selection function sel_0 , at least one instance of $\text{attacker}(p)$ is derivable, since H will contain facts of the form $\text{attacker}(x)$, an instance of which is derivable by $\text{attacker}(a[\])$).

5.4.2.2 Proofs

In this section, we detail the proof of Theorem 5.2. We first need to prove a few preliminary lemmas. The first one shows that two nodes in a derivation can be replaced by one when combining their clauses by resolution.

Lemma 5.1 (Resolution) *Consider a derivation containing a node η' , labeled R' . Let F_0 be a hypothesis of R' . Then there exists a son η of η' , labeled R , such that the edge $\eta' \rightarrow \eta$ is labeled by an instance of F_0 , $R \circ_{F_0} R'$ is defined, and one obtains a derivation of the same fact by replacing the nodes η and η' with a node η'' labeled $R'' = R \circ_{F_0} R'$.*

Proof : This proof is illustrated in Figure 5.9. Let $R' = H' \Rightarrow C'$, H'_1 be the multiset of the labels of the outgoing edges of η' , and C'_1 the label of its incoming edge. We have $R' \sqsupseteq (H'_1 \Rightarrow C'_1)$, so there exists a substitution σ such that $\sigma H' \subseteq H'_1$ and $\sigma C' = C'_1$. Since $F_0 \in H'$, $\sigma F_0 \in H'_1$, so there is an outgoing edge of η' labeled σF_0 . Let η be the node at the end of this edge, let $R = H \Rightarrow C$ be the label of η . We rename the variables of R so that they are distinct from the variables of R' . Let H_1 be the multiset of the labels of the outgoing edges of η . So $R \sqsupseteq (H_1 \Rightarrow \sigma F_0)$. By the above choice of distinct variables, we can then extend σ so that $\sigma H \subseteq H_1$ and $\sigma C = \sigma F_0$.

The edge $\eta' \rightarrow \eta$ is labeled σF_0 , instance of F_0 . Since $\sigma C = \sigma F_0$, the facts C and F_0 are unifiable, so $R \circ_{F_0} R'$ is defined. Let σ' be the most general unifier of C and F_0 , and σ'' such that




Figure 5.9: Merging of nodes of Lemma 5.1

$\sigma = \sigma''\sigma'$. We have $R \circ_{F_0} R' = \sigma'(H \cup (H' \setminus \{F_0\})) \Rightarrow \sigma'C'$. Moreover, $\sigma''\sigma'(H \cup (H' \setminus \{F_0\})) \subseteq H_1 \cup (H'_1 \setminus \{\sigma F_0\})$ and $\sigma''\sigma'C' = \sigma C' = C'_1$. Hence $R'' = R \circ_{F_0} R' \sqsupseteq (H_1 \cup (H'_1 \setminus \{\sigma F_0\})) \Rightarrow C'_1$. The multiset of labels of outgoing edges of η'' is precisely $H_1 \cup (H'_1 \setminus \{\sigma F_0\})$ and the label of its incoming edge is C'_1 , therefore we have obtained a correct derivation by replacing η and η' with η'' . \square

Lemma 5.2 (Subsumption) *If a node η of a derivation D is labeled by R , then one obtains a derivation D' of the same fact as D by relabeling η with a clause R' such that $R' \sqsupseteq R$.*

Proof : Let H be the multiset of labels of outgoing edges of the considered node η , and C be the label of its incoming edge. We have $R \sqsupseteq H \Rightarrow C$. By transitivity of \sqsupseteq , $R' \sqsupseteq H \Rightarrow C$. So we can relabel η with R' . \square

Lemma 5.3 (Saturation) *At the end of saturate, \mathcal{R} satisfies the following properties:*

1. For all $R \in \mathcal{R}_0$, R is subsumed by a clause in \mathcal{R} ;
2. Let $R \in \mathcal{R}$ and $R' \in \mathcal{R}$. Assume that $\text{sel}(R) = \emptyset$ and there exists $F_0 \in \text{sel}(R')$ such that $R \circ_{F_0} R'$ is defined. In this case, $R \circ_{F_0} R'$ is subsumed by a clause in \mathcal{R} .

Proof : To prove the first property, let $R \in \mathcal{R}_0$. We show that, after the addition of R to \mathcal{R} , R is subsumed by a clause in \mathcal{R} .

In the first step of **saturate**, we execute the instruction $\mathcal{R} \leftarrow \text{elim}(\{R\} \cup \mathcal{R})$. After execution of this instruction, R is subsumed by a clause in \mathcal{R} .

Assume that we execute $\mathcal{R} \leftarrow \text{elim}(\{R''\} \cup \mathcal{R})$ for some clause R'' and that, before this execution, R is subsumed by a clause in \mathcal{R} , say R' . If R' is removed by this instruction, there exists a clause R'_1 in \mathcal{R} that subsumes R' , so by transitivity of subsumption, R'_1 subsumes R , hence R is subsumed by the clause $R'_1 \in \mathcal{R}$ after this instruction. If R' is not removed by this instruction, then R is subsumed by the clause $R' \in \mathcal{R}$ after this instruction.

Hence, at the end of **saturate**, R is subsumed by a clause in \mathcal{R} , which proves the first property.

In order to prove the second property, we just need to notice that the fixpoint is reached at the end of **saturate**, so $\mathcal{R} = \text{elim}(\{R \circ_{F_0} R'\} \cup \mathcal{R})$. Hence, $R \circ_{F_0} R'$ is eliminated by **elim**, so it is subsumed by some clause in \mathcal{R} . \square

Proof of Theorem 5.2: Assume that F is derivable from \mathcal{R}_0 and consider a derivation of F from \mathcal{R}_0 . We show that F is derivable from **saturate**(\mathcal{R}_0).

We consider the value of the set of clauses \mathcal{R} at the end of **saturate**. For each clause R in \mathcal{R}_0 , R is subsumed by a clause in \mathcal{R} (Lemma 5.3, Property 1). So, by Lemma 5.2, we can replace all clauses R in the considered derivation with a clause in \mathcal{R} . Therefore, we obtain a derivation D of F from \mathcal{R} .

Next, we build a derivation of F from \mathcal{R}_1 , where $\mathcal{R}_1 = \text{saturate}(\mathcal{R}_0)$. If D contains a node labeled by a clause not in \mathcal{R}_1 , we can transform D as follows. Let η' be a lowest node of D labeled by a clause not in \mathcal{R}_1 . So all sons of η' are labeled by elements of \mathcal{R}_1 . Let R' be the clause labeling η' . Since $R' \notin \mathcal{R}_1$, $\text{sel}(R') \neq \emptyset$. Take $F_0 \in \text{sel}(R')$. By Lemma 5.1, there exists a son of η of η' labeled by R , such that $R \circ_{F_0} R'$ is defined, and we can replace η and η' with a node η'' labeled by $R \circ_{F_0} R'$. Since all sons of η' are labeled by elements of \mathcal{R}_1 , $R \in \mathcal{R}_1$. Hence $\text{sel}(R) = \emptyset$. So, by Lemma 5.3, Property 2, $R \circ_{F_0} R'$ is subsumed by a clause R'' in \mathcal{R} . By Lemma 5.2, we can relabel η'' with R'' . The total number of nodes strictly decreases since η and η' are replaced with a single node η'' .

So we obtain a derivation D' of F from \mathcal{R} , such that the total number of nodes strictly decreases. Hence, this replacement process terminates. Upon termination, all clauses are in \mathcal{R}_1 . So we obtain a derivation of F from \mathcal{R}_1 , which is the expected result.

For the converse implication, notice that, if a fact is derivable from \mathcal{R}_1 , then it is derivable from \mathcal{R} , and that all clauses added to \mathcal{R} do not create new derivable facts: if a fact is derivable by applying the clause $R \circ_{F_0} R'$, then it is also derivable by applying R and R' . \square

5.4.2.3 Optimizations

The resolution algorithm uses several optimizations, in order to speed up resolution. The first two are standard, while the last three are specific to protocols.

Elimination of duplicate hypotheses If a clause contains several times the same hypotheses, the duplicate hypotheses are removed, so that at most one occurrence of each hypothesis remains.

Elimination of tautologies If a clause has a conclusion that is already in the hypotheses, this clause is a tautology: it does not derive new facts. Such clauses are removed.

Elimination of hypotheses $\text{attacker}(x)$ If a clause $H \Rightarrow C$ contains in its hypotheses $\text{attacker}(x)$, where x is a variable that does not appear elsewhere in the clause, then the hypothesis $\text{attacker}(x)$ is removed. Indeed, the attacker always has at least one message, so $\text{attacker}(x)$ is always satisfied for some value of x .

Decomposition of data constructors A data constructor is a constructor f of arity n that comes with associated destructors g_i for $i \in \{1, \dots, n\}$ defined by $g_i(f(x_1, \dots, x_n)) \rightarrow x_i$.

Data constructors are typically used for representing data structures. Tuples are examples of data constructors. For each data constructor f , the following clauses are generated:

$$\text{attacker}(x_1) \wedge \dots \wedge \text{attacker}(x_n) \Rightarrow \text{attacker}(f(x_1, \dots, x_n)) \quad (\text{Rf})$$

$$\text{attacker}(f(x_1, \dots, x_n)) \Rightarrow \text{attacker}(x_i) \quad (\text{Rg})$$

Therefore, $\text{attacker}(f(p_1, \dots, p_n))$ is derivable if and only if $\forall i \in \{1, \dots, n\}$, $\text{attacker}(p_i)$ is derivable. When a fact of the form $\text{attacker}(f(p_1, \dots, p_n))$ is met, it is replaced with $\text{attacker}(p_1) \wedge \dots \wedge \text{attacker}(p_n)$. If this replacement is done in the conclusion of a clause $H \Rightarrow \text{attacker}(f(p_1, \dots, p_n))$, n clauses are created: $H \Rightarrow \text{attacker}(p_i)$ for each $i \in \{1, \dots, n\}$. This replacement is of course done recursively: if p_i itself is a data constructor application, it is replaced again. The clauses (Rf) and (Rg) for data constructors are left unchanged. (When $\text{attacker}(x)$ cannot be selected, the clauses (Rf) and (Rg) for data constructors are in fact not necessary, because they generate only tautologies during resolution. However, when $\text{attacker}(x)$ can be selected, which cannot be excluded with certain extensions, these clauses may become necessary for soundness.)

Secrecy assumptions When the user knows that a fact will not be derivable, he can tell it to the verifier. (When this fact is of the form $\text{attacker}(p)$, the user tells that p remains secret.) The tool then removes all clauses which have this fact in their hypotheses. At the end of the computation, the tool checks that the fact is indeed underivable from the obtained clauses. If the user has given erroneous information, an error message is displayed. Even in this case, the verifier never wrongly claims that a protocol is secure.

Mentioning such underivable facts prunes the search space, by removing useless clauses. This speeds up the resolution algorithm. In most cases, the secret keys of the principals cannot be known by the attacker. So, examples of underivable facts are $\text{attacker}(sk_A[])$, $\text{attacker}(sk_B[])$, \dots

For simplicity, the proofs given in Section 5.4.2.2 do not take into account these optimizations. For a full proof, we refer the reader to [?, Appendix C].

5.4.2.4 Termination

In general, the resolution algorithm may not terminate. (The derivability problem is undecidable.) In practice, however, it terminates in most examples.

Blanchet and Podelski have shown that it always terminates on a large and interesting class of protocols, the *tagged protocols* [?]. They consider protocols that use as cryptographic primitives only public-key encryption and signatures with atomic keys, shared-key encryption, message authentication codes, and hash functions. Basically, a protocol is tagged when each application of a cryptographic primitive is marked with a distinct constant tag. It is easy to transform a protocol into a tagged protocol by adding tags. For instance, our example of protocol can be transformed into a tagged protocol, by adding the tags c_0 , c_1 , c_2 to distinguish the encryptions and signature:

$$\begin{aligned} \text{Message 1. } & A \rightarrow B : \{(c_1, \{(c_0, k)\}_{sk_A})\}_{pk_B} \\ \text{Message 2. } & B \rightarrow A : \{(c_2, s)\}_k \end{aligned}$$

Adding tags preserves the expected behavior of the protocol, that is, the attack-free executions are unchanged. In the presence of attacks, the tagged protocol may be more secure. Hence,

tagging is a feature of good protocol design, as explained e.g. in [?]: the tags are checked when the messages are received; they facilitate the decoding of the received messages and prevent confusions between messages. More formally, tagging prevents type-flaw attacks [?], which occur when a message is taken for another message. However, the tagged protocol is potentially more secure than its untagged version, so, in other words, a proof of security for the tagged protocol does not imply the security of its untagged version. (Note that the tagging scheme considered here differs from the tagging schemes of Section 7.3.)

To illustrate the effect of tagging, we consider the Needham-Schroeder shared-key protocol. The algorithm does not terminate on its original version, which is untagged. It terminates after adding tags. In this protocol, we have two messages of the form:

$$\begin{aligned} \text{Message 4. } B \rightarrow A : & \{N_B\}_K \\ \text{Message 5. } A \rightarrow B : & \{N_B - 1\}_K \end{aligned}$$

where N_B is a nonce. Representing this with a function $\text{minusone}(x) = x - 1$, the algorithm does not terminate.

Indeed, message 5 is represented by a clause of the form:

$$H \wedge \text{attacker}(\text{senc}(n, k)) \Rightarrow \text{attacker}(\text{senc}(\text{minusone}(n), k))$$

where the hypothesis H describes other messages previously received by A . After some resolution steps, we obtain a clause of the form

$$\text{attacker}(\text{senc}(n, K)) \Rightarrow \text{attacker}(\text{senc}(\text{minusone}(n), K)) \quad (\text{Loop})$$

for some term K . The fact $\text{attacker}(\text{senc}(\text{minusone}(N_B), K))$ is also derived, so a resolution step with (Loop) yields: $\text{attacker}(\text{senc}(\text{minusone}(\text{minusone}(N_B)), K))$. This can again be resolved with (Loop), so that we finally have a cycle that derives $\text{attacker}(\text{senc}(\text{minusone}^n(N_B), K))$ for all n .

When tags are added, the rule (Loop) becomes:

$$\text{attacker}(\text{senc}((c_1, n), K)) \Rightarrow \text{attacker}(\text{senc}((c_2, \text{minusone}(n)), K)) \quad (\text{NoLoop})$$

and the previous loop is removed because c_2 does not unify with c_1 . The fact $\text{attacker}(\text{senc}((c_2, \text{minusone}(N_B)), K))$ is derived, but this does not yield a loop.

Other authors have proved related results: Ramanujan and Suresh [?] have shown that secrecy is decidable for tagged protocols. However, their tagging scheme is stronger since it forbids blind copies. A blind copy happens when a protocol participant sends back part of a message he received without looking at what is contained inside this part. On the other hand, they obtain a decidability result, while [?] obtains a termination result for an algorithm which is sound, efficient in practice, but approximate. Arapinis and Duflo [?] extend this result but still forbid blind copies. Comon-Lundh and Cortier [?] show that an algorithm using ordered binary resolution, ordered factorization and splitting terminates on protocols that blindly copy at most one term in each message. In contrast, the result of [?] puts no limit on the number of blind copies, but requires tagging.

For protocols that are not tagged, heuristics have been designed to adapt the selection function in order to obtain termination more often. We refer the reader to [?, Section 8.2] for more details.

It is also possible to obtain termination in all cases at the cost of additional abstractions. For instance, Goubault-Larrecq shows that one can abstract the clauses into clauses in the decidable class \mathcal{H}_1 [?], by losing some relational information on the messages.

5.5 Translation from the Pi Calculus

Given a closed process P_0 in the language of Section 5.3 and a set of names S , ProVerif builds a set of Horn clauses, representing the protocol in parallel with any S -adversary, in the same style as the clauses presented in the previous section. This translation was originally given in [?]. The clauses use *facts* defined by the following grammar:

$F ::=$	facts
$\text{attacker}(p)$	attacker knowledge
$\text{mess}(p, p')$	message on a channel

The fact $\text{attacker}(p)$ means that the attacker may have p , and the fact $\text{mess}(p, p')$ means that the message p' may appear on channel p . The clauses are of the form $F_1 \wedge \dots \wedge F_n \Rightarrow F$, where F_1, \dots, F_n, F are facts. They comprise clauses for the attacker and clauses for the protocol, defined below. These clauses form the set $\mathcal{R}_{P_0, S}$.

5.5.1 Clauses for the Attacker

The abilities of the attacker are represented by the following clauses:

For each $a \in S$, $\text{attacker}(a[])$	(Init)
$\text{attacker}(b_0[])$	(Rn)
For each public constructor f of arity n ,	
$\text{attacker}(x_1) \wedge \dots \wedge \text{attacker}(x_n) \Rightarrow \text{attacker}(f(x_1, \dots, x_n))$	(Rf)
For each destructor g ,	
for each rewrite rule $g(M_1, \dots, M_n) \rightarrow M$ in $\text{def}(g)$,	(Rg)
$\text{attacker}(M_1) \wedge \dots \wedge \text{attacker}(M_n) \Rightarrow \text{attacker}(M)$	
$\text{mess}(x, y) \wedge \text{attacker}(x) \Rightarrow \text{attacker}(y)$	(Rl)
$\text{attacker}(x) \wedge \text{attacker}(y) \Rightarrow \text{mess}(x, y)$	(Rs)

The clause (Init) represents the initial knowledge of the attacker. The clause (Rn) means that the attacker can generate new names. The clauses (Rf) and (Rg) mean that the attacker can apply all operations to all terms it has, (Rf) for constructors, (Rg) for destructors. For (Rg), notice that the rewrite rules in $\text{def}(g)$ do not contain names and that terms without names are also patterns, so the clauses have the required format. Clause (Rl) means that the attacker can listen on all channels it has, and (Rs) that it can send all messages it has on all channels it has.

If $c \in S$, we can replace all occurrences of $\text{mess}(c[], p)$ with $\text{attacker}(p)$ in the clauses. Indeed, these facts are equivalent by the clauses (Rl) and (Rs).

5.5.2 Clauses for the Protocol

When a function ρ associates a pattern with each name and variable, and f is a constructor, we extend ρ as a substitution by $\rho(f(M_1, \dots, M_n)) = f(\rho(M_1), \dots, \rho(M_n))$.

The translation $\llbracket P \rrbracket \rho H$ of a process P is a set of clauses, where ρ is a function that associates a pattern with each name and variable, and H is a sequence of facts of the form

$\text{mess}(p, p')$. The environment ρ maps each variable and name to its associated pattern representation. The sequence H keeps track of messages received by the process, since these may trigger other messages. The empty sequence is denoted by \emptyset ; the concatenation of a fact F to the sequence H is denoted by $H \wedge F$.

$$\begin{aligned}
\llbracket 0 \rrbracket \rho H &= \emptyset \\
\llbracket P \parallel Q \rrbracket \rho H &= \llbracket P \rrbracket \rho H \cup \llbracket Q \rrbracket \rho H \\
\llbracket !P \rrbracket \rho H &= \llbracket P \rrbracket \rho H \\
\llbracket \text{new } a.P \rrbracket \rho H &= \llbracket P \rrbracket (\rho[a \mapsto a[p'_1, \dots, p'_n]]) H \\
&\quad \text{where } H = \text{mess}(p_1, p'_1) \wedge \dots \wedge \text{mess}(p_n, p'_n) \\
\llbracket \text{in}(M, x).P \rrbracket \rho H &= \llbracket P \rrbracket (\rho[x \mapsto x])(H \wedge \text{mess}(\rho(M), x)) \\
\llbracket \text{out}(M, N).P \rrbracket \rho H &= \llbracket P \rrbracket \rho H \cup \{H \Rightarrow \text{mess}(\rho(M), \rho(N))\} \\
\llbracket \text{let } x = g(M_1, \dots, M_n) \text{ in } P \text{ else } Q \rrbracket \rho H &= \bigcup \{ \llbracket P \rrbracket ((\sigma\rho)[x \mapsto \sigma'p']) (\sigma H) \\
&\quad | g(p'_1, \dots, p'_n) \rightarrow p' \text{ is in } \text{def}(g) \text{ and } (\sigma, \sigma') \text{ is a most general pair of} \\
&\quad \text{substitutions such that } \sigma\rho(M_1) = \sigma'p'_1, \dots, \sigma\rho(M_n) = \sigma'p'_n \} \cup \llbracket Q \rrbracket \rho H \\
\llbracket \text{let } x = M \text{ in } P \rrbracket \rho H &= \llbracket P \rrbracket (\rho[x \mapsto \rho(M)]) H \\
\llbracket \text{if } M = N \text{ then } P \text{ else } Q \rrbracket \rho H &= \llbracket P \rrbracket (\sigma\rho)(\sigma H) \cup \llbracket Q \rrbracket \rho H \\
&\quad \text{where } \sigma \text{ is the most general unifier of } \rho(M) \text{ and } \rho(N)
\end{aligned}$$

The translation of a process is a set of Horn clauses that express that it may send certain messages.

- The nil process does nothing, so its translation is empty.
- The clauses for the parallel composition of processes P and Q are the union of clauses for P and Q .
- The replication is ignored, because all Horn clauses are applicable arbitrarily many times.
- For the restriction, we replace the restricted name a in question with the pattern $a[p'_1, \dots, p'_n]$, where p'_1, \dots, p'_n are the previous inputs.
- The sequence H is extended in the translation of an input, with the input in question.
- The translation of an output adds a clause, meaning that the output is triggered when all conditions in H are true.
- The translation of a destructor application is the union of the clauses for the cases where the destructor succeeds (with an appropriate substitution) and where the destructor fails. For simplicity, we assume that the `else` branch of destructors may always be executed; this is sufficient in most cases, since the `else` branch is often empty or just sends an error message. For a more precise treatment, see [?, Section 9.2].
- The conditional `if $M = N$ then P else Q` is in fact equivalent to `let $x = \text{equal}(M, N)$ in P else Q` , where the destructor `equal` is defined by `equal(x, x) $\rightarrow x$` , so the translation of the conditional is a particular case of the destructor application. We give it explicitly since it is particularly simple.

This translation of the protocol into Horn clauses introduces approximations. The actions are considered as implicitly replicated, since the clauses can be applied any number of times. This approximation implies that the tool fails to prove protocols that first need to keep some value secret and later reveal it. For instance, consider the process $\text{new } d.(\text{out}(d, s).\text{out}(c, d) \parallel \text{in}(d, x))$. This process preserves the secrecy of s , because s is output on the private channel d and received by the input on d , before the adversary gets to know d by the output of d on the public channel c . However, the Horn clause method cannot prove this property, because it treats this process like a variant with additional replications $\text{new } d.(!\text{out}(d, s).\text{out}(c, d) \parallel \text{in}(d, x))$, which does not preserve the secrecy of s .

5.5.2.1 Summary and Correctness

Let $\rho = \{a \mapsto a[] \mid a \in \text{fn}(P_0)\}$. We define the clauses corresponding to the process P_0 as:

$$\mathcal{R}_{P_0, S} = \llbracket P_0 \rrbracket \rho \emptyset \cup \{\text{attacker}(a[]) \mid a \in S\} \cup \{(\text{Rn}), (\text{Rf}), (\text{Rg}), (\text{Rl}), (\text{Rs})\}$$

Exercise 5.4 *Translate the process of Section 5.3.2 into clauses. Compare with the clauses given in Section 5.4.*

Exercise 5.5 *Translate the process of Exercise 5.2 into clauses.*

Theorem 5.3 (Correctness of the clauses) *Let P_0 be a closed process. Let M be a closed term and p be the pattern obtained from the term M by replacing all names a with $a[]$. If $\text{attacker}(p)$ is not derivable from $\mathcal{R}_{P_0, S}$, then P_0 preserves the secrecy of M from S .*

The proof of this result relies on a type system to express the soundness of the clauses on P_0 , and on the subject reduction of this type system to show that soundness of the clauses is preserved during all executions of the process. This technique was introduced in [?] where a similar result is proved. [?] also shows an equivalence between an instance of a generic type system for proving secrecy properties of protocols and the Horn clause verification method. This instance is the most precise instance of this generic type system.

By combining Theorem 5.3 with Corollary 5.1, we obtain:

Corollary 5.2 *Let P_0 be a closed process. Let M be a closed term and p be the pattern obtained from the term M by replacing all names a with $a[]$. If $\text{solve}_{\mathcal{R}_{P_0, S}}(\text{attacker}(p)) = \emptyset$, then P_0 preserves the secrecy of M from S .*

5.5.3 Extension to Equational Theories

ProVerif has been extended to handle primitives defined by equational theories [?]. The term algebra consists of constructors equipped with an equational theory, defined by a finite set of equations. For example, we can model a symmetric encryption scheme in which decryption always succeeds (but may return a meaningless message) by the equations

$$\begin{aligned} \text{sdec}(\text{senc}(x, y), y) &= x \\ \text{senc}(\text{sdec}(x, y), y) &= x \end{aligned} \tag{5.4}$$

where senc and sdec are constructors. The first equation is standard; the second one makes it possible to avoid that the equality test $\text{senc}(\text{sdec}(M, N), N) = M$ reveals that M is a ciphertext under N . These equations are satisfied by block ciphers, which are bijective.

We can also model the Diffie-Hellman key agreement [?] using equations. The Diffie-Hellman key agreement relies on the following property of modular exponentiation: $(g^a)^b = (g^b)^a = g^{ab}$ in a cyclic multiplicative subgroup G of \mathbb{Z}_p^* , where p is a large prime number and g is a generator of G , and on the assumption that it is difficult to compute g^{ab} from g^a and g^b , without knowing the random numbers a and b (computational Diffie-Hellman assumption), or on the stronger assumption that it is difficult to distinguish g^a, g^b, g^{ab} from g^a, g^b, g^c without knowing the random numbers a, b , and c (decisional Diffie-Hellman assumption). These properties are exploited to establish a shared key between two participants A and B of a protocol: A chooses randomly a and sends g^a to B ; symmetrically, B chooses randomly b and sends g^b to A . A can then compute $(g^b)^a$, since it has a and receives g^b , while B computes $(g^a)^b$. These two values being equal, they can be used to compute the shared key. The adversary, on the other hand, has g^a and g^b but not a and b so by the computational Diffie-Hellman assumption, it cannot compute the key. (This exchange resists passive attacks only; to resist active attacks, we need additional ingredients, for instance signatures.) We can model the Diffie-Hellman key agreement by the equation [?, ?]

$$\text{exp}(\text{exp}(\mathbf{g}, x), y) = \text{exp}(\text{exp}(\mathbf{g}, y), x) \quad (5.5)$$

where \mathbf{g} is a constant and exp is modular exponentiation. Obviously, this is a basic model: it models the main functional equation but misses many algebraic relations that exist in the group G .

The main idea of our extension to equations is to translate these equations into a set of rewrite rules associated to constructors. For instance, the equations (5.4) are translated into the rewrite rules

$$\begin{array}{ll} \text{senc}(x, y) \rightarrow \text{senc}(x, y) & \text{sdec}(x, y) \rightarrow \text{sdec}(x, y) \\ \text{senc}(\text{sdec}(x, y), y) \rightarrow x & \text{sdec}(\text{senc}(x, y), y) \rightarrow x \end{array} \quad (5.6)$$

while the equation (5.5) is translated into

$$\text{exp}(x, y) \rightarrow \text{exp}(x, y) \quad \text{exp}(\text{exp}(\mathbf{g}, x), y) \rightarrow \text{exp}(\text{exp}(\mathbf{g}, y), x) \quad (5.7)$$

Intuitively, these rewrite rules allow one, by applying them *exactly once* for each constructor, to obtain the various forms of the terms modulo the considered equational theory.² The constructors are then simply evaluated like destructors in the calculus above. One has formally defined when a set of rewrite rules models an equational theory, and designed algorithms that compute from the equations the rewrite rules that model the equational theory in question [?, Section 5]. Then, each trace in the calculus with equational theory corresponds to a trace in the calculus with rewrite rules, and conversely [?, Lemma 1].³ We are then reduced to the simpler case in which there are no equations. The main advantage of this technique is that resolution can still use ordinary syntactic unification (instead of having to use unification modulo the equational theory), and therefore remains efficient.

This extension to equations still has limitations: it does not allow us to model associative operations, such as exclusive or, since this would require an infinite number of rewrite rules.

²The rewrite rules like $\text{sdec}(x, y) \rightarrow \text{sdec}(x, y)$ are necessary so that sdec always succeeds. Thanks to this rule, the evaluation of $\text{sdec}(M, N)$ succeeds and leaves this term unchanged when M is not of the form $\text{senc}(M', N)$.

³More precisely, the inequality tests of (Red Destr 2) must still be performed modulo the equational theory, even in the calculus with rewrite rules.

It may be possible to handle these symbols using unification modulo the equational theory instead of syntactic unification, at the cost of a larger complexity. In the case of a bounded number of sessions, exclusive or is handled in [?, ?] and a more complete theory of modular exponentiation is handled in [?]. A unification algorithm for modular exponentiation is presented in [?]. For an unbounded number of sessions, extensions of the Horn clause approach that can handle XOR and Diffie-Hellman key agreements with more detailed algebraic relations (including equations of the multiplicative group modulo p) have been proposed by Küsters and Truderung: they handle XOR provided one of its two arguments is a constant in the clauses that model the protocol [?] and Diffie-Hellman key agreements provided the exponents are constants in the clauses that model the protocol [?]; they proceed by transforming the initial clauses into richer clauses on which the standard resolution algorithm is applied.

5.5.4 Extension to Scenarios with Several Phases

In some protocol studies, we consider scenarios in which a certain process is executed, then, in a second phase, this process stops and another process starts. For instance, when we model the compromise of long term keys, we consider that, in a first phase, the protocol runs normally, then, in a second phase, some keys are published. We aim at determining which secrets of the sessions of the protocol run in the first phase are preserved even though some keys are compromised. (This is the notion of forward secrecy.)

Such scenarios can be represented in ProVerif thanks to an extension of the syntax: the process **phase** $n.P$ represents a process P that runs in phase number n . The system first runs the processes in phase 0. Then, at some point during execution, it moves to phase 1. At this point, only the processes **phase** $n.P$ for $n \geq 1$ ready to run are kept (that is, the processes in phase 0 are stopped) and the processes **phase** 1. P are executed. Then, we move to phase 2, and so on.

This extension is translated into Horn clauses as follows. We consider predicates attacker_n and mess_n for each phase n , instead of the predicates attacker and mess . The clauses for the protocol use the predicate mess_n to translate the process P in **phase** $n.P$; the clauses for the adversary are repeated for each attacker_n . Moreover, the clauses

$$\text{attacker}_n(x) \Rightarrow \text{attacker}_{n+1}(x) \quad (\text{Rp})$$

for all n transmit the knowledge of the adversary from one phase to the next.

This extension was presented in [?, Section 8] and [?, Section 9.3]. An application of this extension will be mentioned in Section 5.7.1.

5.6 Extension to Correspondences

In this section, we extend the translation of the process calculus to Horn clauses, defined in the previous section for secrecy, to the proof of correspondences.

5.6.1 From Secrecy to Correspondences

In the analysis for secrecy, when $\text{attacker}(p)$ is derivable from the clauses the attacker *may* have p , that is, when $\text{attacker}(p)$ is not derivable from the clauses, we are sure that the attacker cannot have p , but the converse is not true, because the Horn clauses can be applied any number of times, which is not true in general for all actions of the process. Similarly,

when $\text{mess}(p, p')$ is derivable from the clauses, the message p' may be sent on channel p . Hence the analysis overapproximates the execution of actions.

Let us now consider that we want to prove a correspondence, for instance $\text{event}(e_1(x)) \rightsquigarrow \text{event}(e_2(x))$. In order to prove this correspondence, we can overapproximate the executions of event e_1 : if we prove the correspondence with this overapproximation, it will also hold in the exact semantics. So we can easily extend our analysis for secrecy with an additional predicate event , such that $\text{event}(p)$ means that $\text{event}(p)$ may have been executed. We generate clauses $\text{mess}(p_1, p'_1) \wedge \dots \wedge \text{mess}(p_n, p'_n) \Rightarrow \text{event}(p)$ when the process executes $\text{event}(p)$ after receiving p'_1, \dots, p'_n on channels p_1, \dots, p_n respectively. However, such an overapproximation cannot be done for the event e_2 : if we prove the correspondence after overapproximating the execution of e_2 , we are not really sure that e_2 will be executed, so the correspondence may be wrong in the exact semantics. Therefore, we have to use a different method for treating e_2 .

We use the following idea: we fix the exact set \mathcal{E} of allowed events $e_2(p)$ and, in order to prove $\text{event}(e_1(x)) \rightsquigarrow \text{event}(e_2(x))$, we check that only events $e_1(p)$ for p such that $e_2(p) \in \mathcal{E}$ can be executed. If we prove this property for any value of \mathcal{E} , we have proved the desired correspondence. So we introduce a predicate m-event , such that $\text{m-event}(e_2(p))$ is true if and only if $e_2(p) \in \mathcal{E}$. We generate clauses $\text{mess}(p_1, p'_1) \wedge \dots \wedge \text{mess}(p_n, p'_n) \wedge \text{m-event}(e_2(p_0)) \Rightarrow \text{mess}(p, p')$ when the process outputs p' on channel p after executing the event $e_2(p_0)$ and receiving p'_1, \dots, p'_n on channels p_1, \dots, p_n respectively. In other words, the output of p' on channel p can be executed only when $\text{m-event}(e_2(p_0))$ is true, that is, $e_2(p_0) \in \mathcal{E}$. (When the output of p' on channel p is under several events, the clause contains several m-event facts in its hypothesis. We also have similar clauses with $\text{event}(e_1(p))$ instead of $\text{mess}(p, p')$ when the event e_1 is executed after executing e_2 and receiving p'_1, \dots, p'_n on channels p_1, \dots, p_n respectively.)

For instance, if the events $e_2(p_1)$ and $e_2(p_2)$ are executed in a certain trace of the protocol, we define $\mathcal{E} = \{e_2(p_1), e_2(p_2)\}$, so that $\text{m-event}(e_2(p_1))$ and $\text{m-event}(e_2(p_2))$ are true and all other m-event facts are false. Then we show that the only events e_1 that may be executed are $e_1(p_1)$ and $e_1(p_2)$. We prove a similar result for all values of \mathcal{E} , which proves the desired correspondence.

In order to determine whether a fact is derivable from the clauses, we use a resolution-based algorithm. The resolution is performed for an unknown value of \mathcal{E} . So, basically, we keep m-event facts without trying to evaluate them (which we cannot do since \mathcal{E} is unknown). In the vocabulary of resolution, we never select m-event facts. Thus the obtained result holds for any value of \mathcal{E} , which allows us to prove correspondences. In order to prove the correspondence $\text{event}(e_1(x)) \rightsquigarrow \text{event}(e_2(x))$, we show that $\text{event}(e_1(p))$ is derivable only when $\text{m-event}(e_2(p))$ holds. We transform the initial set of clauses into a set of clauses that derives the same facts. If, in the obtained set of clauses, all clauses that conclude $\text{event}(e_1(p))$ contain $\text{m-event}(e_2(p))$ in their hypotheses, then $\text{event}(e_1(p))$ is derivable only when $\text{m-event}(e_2(p))$ holds, so the desired correspondence holds.

We still have to solve one problem: the reasoning above does not take into account that patterns p which occur in clauses differ from terms M , which represent messages. In patterns, we use a special encoding of names: a name a created by a restriction $\text{new } a$ is represented by a function $a[p_1, \dots, p_n]$ of the messages p_1, \dots, p_n received above the restriction, so that names created after receiving different messages are distinguished in the analysis (which is important for the precision of the analysis). However, this encoding still merges names created by the same restriction after receiving the same messages. For example, in the process $\text{!in}(c, x) \text{new } a$, the names created by $\text{new } a$ are represented by $a[x]$, so several names created for the same

value of x are merged. This merging is not acceptable for the verification of correspondences, because when we prove $\text{event}(e_1(x)) \rightsquigarrow \text{event}(e_2(x))$, we must make sure that x contains exactly the same names in $e_1(x)$ and in $e_2(x)$. In order to solve this problem, we label each replication with a *session identifier* i , which is an integer that takes a different value for each copy of the process generated by the replication. We add session identifiers as arguments to our encoding of names, which becomes $a[M_1, \dots, M_n, i_1, \dots, i_{n'}]$ where $i_1, \dots, i_{n'}$ are the session identifiers of the replications above the restriction $\text{new } a$. For example, in the process $!\text{in}(c, x) \text{new } a$, the names created by $\text{new } a$ are represented by $a[x, i]$. Each execution of the restriction is then associated with a distinct value of the session identifiers $i_1, \dots, i_{n'}$, so each name has a distinct encoding. We detail and formalize this encoding in Section 5.6.2.

5.6.2 Instrumented Processes

We consider a closed process P_0 representing the protocol we wish to verify. We assume that the bound names of P_0 have been renamed so that they are pairwise distinct and distinct from names in $S \cup \text{fn}(P_0)$ and in the correspondence to prove. We denote by Q a particular adversary; below, we prove the correspondence properties for any Q . Furthermore, we assume that, in the initial configuration $E_0, \{P_0, Q\}$, the names of E_0 not in $S \cup \text{fn}(P_0)$ or in the correspondence to prove have been renamed to fresh names, and the bound names of Q have been renamed so that they are pairwise distinct and fresh. (These renamings do not change the satisfied correspondences, since $\text{new } a.P$ and the renamed process $\text{new } a'.P\{a'/a\}$ reduce to the same configuration by (Red Res).)

In order to define formally the patterns associated with a name, we use a notion of instrumented processes. The syntax of instrumented processes is defined as follows:

- The replication $!P$ is labeled with a variable i in the set V_s of variable session identifiers (disjoint from the set V_o of ordinary variables): $!^i P$. The process $!^i P$ represents copies of P for a countable number of values of i . The variable i is a session identifier. It indicates which copy of P , that is, which session, is executed.
- The restriction $\text{new } a.P$ is labeled with a restriction label ℓ : $\text{new } a : \ell.P$, where ℓ is either $a[M_1, \dots, M_n, i_1, \dots, i_{n'}]$ for restrictions in honest processes or $b_0[a[i_1, \dots, i_{n'}]]$ for restrictions in the adversary. The symbol b_0 is a special name function symbol, distinct from all other such symbols. Using a specific instrumentation for the adversary is helpful so that all names generated by the adversary are encoded by instances of $b_0[x]$. They are therefore easy to generate. This labeling of restrictions is similar to a Church-style typing: ℓ can be considered as the type of a . (This type is polymorphic since it can contain variables.)

The instrumented processes are then generated by the following grammar:

$P, Q ::=$	instrumented processes
$!^i P$	replication
$\text{new } a : \ell.P$	restriction
\dots (as in the standard calculus)	

For instrumented processes, a semantic configuration S, E, \mathcal{P} consists of a set S of session identifiers that have not yet been used by \mathcal{P} , an environment E that is a mapping from names to closed patterns of the form $a[\dots]$, and a finite multiset of instrumented processes \mathcal{P} .

The first semantic configuration uses any countable set of session identifiers S_0 . The domain of E must always contain all free names of processes in \mathcal{P} , and the initial environment maps all names a to the pattern $a[]$. The semantic rules (Red Repl) and (Red Res) become:

$$S, E, \mathcal{P} \cup \{!^i P\} \rightarrow S \setminus \{\lambda\}, E, \mathcal{P} \cup \{P\{\lambda/i\}, !^i P\} \text{ where } \lambda \in S \quad (\text{Red Repl})$$

$$\begin{aligned} S, E, \mathcal{P} \cup \{\text{new } a : \ell.P\} \\ \rightarrow S, E[a' \mapsto E(\ell)], \mathcal{P} \cup \{P\{a'/a\}\} \text{ if } a' \notin \text{dom}(E) \end{aligned} \quad (\text{Red Res})$$

where the mapping E is extended to all terms as a substitution by $E(f(M_1, \dots, M_n)) = f(E(M_1), \dots, E(M_n))$ and to restriction labels by $E(a[M_1, \dots, M_n, i_1, \dots, i_{n'}]) = a[E(M_1), \dots, E(M_n), i_1, \dots, i_{n'}]$ and $E(b_0[a[i_1, \dots, i_{n'}]]) = b_0[a[i_1, \dots, i_{n'}]]$, so that it maps terms and restriction labels to patterns. The rule (Red Repl) takes an unused constant session identifier λ in S , and creates a copy of P with session identifier λ . The rule (Red Res) creates a fresh name a' , substitutes it for a in P , and adds to the environment E the mapping of a' to its encoding $E(\ell)$. Other semantic rules $E, \mathcal{P} \rightarrow E, \mathcal{P}'$ simply become $S, E, \mathcal{P} \rightarrow S, E, \mathcal{P}'$.

The instrumented process $P'_0 = \text{instr}(P_0)$ associated with the process P_0 is built from P_0 as follows:

- We label each replication $!P$ of P_0 with a distinct, fresh session identifier i , so that it becomes $!^i P$.
- We label each restriction $\text{new } a$ of P_0 with $a[t, s]$, so that it becomes $\text{new } a : a[t, s]$, where s is the sequence of session identifiers that label replications above $\text{new } a$ in the abstract syntax tree of P'_0 , in the order from top to bottom; t is the sequence of variables x that store received messages in inputs $\text{in}(M, x)$ above $\text{new } a$ in P_0 and results of non-deterministic destructor applications $\text{let } x = g(\dots) \text{ in } P \text{ else } Q$ above $\text{new } a$ in P_0 . (A destructor is said to be non-deterministic when it may return several different results for the same arguments. Adding the result of destructor applications to t is useful to improve precision, only for non-deterministic destructors. For deterministic destructors, the result of the destructor can be uniquely determined from the other elements of t , so the addition is useless. If we add the result of non-deterministic destructors to t , we can show that the relative completeness result of [?] still holds in the presence of non-deterministic destructors. This result shows that, for secrecy, the Horn clause approach is at least as precise as a large class of type systems.)

Hence names are represented by functions $a[t, s]$ of the inputs and results of destructor applications in t and the session identifiers in s . In each trace of the process, at most one name corresponds to a given $a[t, s]$, since different copies of the restriction have different values of session identifiers in s . Therefore, different names are not merged by the verifier.

For the adversary, we use a slightly different instrumentation. We build the instrumented process $Q' = \text{instrAdv}(Q)$ as follows:

- We label each replication $!P$ of Q with a distinct, fresh session identifier i , so that it becomes $!^i P$.
- We label each restriction $\text{new } a$ of Q with $b_0[a[s]]$, so that it becomes $\text{new } a : b_0[a[s]]$, where s is the sequence of session identifiers that label replications above $\text{new } a$ in Q' .

(Including the session identifiers as arguments of nonces is necessary for soundness, as discussed in Section 5.6.1. Including the messages previously received as arguments of nonces is important for precision in the case of honest processes, in order to relate the nonces to these messages. It is however useless for the adversary: since we consider any S -adversary Q , we have no definite information on the relation between nonces generated by the adversary and messages previously received by the adversary.)

Example 5.1 *The instrumentation of the example process of Section 5.3.4.2 is*

$$\begin{aligned}
P_A(sk_A, pk_A) &= !^{i_A} \text{in}(c, x_pk_B).\text{event}(e_A(x_pk_B)).\text{out}(c, pk_A).\text{in}(c, x_b). \\
&\quad \text{out}(c, \text{sign}((pk_A, x_pk_B, x_b), sk_A)) \\
P_B(sk_B, pk_B, pk_A) &= !^{i_B} \text{in}(c, x_pk_A).\text{new } b : b[x_pk_A, i_B].\text{out}(c, b).\text{in}(c, m). \\
&\quad \text{if } (x_pk_A, pk_B, b) = \text{check}(m, x_pk_A) \text{ then} \\
&\quad \text{if } x_pk_A = pk_A \text{ then event}(e_B(pk_B)) \\
P &= \text{new } sk_A : sk_A[].\text{new } sk_B : sk_B[].\text{let } pk_A = \text{pk}(sk_A) \text{ in let } pk_B = \text{pk}(sk_B) \text{ in} \\
&\quad \text{out}(c, pk_A).\text{out}(c, pk_B).(P_A(sk_A, pk_A) \parallel P_B(sk_B, pk_B, pk_A))
\end{aligned}$$

The replications are instrumented with session identifiers i_A and i_B respectively, and the restriction $\text{new } b$ is instrumented with the pattern $b[x_pk_A, i_B]$.

The semantics of instrumented processes allows exactly the same communications and events as the one of standard processes. More precisely, let \mathcal{P} be a multiset of instrumented processes. We define $\text{unInstr}(\mathcal{P})$ as the multiset of processes of \mathcal{P} without the instrumentation. We define that an instrumented trace executes an event $\text{event}(M)$ (at step τ) by naturally extending Definitions 5.4 and 5.7. Thus we have:

Proposition 5.1 *If $E_0, \{P_0, Q\} \rightarrow^* E_1, \mathcal{P}_1$, then there exist E'_1 and \mathcal{P}'_1 such that for any S , countable set of session identifiers, there exists S' such that $S, \{a \mapsto a[] \mid a \in E_0\}, \{\text{instr}(P_0), \text{instrAdv}(Q)\} \rightarrow^* S', E'_1, \mathcal{P}'_1$, $\text{dom}(E'_1) = E_1$, $\text{unInstr}(\mathcal{P}'_1) = \mathcal{P}_1$, and both traces execute the same events at the same steps.*

Conversely, if $S, \{a \mapsto a[] \mid a \in E_0\}, \{\text{instr}(P_0), \text{instrAdv}(Q)\} \rightarrow^ S', E'_1, \mathcal{P}'_1$, then $E_0, \{P_0, Q\} \rightarrow^* \text{dom}(E'_1), \text{unInstr}(\mathcal{P}'_1)$, and both traces execute the same events at the same steps.*

Proof : This is an easy proof by induction on the length of the traces. The reduction rules applied in both traces are rules with the same name. \square

We say that $\text{event}(p)$ is executed (at step τ) in the instrumented trace $\mathcal{T} = S_0, E_0, \mathcal{P}_0 \rightarrow^* S', E', \mathcal{P}'$ when there exists a term M such that $\text{event}(M)$ is executed (at step τ) in \mathcal{T} and $E'(M) = p$.

5.6.3 Generation of Horn Clauses and Resolution

Given a closed process P_0 and a set of names S , ProVerif first instruments P_0 to obtain $P'_0 = \text{instr}(P_0)$, then it builds a set of Horn clauses, representing the protocol in parallel with any S -adversary. The clauses use the facts:

$F ::=$	facts
$\text{attacker}(p)$	attacker knowledge
$\text{mess}(p, p')$	message on a channel
$\text{m-event}(p)$	must-event
$\text{event}(p)$	may-event

The predicates **attacker** and **mess** are as before. The fact **m-event**(p) means that **event**(M) must have been executed with M corresponding to p , and **event**(p) that **event**(M) may have been executed with M corresponding to p . The clauses are of the form $F_1 \wedge \dots \wedge F_n \Rightarrow F$, where F_1, \dots, F_n, F are facts. They comprise clauses for the attacker and clauses for the protocol, defined below. These clauses form the set $\mathcal{R}_{P'_0, S}$.

The clauses describing the attacker are almost the same as for the verification of secrecy in Section 5.5.1. The only difference is that, here, the attacker is given an infinite set of fresh names $b_0[x]$, instead of only one fresh name $b_0[]$. Indeed, we cannot merge all fresh names created by the attacker, since we have to make sure that different terms are represented by different patterns for the verification of correspondences to be correctly implemented, as seen in Section 5.6.1. So the clause (Rn) becomes **attacker**($b_0[x]$).

The clauses for the protocol are modified as follows:

$$\begin{aligned}
\llbracket !^i P \rrbracket \rho H &= \llbracket P \rrbracket (\rho[i \mapsto i]) H \\
\llbracket \text{new } a : a[M_1, \dots, M_n, i_1, \dots, i_{n'}]. P \rrbracket \rho H &= \\
&\llbracket P \rrbracket (\rho[a \mapsto a[\rho(M_1), \dots, \rho(M_n), \rho(i_1), \dots, \rho(i_{n'})]]) H \\
\llbracket \text{event}(M). P \rrbracket \rho H &= \llbracket P \rrbracket \rho (H \wedge \text{m-event}(\rho(M))) \cup \{H \Rightarrow \text{event}(\rho(M))\}
\end{aligned}$$

The other cases are the same as for secrecy.

- The replication only inserts the new session identifier i in the environment ρ . It is otherwise ignored, because all Horn clauses are applicable arbitrarily many times.
- For the restriction, we replace the restricted name a in question with the pattern $a[\rho(M_1), \dots, \rho(M_n), \rho(i_1), \dots, \rho(i_{n'})]$. By definition of the instrumentation, this pattern contains the previous inputs, results of non-deterministic destructor applications, and session identifiers.
- The translation of an event adds the hypothesis **m-event**($\rho(M)$) to H , meaning that P can be executed only if the event has been executed first. Furthermore, it adds a clause, meaning that the event is triggered when all conditions in H are true.

Remark 5.2 Depending on the form of the correspondences we want to prove, we can sometimes simplify the clauses generated for events. Suppose that all arguments of events in the process and in correspondences are of the form $f(M_1, \dots, M_n)$ for some function symbol f .

If, for a certain function symbol f , events **event**($f(\dots)$) occur only before \rightsquigarrow in the desired correspondences, then it is easy to see in the following theorems that hypotheses of the form **m-event**($f(\dots)$) in clauses can be removed without changing the result, so the clauses generated by the event **event**(M) when M is of the form $f(\dots)$ can be simplified into:

$$\llbracket \text{event}(M). P \rrbracket \rho H = \llbracket P \rrbracket \rho H \cup \{H \Rightarrow \text{event}(\rho(M))\}$$

(Intuitively, since the events $\text{event}(f(\dots))$ occur only before \rightsquigarrow in the desired correspondences, we never prove that an event $\text{event}(f(\dots))$ has been executed, so the facts $\text{m-event}(f(\dots))$ are useless.)

Similarly, if $\text{event}(f(\dots))$ occurs only after \rightsquigarrow in the desired correspondences, then clauses that conclude a fact of the form $\text{event}(f(\dots))$ can be removed without changing the result, so the clauses generated by the event $\text{event}(M)$ when M is of the form $f(\dots)$ can be simplified into:

$$\llbracket \text{event}(M).P \rrbracket \rho H = \llbracket P \rrbracket \rho (H \wedge \text{m-event}(\rho(M)))$$

(Intuitively, since the events $\text{event}(f(\dots))$ occur only after \rightsquigarrow in the desired correspondences, we never prove properties of the form “if $\text{event}(f(\dots))$ has been executed, then ...”, so clauses that conclude $\text{event}(f(\dots))$ are useless.)

Let $\rho = \{a \mapsto a[] \mid a \in \text{fn}(P'_0)\}$. We define the clauses corresponding to the instrumented process P'_0 as:

$$\mathcal{R}_{P'_0, S} = \llbracket P'_0 \rrbracket \rho \emptyset \cup \{\text{attacker}(a[]) \mid a \in S\} \cup \{(\text{Rn}), (\text{Rf}), (\text{Rg}), (\text{Rl}), (\text{Rs})\}$$

Exercise 5.6 Give the clauses for the process of Section 5.3.4.2.

Exercise 5.7 Model mutual authentication in the Needham-Schroeder public-key protocol (see Figure 2.1), by correspondences in a process. Give the corresponding Horn clauses.

Theorem 5.4 (Correctness of the clauses) Let P_0 be a closed process and Q be an S -adversary. Let $P'_0 = \text{instr}(P_0)$ and $Q' = \text{instrAdv}(Q)$. Consider a trace $\mathcal{T} = S_0, E_0, \{P'_0, Q'\} \rightarrow^* S', E', \mathcal{P}'$, with $\text{fn}(P'_0) \cup S \subseteq \text{dom}(E_0)$ and $E_0(a) = a[]$ for all $a \in \text{dom}(E_0)$. Assume that, if \mathcal{T} executes $\text{event}(p)$, then $\text{m-event}(p) \in \mathcal{F}_{\text{me}}$. Finally, assume that \mathcal{T} executes $\text{event}(p')$. Then $\text{event}(p')$ is derivable from $\mathcal{R}_{P'_0, S} \cup \mathcal{F}_{\text{me}}$.

This result shows that, if the only executed events are those allowed in \mathcal{F}_{me} and an event $\text{event}(p')$ is executed, then $\text{event}(p')$ is derivable from the clauses. It is proved in [?, Appendix B], using a technique similar to that for secrecy.

We use the same resolution algorithm as for secrecy, except that we never select facts $\text{m-event}(p)$. So the selection function becomes

$$\text{sel}_1(H \Rightarrow C) = \begin{cases} \emptyset & \text{if } \forall F \in H, F = \text{attacker}(x) \text{ for some variable } x \text{ or} \\ & F = \text{m-event}(p) \text{ for some pattern } p \\ \{F_0\} & \text{where } F_0 \in H \text{ and } F_0 \neq \text{attacker}(x) \text{ for any variable } x \text{ and} \\ & F_0 \neq \text{m-event}(p) \text{ for any pattern } p \end{cases} \quad (5.8)$$

By adapting the proof of Theorem 5.2, taking into account that we never select $\text{m-event}(p)$, we obtain:

Theorem 5.5 (Correctness of saturate) Let F be a closed fact. F is derivable from $\mathcal{R}_0 \cup \mathcal{F}_{\text{me}}$ if and only if F is derivable from $\text{saturate}(\mathcal{R}_0) \cup \mathcal{F}_{\text{me}}$.

Corollary 5.3 Let $\text{solve}_{\mathcal{R}_0}(\text{pred}(p_1, \dots, p_n)) = \{H \Rightarrow \text{pred}(p'_1, \dots, p'_n) \mid H \Rightarrow \text{pred}'(p'_1, \dots, p'_n) \in \text{saturate}(\mathcal{R}'_0)\}$, where pred' is a new predicate and $\mathcal{R}'_0 = \mathcal{R}_0 \cup \{\text{pred}(p_1, \dots, p_n) \Rightarrow \text{pred}'(p_1, \dots, p_n)\}$.

The fact $\sigma \text{pred}(p_1, \dots, p_n)$ is derivable from $\mathcal{R}_0 \cup \mathcal{F}_{\text{me}}$ if and only if there exists a clause $H \Rightarrow \text{pred}(p'_1, \dots, p'_n)$ in $\text{solve}_{\mathcal{R}_0}(\text{pred}(p_1, \dots, p_n))$ and a substitution σ' such that $\sigma' \text{pred}(p'_1, \dots, p'_n) = \sigma \text{pred}(p_1, \dots, p_n)$ and $\sigma' H$ is derivable from $\mathcal{R}'_0 \cup \mathcal{F}_{\text{me}}$.

Theorem 5.6 below states formally how we can interpret the result of resolution in terms of executed events. It is proved by combining Corollary 5.3 and Theorem 5.4.

Theorem 5.6 (Main theorem for events) *Let P_0 be a closed process and $P'_0 = \text{instr}(P_0)$. Let Q be an S -adversary and $Q' = \text{instrAdv}(Q)$.*

Consider a trace $\mathcal{T} = S_0, E_0, \{P'_0, Q'\} \rightarrow^ S', E', P'$, with $\text{fn}(P'_0) \cup S \subseteq \text{dom}(E_0)$ and $E_0(a) = a[]$ for all $a \in \text{dom}(E_0)$.*

If \mathcal{T} executes an instance $\text{event}(p')$ of $\text{event}(p)$, then there exist a clause $H \Rightarrow C \in \text{solve}_{\mathcal{R}_{P'_0, S}}(\text{event}(p))$ and a substitution σ such that $\text{event}(p') = \sigma C$ and, for all $\text{m-event}(p'')$ in σH , \mathcal{T} executes $\text{event}(p'')$.

Proof : Let $\mathcal{F}_{\text{me}} = \{\text{m-event}(p'') \mid \mathcal{T} \text{ executes } \text{event}(p'')\}$. By Theorem 5.4, since \mathcal{T} executes $\text{event}(p')$, $\text{event}(p')$ is derivable from $\mathcal{R}_{P'_0, S} \cup \mathcal{F}_{\text{me}}$. By Corollary 5.3, there exist a clause $R = H \Rightarrow C$ in $\text{solve}_{\mathcal{R}_{P'_0, S}}(\text{event}(p))$ and a substitution σ such that $\sigma C = \text{event}(p')$ and all elements of σH are derivable from $\mathcal{R}'_0 \cup \mathcal{F}_{\text{me}}$. For all $\text{m-event}(p)$ in σH , $\text{m-event}(p)$ is derivable from $\mathcal{R}'_0 \cup \mathcal{F}_{\text{me}}$. Since no clause in \mathcal{R}'_0 has a conclusion of the form $\text{m-event}(_)$, $\text{m-event}(p) \in \mathcal{F}_{\text{me}}$. Given the choice of \mathcal{F}_{me} , this means that \mathcal{T} executes $\text{event}(p)$. \square

5.6.4 Non-injective correspondences

Theorem 5.7 *Let P_0 be a closed process and $P'_0 = \text{instr}(P_0)$. Let M_k ($k \in \{1, \dots, l\}$) and M be terms. Let p_k, p be the patterns obtained by replacing names a with patterns $a[]$ in the terms M_k, M respectively. Assume that, for all clauses R in $\text{solve}_{\mathcal{R}_{P'_0, S}}(\text{event}(p))$, there exist σ' and H such that $R = H \wedge \text{m-event}(\sigma' p_1) \wedge \dots \wedge \text{m-event}(\sigma' p_l) \Rightarrow \sigma' \text{event}(p)$.*

Then P_0 satisfies the correspondence $\text{event}(M) \rightsquigarrow \bigwedge_{k=1}^l \text{event}(M_k)$ against S -adversaries.

Proof : Let Q be an S -adversary and $Q' = \text{instrAdv}(Q)$. Consider a trace $\mathcal{T} = E_0, \{P_0, Q\} \rightarrow^* E', P'$ with $\text{fn}(P'_0) \cup S \cup \text{fn}(M) \cup \bigcup_k \text{fn}(M_k) \subseteq E_0$ that executes $\sigma \text{event}(M)$ for some substitution σ . Let $\mathcal{T}' = S_0, \{a \mapsto a[] \mid a \in E_0\}, \{P'_0, Q'\} \rightarrow^* S'_1, E'_1, P'_1$ be the corresponding instrumented trace, which executes $\text{event}(E'_1(\sigma M))$, instance of $\text{event}(p)$. By Theorem 5.6, there exist $R = H' \Rightarrow C' \in \text{solve}_{\mathcal{R}_{P'_0, S}}(\text{event}(p))$ and σ'' such that $\text{event}(E'_1(\sigma M)) = \sigma'' C'$ and for all $\text{m-event}(p'')$ in $\sigma'' H'$, \mathcal{T}' executes $\text{event}(p'')$. All clauses R in $\text{solve}_{\mathcal{R}_{P'_0, S}}(\text{event}(p))$ are of the form $H \wedge \text{m-event}(\sigma' p_1) \wedge \dots \wedge \text{m-event}(\sigma' p_l) \Rightarrow \sigma' \text{event}(p)$ for some σ' . So, there exists σ' such that for all $k \in \{1, \dots, l\}$, $\text{m-event}(\sigma' p_k) \in H'$ and $C' = \sigma' \text{event}(p)$. Hence $\text{event}(E'_1(\sigma M)) = \sigma'' C' = \sigma'' \sigma' \text{event}(p)$ and for all $k \in \{1, \dots, l\}$, $\text{m-event}(\sigma'' \sigma' p_k) \in \sigma'' H'$, so \mathcal{T}' executes $\text{event}(\sigma'' \sigma' p_k)$, hence \mathcal{T} executes $\text{event}(M'_k)$ for some M'_k such that $E'_1(M'_k) = \sigma'' \sigma' p_k = \sigma'' \sigma' E'_1(M_k)$. We also have $E'_1(\sigma M) = \sigma'' \sigma' p = \sigma'' \sigma' E'_1(M)$. Since E'_1 defines a suitable relation between terms and patterns (more precisely, using [?, Lemma 9]), there exists a substitution σ_0 such that for all $k \in \{1, \dots, l\}$, $M'_k = \sigma_0 M_k$ and $\sigma M = \sigma_0 M$. So $\sigma M = \sigma_0 M$ and for all $k \in \{1, \dots, l\}$, \mathcal{T} executes $\text{event}(M'_k) = \text{event}(\sigma_0 M_k)$, so we have the result. \square

Example 5.2 *For the process P of Section 5.3.4.2, letting $P' = \text{instr}(P)$ and $S = \{c\}$, we have $\text{solve}_{\mathcal{R}_{P', S}}(\text{event}(e_B(x))) = \{\text{m-event}(e_A(\text{pk}(sk_B[]))) \Rightarrow \text{event}(e_B(\text{pk}(sk_B[])))\}$ so we have proved the correspondence $\text{event}(e_B(x)) \rightsquigarrow \text{event}(e_A(x))$.*

5.6.5 Sketch for injective correspondences

Let P_0 be a closed process and $P'_0 = \text{instr}'(P_0)$. We adapt the generation of clauses as follows: the set of clauses $\mathcal{R}'_{P'_0, S}$ is defined as $\mathcal{R}_{P'_0, S}$ except that

$$\begin{aligned} \llbracket \text{out}(M, N).P \rrbracket \rho H &= \llbracket P \rrbracket \rho H \cup \{H\{\rho|_{V_o \cup V_s} / \square\} \Rightarrow \text{mess}(\rho(M), \rho(N))\} \\ \llbracket !^i P \rrbracket \rho H &= \llbracket P \rrbracket (\rho[i \mapsto i])(H\{\rho|_{V_o \cup V_s} / \square\}) \\ \llbracket \text{event}(M, i).P \rrbracket \rho H &= \llbracket P \rrbracket \rho(H \wedge \text{m-event}(\rho(M), \square)) \cup \{H\{\rho|_{V_o \cup V_s} / \square\} \Rightarrow \text{event}(\rho(M), i)\} \end{aligned}$$

where \square is a special variable. The predicate **event** has as additional argument the session identifier in which the event is executed. The predicate **m-event** has as additional argument an environment ρ that contains the variables that have a single value for each execution of the considered event $\text{event}(M, i)$ and that are defined when we generate the clause. These variables are the variables that are defined above the first replication that follows $\text{event}(M, i)$ and above the output or event that generates the clause. The special variable \square is a placeholder for this environment ρ .

Proposition 5.2 (Injective correspondences) *Let P_0 be a closed process, $P'_0 = \text{instr}'(P_0)$. We assume that, in P_0 , all events are of the form $\text{event}(f(M_1, \dots, M_n))$ and that different occurrences of event have different root function symbols.*

Let M_k ($k \in \{1, \dots, l\}$) and M be terms. Let p_k, p be the patterns obtained by replacing names a with patterns $a[]$ in the terms M_k, M respectively.

We also assume that p is of the form $f(\dots)$ for some function symbol f and that, for all k , $p_k = f_k(\dots)$ for some function symbol f_k .

Let $\text{solve}_{\mathcal{R}'_{P'_0, S}}(\text{event}(p, i)) = \{R_r \mid r \in \{1, \dots, n\}\}$. Assume that there exist x_k, i_r , and ρ_{rk} ($r \in \{1, \dots, n\}, k \in \{1, \dots, l\}$) such that

- *For all $r \in \{1, \dots, n\}$, there exist H and σ such that $R_r = H \wedge \text{m-event}(\sigma p_1, \rho_{r1}) \wedge \dots \wedge \text{m-event}(\sigma p_l, \rho_{rl}) \Rightarrow \text{event}(\sigma p, i_r)$.*
- *For all r and r' in $\{1, \dots, n\}$, for all $k \in \{1, \dots, l\}$, $\rho_{rk}(x_k)\{\lambda/i_r\}$ does not unify with $\rho_{r'k}(x_k)\{\lambda'/i_{r'}\}$ when $\lambda \neq \lambda'$.*

Then P_0 satisfies the correspondence

$$\text{event}(M) \rightsquigarrow \bigwedge_{k=1}^l \text{inj event}(M_k)$$

against S -adversaries.

By Theorem 5.7, after deleting session identifiers and environments, the first item shows that P'_0 satisfies the correspondence

$$\text{event}(M) \rightsquigarrow \bigwedge_{k=1}^l \text{event}(M_k) \tag{5.9}$$

The environments and session identifiers as well as the second item serve in proving injectivity. Denote by $_$ an unknown term. If two instances of $\text{event}(p, i)$ are executed in P'_0 , by

the first item, they are instances of $\text{event}(\sigma_r p, i_r)$ for some r , so they are $\text{event}(\sigma'_1 \sigma_{r_1} p, \sigma'_1 i_{r_1})$ and $\text{event}(\sigma'_2 \sigma_{r_2} p, \sigma'_2 i_{r_2})$ for some σ'_1 and σ'_2 . Furthermore, there is only one occurrence of $\text{event}(f(\dots), i)$ in P'_0 , so the event $\text{event}(f(\dots), i)$ can be executed at most once for each value of the session identifier i , so $\sigma'_1 i_{r_1} \neq \sigma'_2 i_{r_2}$. Then, by the first item, corresponding events $\text{event}(\sigma'_1 \sigma_{r_1} p_k, -)$ and $\text{event}(\sigma'_2 \sigma_{r_2} p_k, -)$ have been executed, with associated environments $\sigma'_1 \rho_{r_1 k}$ and $\sigma'_2 \rho_{r_2 k}$. By the second item, $\rho_{r_1 k}(x_k)\{\lambda_1/i_{r_1}\}$ does not unify with $\rho_{r_2 k}(x_k)\{\lambda_2/i_{r_2}\}$ for different values $\lambda_1 = \sigma'_1 i_{r_1}$ and $\lambda_2 = \sigma'_2 i_{r_2}$ of the session identifier. (In this condition, r_1 can be equal to r_2 , and when $r_1 = r_2 = r$, the condition simply means that i_r occurs in ρ_{rk} .) So $\sigma'_1 \rho_{r_1 k}(x_k) \neq \sigma'_2 \rho_{r_2 k}(x_k)$, so the events $\text{event}(\sigma'_1 \sigma_{r_1} p_k, -)$ and $\text{event}(\sigma'_2 \sigma_{r_2} p_k, -)$ are distinct, which shows injectivity. This point is very similar to the fact that injective agreement is implied by non-injective agreement when the parameters of events contain nonces generated by the agent to whom authentication is being made, because the event can be executed at most once for each value of the nonce. (The session identifier i_r in our theorem plays the role of the nonce.) [Andrew Gordon, personal communication].

For the detailed proof of this result and more general results on correspondences, we refer to [?], in particular Section 7.2.

Example 5.3 For the process P of Section 5.3.4.2, letting $P' = \text{instr}(P)$ and $S = \{c\}$, we obtain $\text{solve}_{\mathcal{R}'_{P', S}}(\text{event}(e_B(x, i_B))) = \{\text{m-event}(e_A(\text{pk}(\text{sk}_B[])), \{x.b \mapsto b[\text{pk}(\text{sk}_A[]), i_B], x.pk_B \mapsto \text{pk}(\text{sk}_B[])\}) \Rightarrow \text{event}(e_B(\text{pk}(\text{sk}_B[]), i_B)\}$, so we have proved the correspondence $\text{event}(e_B(x)) \rightsquigarrow \text{inj event}(e_A(x))$ because i_B occurs in $\rho = \{x.b \mapsto b[\text{pk}(\text{sk}_A[]), i_B], x.pk_B \mapsto \text{pk}(\text{sk}_B[])\}$. (If distinct events e_B are executed, they have different values of i_B , so they correspond to different values of ρ , and to distinct events e_A .)

5.7 Extension to the Proof of Observational Equivalences

We now present the extension of ProVerif to the proof of observational equivalences, which was introduced in [?]. This section is adapted from that paper.

We focus on proving equivalences $P \approx Q$ in which P and Q are two processes that differ only in the choice of some terms. These equivalences arise often in applications. ProVerif proves these equivalences by proving a stronger equivalence that can be expressed as a property on the traces of a process that represents P and Q at the same time.

We first introduce a minor extension of the syntax of processes:

$D ::=$	term evaluations
M	term
$f(D_1, \dots, D_n)$	constructor application
$g(D_1, \dots, D_n)$	destructor evaluation
$P, Q ::=$	processes
\dots	(as in Figure 5.2)
$\text{let } x = D \text{ in } P \text{ else } Q$	term evaluation

This extension allows to evaluate several destructors in one step. This can help the proof of equivalences succeed more often: if a destructor g_1 fails in P but succeeds in Q , and a subsequent destructor g_2 fails in Q , the proof of equivalence will succeed if we evaluate g_1 and g_2 in one step (the term evaluation fails in both P and Q), but fail if we evaluate them separately, even if the processes P and Q are actually equivalent.

$$\begin{aligned}
& \text{out}(N, M).Q \parallel \text{in}(N', x).P \rightarrow Q \parallel P\{M/x\} && \text{(Red I/O)} \\
& \text{if } \text{fst}(N) = \text{fst}(N') \text{ and } \text{snd}(N) = \text{snd}(N') \\
& \text{let } x = D \text{ in } P \text{ else } Q \rightarrow P\{\text{diff}[M_1, M_2]/x\} && \text{(Red Fun 1)} \\
& \text{if } \text{fst}(D) \Downarrow M_1 \text{ and } \text{snd}(D) \Downarrow M_2 \\
& \text{let } x = D \text{ in } P \text{ else } Q \rightarrow Q && \text{(Red Fun 2)} \\
& \text{if there is no } M_1 \text{ such that } \text{fst}(D) \Downarrow M_1 \text{ and} \\
& \text{there is no } M_2 \text{ such that } \text{snd}(D) \Downarrow M_2
\end{aligned}$$

Figure 5.10: Generalized rules for biprocesses

The semantics is defined as in 5.3, except that (Red Destr 1) and (Red Destr 2) are replaced with (Red Fun 1) and (Red Fun 2) defined below. They use as auxiliary relation the term evaluation relation $D \Downarrow M$, which means that D evaluates to the term M , and is formally defined as follows:

$$\begin{aligned}
& M \Downarrow M \\
& f(D_1, \dots, D_n) \Downarrow f(N_1, \dots, N_n) \text{ if for all } i, D_i \Downarrow N_i \\
& g(D_1, \dots, D_n) \Downarrow \sigma N \text{ if } g(N_1, \dots, N_n) \rightarrow N \in \text{def}(g) \text{ and, for all } i, D_i \Downarrow \sigma N_i \\
& \text{let } x = D \text{ in } P \text{ else } Q \rightarrow P\{M/x\} \\
& \text{if } D \Downarrow M && \text{(Red Fun 1)} \\
& \text{let } x = D \text{ in } P \text{ else } Q \rightarrow Q \\
& \text{if there is no } M \text{ such that } D \Downarrow M && \text{(Red Fun 2)}
\end{aligned}$$

Next we introduce a new calculus that can represent pairs of processes that have the same structure and differ only by the terms and term evaluations that they contain. We call such a pair of processes a *biprocess*. The grammar for the calculus is a simple extension of the grammar of Figure 5.2, with additional cases so that $\text{diff}[M, M']$ is a term and $\text{diff}[D, D']$ is a term evaluation. We also extend the definition of contexts to permit the use of diff , and sometimes refer to contexts without diff as plain contexts.

Given a biprocess P , we define two processes $\text{fst}(P)$ and $\text{snd}(P)$, as follows: $\text{fst}(P)$ is obtained by replacing all occurrences of $\text{diff}[M, M']$ with M and $\text{diff}[D, D']$ with D in P , and similarly, $\text{snd}(P)$ is obtained by replacing $\text{diff}[M, M']$ with M' and $\text{diff}[D, D']$ with D' in P . We define $\text{fst}(D)$, $\text{fst}(M)$, $\text{snd}(D)$, and $\text{snd}(M)$ similarly. Our goal is to show that the processes $\text{fst}(P)$ and $\text{snd}(P)$ are observationally equivalent:

Definition 5.16 *Let P be a closed biprocess. We say that P satisfies observational equivalence when $\text{fst}(P) \approx \text{snd}(P)$.*

The semantics for biprocesses is defined as in Figure 5.3 with generalized rules (Red I/O), (Red Fun 1), and (Red Fun 2) given in Figure 5.10. Reductions for biprocesses bundle those for processes: if $P \rightarrow Q$ then $\text{fst}(P) \rightarrow \text{fst}(Q)$ and $\text{snd}(P) \rightarrow \text{snd}(Q)$. Conversely, however, reductions in $\text{fst}(P)$ and $\text{snd}(P)$ need not correspond to any biprocess reduction, in particular when they do not match up. Our first theorem shows that the processes are equivalent when this does not happen.

Definition 5.17 *We say that the biprocess P is uniform when $\text{fst}(P) \rightarrow Q_1$ implies that $P \rightarrow Q$ for some biprocess Q with $\text{fst}(Q) \equiv Q_1$, and symmetrically for $\text{snd}(P) \rightarrow Q_2$.*

Theorem 5.8 *Let P_0 be a closed biprocess. If, for all plain evaluation contexts C and reductions $C[P_0] \rightarrow^* P$, the biprocess P is uniform, then P_0 satisfies observational equivalence.*

Proof : Let P be a closed biprocess such that $C[P] \rightarrow^* \equiv Q$ always yields a uniform biprocess Q , and consider the relation

$$\mathcal{R} = \{(Q_1, Q_2) \mid \exists Q, C \text{ such that } Q_1 \equiv \text{fst}(Q), Q_2 \equiv \text{snd}(Q), \text{ and } C[P] \rightarrow^* \equiv Q\}$$

In particular, we have $\text{fst}(P) \mathcal{R} \text{snd}(P)$, so we can show that P satisfies observational equivalence by establishing that the relation $\mathcal{R}' = \mathcal{R} \cup \mathcal{R}^{-1}$ meets the three conditions of Definition 5.10. By symmetry, we focus on \mathcal{R} . Assume $Q_1 \mathcal{R} Q_2$. We have $Q_1 \equiv \text{fst}(Q)$, $Q_2 \equiv \text{snd}(Q)$, and $C[P] \rightarrow^* \equiv Q$ for some C and Q . We show that

1. if $P \downarrow_M$ then $Q \downarrow_M$;
2. if $P \rightarrow P'$ then $Q \rightarrow Q'$ and $P' \mathcal{R} Q'$ for some Q' ;
3. $C[P] \mathcal{R} C[Q]$ for all evaluation contexts C .

It is then easy to conclude the three conditions of Definition 5.10.

1. Assume $Q_1 \downarrow_M$, and let $T_M = \text{in}(M, x).\text{out}(c, c)$ for some fresh name c . As usual in the pi calculus, the predicate \downarrow_M tests the ability to send any message on M , hence for any plain process Q_i , we have $Q_i \downarrow_M$ if and only if $Q_i \parallel T_M \rightarrow R_i \parallel \text{out}(c, c)$ for some R_i .
Here, we have $Q_1 \parallel T_M \rightarrow R_1 \parallel \text{out}(c, c)$ for some R_1 . Since $Q_1 \equiv \text{fst}(Q)$, $\text{fst}(Q) \parallel T_M \rightarrow R_1 \parallel \text{out}(c, c)$. The reductions $C[P] \rightarrow^* \equiv Q$ imply $C[P] \parallel T_M \rightarrow^* \equiv Q \parallel T_M$. By hypothesis (with the context $C[_] \parallel T_M$), $Q \parallel T_M$ is uniform, hence $Q \parallel T_M \rightarrow Q' \parallel T_M$ for some Q' with $\text{fst}(Q') \equiv R_1 \parallel \text{out}(c, c)$. Since c does not occur anywhere in Q , by case analysis on this reduction step with our semantics for biprocesses we obtain $Q' \equiv R \parallel \text{out}(c, c)$ for some biprocess R . Thus, we obtain $\text{snd}(Q) \parallel T_M \rightarrow \text{snd}(R) \parallel \text{out}(c, c)$, so $Q_2 \parallel T_M \rightarrow \text{snd}(R) \parallel \text{out}(c, c)$, and finally $\text{snd}(Q) \downarrow_M$.
2. If $Q_1 \rightarrow Q'_1$, then $\text{fst}(Q) \rightarrow Q'_1$, so by uniformity, we have $Q \rightarrow Q'$ with $\text{fst}(Q') \equiv Q'_1$. Thus, $C[P] \rightarrow^* \equiv Q'$ and, by definition of \mathcal{R} , we obtain $Q'_1 \mathcal{R} \text{snd}(Q')$. Finally, by definition of the semantics of biprocesses, $Q \rightarrow Q'$ implies $\text{snd}(Q) \rightarrow \text{snd}(Q')$, so $Q_2 \rightarrow \text{snd}(Q')$.
3. Let C' be a plain evaluation context. By definition of the semantics of biprocesses, $C[P] \rightarrow^* \equiv Q$ always implies $C'[C[P]] \rightarrow^* \equiv C'[Q]$. Moreover, $C'[Q_1] \equiv C'[\text{fst}(Q)] = \text{fst}(C'[Q])$ and $C'[Q_2] \equiv C'[\text{snd}(Q)] = \text{snd}(C'[Q])$, hence $C'[Q_1] \mathcal{R} C'[Q_2]$. \square

Our plan is to establish the hypothesis of Theorem 5.8 by automatically verifying that all the biprocesses P in question meet conditions that imply uniformity. The next corollary details those conditions, which guarantee that a communication and an evaluation, respectively, succeed in $\text{fst}(P)$ if and only if they succeed in $\text{snd}(P)$:

Corollary 5.4 *Let P_0 be a closed biprocess. Suppose that, for all plain evaluation contexts C , all evaluation contexts C' , and all reductions $C[P_0] \rightarrow^* P$,*

1. if $P \equiv C'[\text{out}(N, M).Q \parallel \text{in}(N', x).R]$, then $\text{fst}(N) = \text{fst}(N')$ if and only if $\text{snd}(N) = \text{snd}(N')$,
2. if $P \equiv C'[\text{let } x = D \text{ in } Q \text{ else } R]$, then there exists M_1 such that $\text{fst}(D) \Downarrow M_1$ if and only if there exists M_2 such that $\text{snd}(D) \Downarrow M_2$.

Then P_0 satisfies observational equivalence.

Proof : We show that P is uniform, then we conclude by Theorem 5.8. Let us show that, if $\text{fst}(P) \rightarrow P'_1$ then there exists a biprocess P' such that $P \rightarrow P'$ and $\text{fst}(P') \equiv P'_1$. The case for $\text{snd}(P) \rightarrow P'_2$ is symmetric.

By induction on the derivation of $\text{fst}(P) \rightarrow P'_1$, we first show that there exist C , Q , and Q'_1 such that $P \equiv C[Q]$, $P'_1 \equiv \text{fst}(C)[Q'_1]$, and $\text{fst}(Q) \rightarrow Q'_1$ using one of the four process rules (Red I/O), (Red Fun 1), (Red Fun 2), or (Red Repl): every step in this derivation trivially commutes with fst , except for structural steps that involve a parallel composition and a restriction, in case $a \in \text{fn}(P)$ but $a \notin \text{fn}(\text{fst}(P))$. In that case, we use a preliminary renaming from a to some fresh $a' \notin \text{fn}(P)$.

For each of these four rules, relying on a hypothesis of Corollary 5.4, we find Q' such that $\text{fst}(Q') = Q'_1$ and $Q \rightarrow Q'$ using the corresponding biprocess rule:

(Red I/O): We have $Q = \text{out}(N, M).R \parallel \text{in}(N', x).R'$ with $\text{fst}(N) = \text{fst}(N')$ and $Q'_1 = \text{fst}(R) \parallel \text{fst}(R')\{\text{fst}(M)/x\}$. For $Q' = R \parallel R'\{M/x\}$, we have $\text{fst}(Q') = Q'_1$ and, by hypothesis 1, $\text{snd}(N) = \text{snd}(N')$, hence $Q \rightarrow Q'$.

(Red Fun 1): We have $Q = \text{let } x = D \text{ in } R \text{ else } R'$ with $\text{fst}(D) \Downarrow M_1$ and $Q'_1 = \text{fst}(R)\{M_1/x\}$. By hypothesis 2, $\text{snd}(D) \Downarrow M_2$ for some M_2 . We take $Q' = R\{\text{diff}[M_1, M_2]\}$, so that $\text{fst}(Q') = Q'_1$ and $Q \rightarrow Q'$.

(Red Fun 2): We have $Q = \text{let } x = D \text{ in } R \text{ else } R'$ with no M_1 such that $\text{fst}(D) \Downarrow M_1$ and $Q'_1 = \text{fst}(R')$. By hypothesis 2, there is no M_2 such that $\text{snd}(D) \Downarrow M_2$. We obtain $Q \rightarrow Q'$ for $Q' = R'$.

(Red Repl): We have $Q = !R$ and $Q'_1 = \text{fst}(R) \parallel !\text{fst}(R)$. We take $Q' = R \parallel !R$, so that $\text{fst}(Q') = Q'_1$ and $Q \rightarrow Q'$.

To conclude, we take the biprocess $P' = C[Q']$ and the reduction $P \rightarrow P'$. □

Thus, we have a sufficient condition for observational equivalence of biprocesses. This condition is essentially a reachability condition on biprocesses. We adapt existing techniques for reasoning about processes in order to prove this condition. The condition is however not necessary: if $P \approx Q$, then if $\text{diff}[\text{true}, \text{false}] = \text{true}$ then $P \text{ else } Q$ satisfies observational equivalence, but Theorem 5.8 and Corollary 5.4 will not enable us to prove this fact.

Example 5.4 *Probabilistic encryption is a variant of public-key encryption that further protects the secrecy of the plaintext by embedding some additional, fresh value in each encryption. It can be modeled using three functions for public-key decryption, public-key encryption, and public-key derivation, linked by the equation*

$$\text{adec}(\text{aenc}(x, \text{pk}(y), z), y) = x$$

where z is the additional random parameter for the encryption. A key property of probabilistic encryption is that, without knowledge of the decryption key, ciphertexts appear to be unrelated to the plaintexts, even if the attacker knows the plaintexts and the encryption key. A strong version of this property is that the ciphertexts cannot be distinguished from freshly generated random values. Formally, we state that

$$\text{new } s.(\text{out}(c, \text{pk}(s)) \parallel !\text{in}(c', x).\text{new } a.\text{out}(c, \text{diff}[\text{aenc}(x, \text{pk}(s), a), a]))$$

satisfies equivalence. This biprocess performs a first output to reveal the public key $\text{pk}(s)$ (but not $s!$), then repeatedly inputs a term x from the environment and either outputs its encryption under $\text{pk}(s)$ or outputs a fresh, unrelated name. Thus, a single biprocess represents the family of static equivalences that relate a series of probabilistic encryptions for any series of plaintext to a series of fresh, independent names. (Formally, each such equivalence can be obtained as a corollary of this biprocess equivalence, by applying the congruence property of equivalence for the particular context that sends the plaintexts of values on channel c' and reads the encryption key and encryptions on channel c .)

As usual, we can prove the reachability properties that occur as assumptions of Corollary 5.4 by translating the process into Horn clauses. Clauses are built from the following predicates:

$F ::=$	facts
$\text{attacker}'(p, p')$	attacker knowledge
$\text{mess}'(p_1, p_2, p'_1, p'_2)$	output message p_2 on p_1 (resp. p'_2 on p'_1)
$\text{input}'(p, p')$	input on p (resp. p')
$\text{nounif}(p, p')$	impossible unification
bad	bad

Informally, $\text{attacker}'(p, p')$ means that the attacker may obtain p in $\text{fst}(P)$ and p' in $\text{snd}(P)$ by the same operations; $\text{mess}'(p_1, p_2, p'_1, p'_2)$ means that message p_2 may appear on channel p_1 in $\text{fst}(P)$ and that message p'_2 may appear on channel p'_1 in $\text{snd}(P)$ after the same reductions; $\text{input}'(p, p')$ means that an input may be executed on channel p in $\text{fst}(P)$ and on channel p' in $\text{snd}(P)$, thus enabling the attacker to infer whether p (resp. p') is equal to another channel used for output; $\text{nounif}(p, p')$ means that p and p' cannot be unified by substituting elements of $GVar$ with patterns; finally, bad serves in detecting violations of observational equivalence: when bad is not derivable, we have observational equivalence.

An evident difference with respect to previous translations from processes to clauses is that predicates have twice as many arguments: we use the binary predicate $\text{attacker}'$ instead of the unary one attacker and the 4-ary predicate mess' instead of the binary one mess . This extension allows us to represent information for both variants of a biprocess.

The predicate nounif is not defined by clauses, but by special simplification steps in the solver. We omit the details of the generation of the clauses and of the resolution algorithm, which can be found in [?].

Example 5.5 *The biprocess of Example 5.4 yields the clauses:*

$$\begin{aligned} &\text{mess}'(c, \text{pk}(s), c, \text{pk}(s)) \\ &\text{mess}'(c', x, c', x') \Rightarrow \text{mess}'(c, \text{aenc}(x, \text{pk}(s), a[i, x]), c, a[i, x']) \end{aligned}$$

The first clause corresponds to the output of the public key $\text{pk}(s)$. The second clause corresponds to the other output: if a message x (resp. x') is received on channel c' , then the message $\text{aenc}(x, \text{pk}(s), a[i, x])$ in the first variant (resp. $a[i, x']$ in the second variant) is sent on channel c . The fresh name a is encoded as a pattern $a[i, x]$.

In a similar way as for previous security properties, we can prove:

Theorem 5.9 *If bad is not a logical consequence of the clauses, then P_0 satisfies observational equivalence.*

In order to determine whether bad is a logical consequence of the clauses, we use again an algorithm based on resolution with free selection, adapting the algorithm presented previously, in particular with specific simplification steps for nounif .

5.7.1 Weak Secrets

A *weak secret* represents a secret value with low entropy, such as a human-memorizable password. Protocols that rely on weak secrets are often subject to guessing attacks, whereby an attacker guesses a weak secret, perhaps using a dictionary, and verifies its guess. The guess verification may rely on interaction with protocol participants or on computations on intercepted messages (e.g., $[?, ?, ?]$). With some care in protocol design, however, those attacks can be prevented:

- On-line guessing attacks can be mitigated by limiting the number of retries that participants allow. An attacker that repeatedly attempts to guess the weak secret should be eventually detected and stopped if it tries to verify its guesses by interacting with other participants.
- Off-line guessing attacks can be prevented by making sure that, even if the attacker (systematically) guesses the weak secret, it cannot verify whether its guess is correct by computing on intercepted traffic.

Off-line guessing attacks can be explained and modeled in terms of a 2-phase scenario. In phase 0, on-line attacks are possible, but the weak secret is otherwise unguessable. In phase 1, the attacker obtains a possible value for the weak secret (intuitively, by guessing it). The absence of off-line attacks is characterized by an equivalence: the attacker cannot distinguish the weak secret used in phase 0 from an unrelated fresh value.

In our calculus, we arrive at the following definition:

Definition 5.18 (Weak secrecy) *Let P be a closed process with no phase prefix. We say that P prevents off-line attacks against w when $\text{new } w.(\text{phase } 0.P \parallel \text{phase } 1.\text{new } w'.\text{out}(c, \text{diff}[w, w'])))$ satisfies observational equivalence.*

This definition is in line with the work of Cohen, Corin et al., Delaune and Jacquemard, Drielsma et al., and Lowe $[?, ?, ?, ?, ?, ?]$. Lowe uses the model-checker FDR to handle a bounded number of sessions, while Delaune and Jacquemard give a decision procedure in this case. Corin et al. give a definition based on equivalence like ours, but do not consider the first, active phase; they analyze only one session.

As a first example, assume that a principal attempts to prove knowledge of a shared password w to a trusted server by sending a hash of this password encrypted under the

server's public key. (For simplicity, the protocol does not aim to provide freshness guarantees, so anyone may replay this proof.) Omitting the code for the server, a first protocol may be written:

$$P = \text{new } s.\text{out}(c, pk(s)).\text{out}(c, \text{aenc}(h(w), pk(s)))$$

The first output reveals the public key of the server; the second output communicates the proof of knowledge of w . This protocol does not prevent off-line attacks against w . ProVerif finds an attack that corresponds to the following adversary:

$$\begin{aligned} A = & \text{phase } 0.\text{in}(c, pk).\text{in}(c, e). \\ & \text{phase } 1.\text{in}(c, w).\text{if } e = \text{aenc}(h(w), pk) \text{ then out}(\text{Guessed}, ()) \end{aligned}$$

A corrected protocol uses probabilistic encryption (see Example 5.4):

$$P = \text{new } s, a.\text{out}(c, pk(s)).\text{out}(c, \text{aenc}(h(w), pk(s), a))$$

ProVerif automatically produces a proof for this corrected protocol.

5.7.2 Authenticity

Abadi and Gordon [?] use equivalences for characterizing authenticity properties, and treat a variant of the Wide-Mouth-Frog protocol as an example. Essentially, authenticity is defined as an equivalence between the protocol and a specification. The technique presented here automatically proves authenticity for the one-session version of this protocol [?, Section 3.2.2], thereby eliminating the need for a laborious manual proof. (Authenticity properties are often formulated as correspondence assertions on behaviors, rather than as equivalences. Those assertions can also be verified with ProVerif, as shown in Section 5.6.)

5.8 Applications

The automatic protocol verifier ProVerif is available at <http://www.proverif.ens.fr/>. It was successfully applied to many protocols of the literature, to prove secrecy and authentication properties: flawed and corrected versions of the Needham-Schroeder public-key [?, ?] and shared-key [?, ?, ?], Woo-Lam public-key [?, ?] and shared-key [?, ?, ?, ?, ?], Denning-Sacco [?, ?], Yahalom [?], Otway-Rees [?, ?, ?], and Skeme [?] protocols. No false attack occurred in these tests and the only non-termination cases were some flawed versions of the Woo-Lam shared-key protocol. The other protocols were verified in less than one second each on a Pentium M 1.8 GHz [?].

Moreover, ProVerif was also used in more substantial case studies:

- Abadi and Blanchet [?] applied it to the verification of a certified email protocol [?]. They use correspondence properties to prove that the receiver receives the message if and only if the sender has a receipt for the message. (They use simple manual arguments to take into account that the reception of sent messages is guaranteed.) One of the tested versions includes the SSH transport layer in order to establish a secure channel. (Total runtime: 6 min on a Pentium M 1.8 GHz.)
- Abadi, Blanchet, and Fournet [?] studied the JFK protocol (*Just Fast Keying*) [?], which was one of the candidates to the replacement of IKE as key exchange protocol in IPSec. They combined manual proofs and ProVerif to prove correspondences and equivalences. (Total runtime: 3 min on a Pentium M 1.8 GHz.)

- Blanchet and Chaudhuri [?] studied the secure filesystem Plutus [?] with ProVerif, which allowed them to discover and fix weaknesses of the initial system.
- Bhargavan et al. [?, ?, ?] use it to build the Web services verification tool TulaFale: Web services are protocols that send XML messages; TulaFale translates them into the input format of ProVerif and uses ProVerif to prove the desired security properties.
- Bhargavan et al. [?, ?, ?] use ProVerif for verifying implementations of protocols in F# (a functional language of the Microsoft .NET environment): a subset of F# large enough for expressing security protocols is translated into the input format of ProVerif. The TLS protocol, in particular, was studied using this technique [?].
- Canetti and Herzog [?] use ProVerif for verifying protocols in the computational model: they show that, for a restricted class of protocols that use only public-key encryption, a proof in the Dolev-Yao model implies security in the computational model, in the universal composability framework. Authentication is verified using correspondences, while secrecy of keys corresponds to strong secrecy.
- ProVerif was also used for verifying a certified email web service [?], a certified mailing-list protocol [?], e-voting protocols [?, ?], the ad-hoc routing protocol ARAN (*Authenticated Routing for Adhoc Networks*) [?], and zero-knowledge protocols [?].

Finally, Goubault-Larrecq and Parrennes [?] also use the Horn clause method for analyzing implementations of protocols written in C. However, they translate protocols into clauses of the \mathcal{H}_1 class and use the \mathcal{H}_1 prover by Goubault-Larrecq [?] rather than ProVerif to prove secrecy properties of the protocol.

5.9 Further Readings

The verifier ProVerif has been presented in several research papers. The translation from the spi calculus to Horn clauses for secrecy and its relation with typing was presented in [?]. Its extension to correspondence assertions was presented in [?]. Its extension to the proof of equivalences, the treatment of equational theories, and the extension to scenarios with several phases (named stages in [?]) was presented in [?]. ProVerif can also reconstruct attacks against protocols [?], although we did not detail this point.

5.10 Exercises

Exercise 5.8 *We assume the formal model of cryptography (the so-called Dolev-Yao model) and we consider the following protocol by Otway and Rees (1987).*

- | | |
|--|----------------|
| Message 1. $A \rightarrow B : M, A, B, \{N_a, M, A, B\}_{K_{as}}$ | M, N_a fresh |
| Message 2. $B \rightarrow S : M, A, B, \{N_a, M, A, B\}_{K_{as}}, \{N_b, M, A, B\}_{K_{bs}}$ | N_b fresh |
| Message 3. $S \rightarrow B : M, \{N_a, K_{ab}\}_{K_{as}}, \{N_b, K_{ab}\}_{K_{bs}}$ | K_{ab} fresh |
| Message 4. $B \rightarrow A : M, \{N_a, K_{ab}\}_{K_{as}}$ | |
| Message 5. $A \rightarrow B : \{s\}_{K_{ab}}$ | |

We recall that the notation above describes the messages exchanged in a correct execution of the protocol. Three participants are involved in this protocol: A and B aim at establishing a session, using an authentication server S . In the protocol, A and B are the identities of A and B ; K_{as} is a long-term key shared between A and S ; similarly, K_{bs} is a key shared between B and S ; M , N_a , and N_b are nonces; K_{ab} is the session key between A and B , established by the protocol (it is chosen by the server S); s is a message that A wants to send to B , with the guarantee that the adversary will not be able to obtain it. The notation $\{M\}_K$ represents the shared-key encryption of the message M under the key K .

One shall consider that A is willing to execute the protocol with B , but also with other participants, which are included in the adversary. However, Message 5 is sent by A only if its interlocutor is B . (This message is used in order to test that the key K_{ab} shared between two honest participants A and B is really secret.) B is also willing to execute the protocol not only with A , but also with other participants. The server S is willing to interact with any participants. The server S has a table that relates the identity of each participant X (A , B , ...) to the secret key K_{xs} shared between X and S .

1. Explain the actions of the three participants of the protocol: one shall explain how each participant verifies the messages that he receives and how he constructs the messages that he sends.
2. Code this protocol in the variant of the spi calculus presented during the course. (One shall give both the process and the definition of the destructors used in the protocol.)

Indication: In order to represent the table that S has and that relates the identity of each participant X with the secret key K_{xs} shared between X and S , one may consider that the identity X is coded by the term $\text{host}(K_{xs})$, and that the server uses a function getkey such that $\text{getkey}(X) = \text{getkey}(\text{host}(K_{xs})) = K_{xs}$ in order to reconstruct the key K_{xs} from the identity X .

3. Give the Horn clauses associated to this protocol, in order to show that s is secret. Give a succinct intuitive explanation of each of these clauses (or groups of clauses).
(Several variants are possible.)

Exercise 5.9 Similar exercise with the following protocol by Carlsen:

Message 1. $A \rightarrow B : A, N_a$	N_a fresh
Message 2. $B \rightarrow S : A, N_a, B, N_b$	N_b fresh
Message 3. $S \rightarrow B : \{K_{ab}, N_a, A\}_{K_{bs}}, \{N_a, B, K_{ab}\}_{K_{as}}$	K_{ab} fresh
Message 4. $B \rightarrow A : \{N_a, B, K_{ab}\}_{K_{as}}, \{N_a\}_{K_{ab}}, N'_b$	N'_b fresh
Message 5. $A \rightarrow B : \{N'_b\}_{K_{ab}}$	
Message 6. $A \rightarrow B : \{s\}_{K_{ab}}$	

In this protocol, A and B are the identities of A and B ; K_{as} is a long-term key shared between A and S ; similarly, K_{bs} is a key shared between B and S ; N_a , N_b , and N'_b are nonces; K_{ab} is the session key between A and B , established by the protocol (it is chosen by the server S); s is a message that A wants to send to B , with the guarantee that the adversary will not be able to obtain it. The notation $\{M\}_K$ represents the shared-key encryption of the message M under the key K .

Message 6 is sent by A only if its interlocutor is B. (This message is used in order to test that the key K_{ab} shared between two honest participants A and B is really secret.)

Chapter 6

Equivalence based properties (static equivalence)

6.1 Definitions and Applications

In the first part we have seen how to verify trace properties such as secrecy, formulated as non-deducibility, and authentication properties in a symbolic model. Now we will investigate *equivalence* properties. These properties formalize the notion of indistinguishability in a symbolic model.

6.1.1 Static equivalence

Remember that messages can be modelled as terms over an abstract algebra: given a set of function symbols \mathcal{F} , a set of names \mathcal{N} and a set of variables \mathcal{X} we denote the set of all terms constructed over these sets by $\mathcal{T}(\mathcal{F}, \mathcal{N} \uplus \mathcal{X})$. We consider here a slightly simplified setting where \mathcal{F} only contains public symbols, *i.e.*, $\mathcal{F} = \mathcal{F}_{\text{pub}}$. Moreover, the algebra will be equipped with an equational theory. Let us now be a bit more precise.

Definition 6.1 *An equational theory \mathcal{E} is a set of pairs $\{(M, N) \mid M, N \in \mathcal{T}(\mathcal{F}, \mathcal{N} \uplus \mathcal{X})\}$. We define equality modulo \mathcal{E} denoted $=_{\mathcal{E}}$ to be the smallest equivalence relation such that*

- $(M, N) \in \mathcal{E}$ implies that $M =_{\mathcal{E}} N$,
- $=_{\mathcal{E}}$ is closed under substitutions of terms for variables,
- $=_{\mathcal{E}}$ is closed under application of function symbols, *i.e.* $M_1 =_{\mathcal{E}} N_1, \dots, M_k =_{\mathcal{E}} N_k$ implies that $f(M_1, \dots, M_k) =_{\mathcal{E}} f(N_1, \dots, N_k)$ for all $f \in \mathcal{F}$ of arity k ,
- $=_{\mathcal{E}}$ is closed under bijective renaming of names.

For convenience in examples we will often just say that \mathcal{E} is defined by the equations

$$M_1 = N_1 \quad \dots \quad M_n = N_n$$

rather than writing $\mathcal{E} = \{(M_1, N_1), \dots, (M_n, N_n)\}$. So we will for instance define the equational theory `enc` which models symmetric encryption and pairings by the equations

$$\pi_1(\langle x, y \rangle) = x \quad \pi_2(\langle x, y \rangle) = y \quad \text{sdec}(\{x\}_y, y) = x$$

Hence we have that $\pi_2(\langle n, \text{sdec}(\{s\}_k, k) \rangle) =_{\text{enc}} s$.

In Section 3.1.4 we have already introduced static equivalence. Let us rephrase this definition now in this slightly simplified setting. We write $\text{new } \bar{n}_1.\sigma_1 =_\alpha \text{new } \bar{n}_2.\sigma_2$ when these two frames are the same up to α -conversion, i.e. up to a bijective renaming of the bound names. Before defining static equivalence between frames we define what it means for two terms to be equal in a frame. We say that $(M =_\mathcal{E} N) \text{new } \bar{n}.\sigma$ iff $\text{new } \bar{n}.\sigma =_\alpha \text{new } \bar{n}'.\sigma'$ and $\text{names}(M, N) \cap \bar{n}' = \emptyset$ and $M\sigma' =_\mathcal{E} N\sigma'$.

Definition 6.2 *Two frames $\text{new } \bar{n}_1.\sigma_1$ and $\text{new } \bar{n}_2.\sigma_2$ are statically equivalent, written $\text{new } \bar{n}_1.\sigma_1 \sim_\mathcal{E} \text{new } \bar{n}_2.\sigma_2$ iff $\text{Dom}(\sigma_1) = \text{Dom}(\sigma_2)$ and for all terms M, N we have that $(M =_\mathcal{E} N) \text{new } \bar{n}_1.\sigma_1 \Leftrightarrow (M =_\mathcal{E} N) \text{new } \bar{n}_2.\sigma_2$.*

Note that it follows directly from the definition that static equivalence is closed under α -conversion.

6.1.2 Applications of static equivalence

We have already seen some uses of static equivalence in Section 3.3. There we used static equivalence to model guessing attacks and to define equivalence properties. We will give here some additional examples.

Password protocols. It is not always possible to rely on previously shared keys or on an existing public key infrastructure. In these cases protocols often rely on shared passwords. However, such passwords are *weak secrets* and can be subject to brute-force attacks. Brute-force attacks on passwords are also referred to as *dictionnary attacks* (referring to a dictionnary containing all passwords) or *guessing attacks* (the search space is small enough that the attacker can guess the password). As we have seen before (Definition 3.3) resistance against offline guessing attacks can be modelled by the means of static equivalence.

Example 6.1 *Let us now consider an example of a password protocol, which indeed resists offline guessing attacks. The EKE protocol [?] can be described by the following 5 steps.*

1. $A \rightarrow B : \{(\text{pk}(k))\}_w$ (EKE.1)
2. $B \rightarrow A : \{\text{aenc}(r, \text{pk}(k))\}_w$ (EKE.2)
3. $A \rightarrow B : \{na\}_r$ (EKE.3)
4. $B \rightarrow A : \{\langle na, nb \rangle\}_r$ (EKE.4)
5. $A \rightarrow B : \{nb\}_r$ (EKE.5)

This protocol uses both ciphers and public key encryption. It does however not rely on a public key infrastructure, i.e. we do not assume that public keys are known and can be associated to a particular identity. In the first step (EKE.1) A generates a new private key k and sends the corresponding public key $\text{pk}(k)$ to B, encrypted (using symmetric encryption) with the shared password w . Then, B generates a fresh session key r , which he encrypts (using asymmetric encryption) with the previously received public key $\text{pk}(k)$. Finally, he encrypts the resulting ciphertext with the password w and sends the result to A (EKE.2). The last three steps (EKE.3-5) perform a handshake to avoid replay attacks. One may note that this is a password-only protocol. A new private and public key are used for each session and the only shared secret between different sessions is the password w .

We use the equational theory EKE defined by the equations

$$\text{adec}(\text{aenc}(x, \text{pk}(y)), y) = x \quad \text{sdec}(\text{senc}(x, y), y) = x \quad \text{senc}(\text{sdec}(x, y), y) = x \quad \pi_i(\langle x_1, x_2 \rangle) = x_i$$

where $i \in \{1, 2\}$. Note the modelling of ciphers with the additional equation $\text{senc}(\text{sdec}(x, y), y) = x$. The fact that decryption always succeeds is one of the differences between a cipher and a symmetric encryption. The protocol can be modeled in our process calculus by $\text{new } w.(P_A \parallel P_B)$ where

$$\begin{array}{ll} P_A = & \text{new } k, na. \\ & \text{out}(c, \{\text{pk}(k)\}_w). \\ & \text{in}(c, x_1). \\ & \text{let } ra = \text{adec}(\text{sdec}(x_1, w), k). \\ & \text{out}(c, \{na\}_{ra}) \\ & \text{in}(c, x_2). \\ & \text{if } \pi_1(\text{sdec}(x_2, ra)) = na \text{ then} \\ & \text{out}(c, \{\pi_2(\text{sdec}(x_2, ra))\}_{ra}) \\ P_B = & \text{new } r, nb. \\ & \text{in}(c, y_1). \\ & \text{out}(c, \{\text{aenc}(r, \text{sdec}(y_1, w))\}_w). \\ & \text{in}(c, y_2). \\ & \text{out}(c, \{\langle \text{sdec}(y_2, r), nb \rangle\}_r). \\ & \text{in}(c, y_3) \\ & \text{if } \text{sdec}(y_3, r) = nb \text{ then} \\ & 0 \end{array}$$

We use the *let* construction for readability, with the obvious meaning, i.e., $\text{let } x = M.P$ stands for $P\{M/x\}$. An honest execution of this protocol yields the frame $\text{new } w, k, r, na, nb.\sigma$ where

$$\sigma = \{x_1 \mapsto \{\text{pk}(k)\}_w, x_2 \mapsto \{\text{aenc}(r, \text{pk}(k))\}_w, x_3 \mapsto \{na\}_r, x_4 \mapsto \{\langle na, nb \rangle\}_r, x_5 \mapsto \{nb\}_r\}$$

We indeed have that $\text{new } w, k, r, na, nb.\sigma \uplus \{x_w \mapsto w\} \sim_{\text{EKE}} \text{new } w, w', k, r, na, nb.\sigma \uplus \{x_w \mapsto w'\}$. Using the tool ProVerif [?] it is possible to show that this equivalence holds for all reachable frames declaring w as *weaksecret*.

Remark 6.1 We have used static equivalence to model resistance against guessing attacks. One can note that the same modelization captures real-or-random properties in general.

Anonymity in electronic voting. Consider the following toy protocol for electronic voting where a voter sends his vote to an administrator, encrypted with the administrator's public key and signed with his private key. When the administrator has received all the votes he decrypts them and publishes the result. This can be modelled by the following voter and administrator processes. We consider the same equational theory for asymmetric encryption as before, with the slight difference that *aenc* will be a ternary function symbol, modelling randomization of the encryption.

$$V = \text{new } r; \text{out}(c_1, \text{sign}(\text{aenc}(v, r, \text{pk}(skA)), skV))$$

$$A = \text{in}(c_1, y); \text{out}(c_2, \text{adec}(\text{check}(y, \text{pk}(skV)), skA))$$

To model anonymity we consider two situations each involving two voters. In the first situation V_1 votes v_1 and V_2 votes v_2 ; in the second situation V_1 votes v_2 and V_2 votes v_1 , i.e. the two voters swap their vote. The protocol provides anonymity if these two situations are indistinguishable. Let $\sigma = \{z_1 \mapsto \text{pk}(skA), z_2 \mapsto \text{pk}(skV_1), z_3 \mapsto \text{pk}(skV_2)\}$ model the knowledge of public keys. The previous scenario can then be modelled by the following two general

process:

$$\begin{aligned}
VP_1 &= \text{new } skA, skV_1, skV_2; \\
&\quad (\sigma \| V\{skV_1/skV\}\{v_1/v\}\{r_1/r\} \| V\{skV_2/skV\}\{v_2/v\}\{r_2/r\} \| A\{skV_1/skV\} \| A\{skV_2/skV\}) \\
VP_2 &= \text{new } skA, skV_1, skV_2; \\
&\quad (\sigma \| V\{skV_1/skV\}\{v_2/v\}\{r_1/r\} \| V\{skV_2/skV\}\{v_1/v\}\{r_2/r\} \| A\{skV_1/skV\} \| A\{skV_2/skV\})
\end{aligned}$$

Two frames which can be derived from this protocol are

$$\begin{aligned}
\phi_1 &= \text{new } skA, skV_1, skV_2, v_1, r_1, v_2, r_2; \sigma \uplus \\
&\quad \{x_1 \mapsto \text{sign}(\text{aenc}(v_1, r_1, \text{pk}(skA)), skV_1); x_2 \mapsto \text{sign}(\text{aenc}(v_2, r_2, \text{pk}(skA)), skV_2); \\
&\quad x_3 \mapsto v_1; x_4 \mapsto v_2\}
\end{aligned}$$

$$\begin{aligned}
\phi_2 &= \text{new } skA, skV_1, skV_2, v_1, r_1, v_2, r_2; \sigma \uplus \\
&\quad \{x_1 \mapsto \text{sign}(\text{aenc}(v_2, r_1, \text{pk}(skA)), skV_1); x_2 \mapsto \text{sign}(\text{aenc}(v_1, r_2, \text{pk}(skA)), skV_2); \\
&\quad x_3 \mapsto v_1; x_4 \mapsto v_2\}
\end{aligned}$$

and anonymity can be modelled as $\phi_1 \sim_{\mathcal{E}} \phi_2$.

6.1.3 Some properties of static equivalence

We will now show some properties of static equivalence.

Proposition 6.1

$$\text{new } \bar{n}_1.\sigma_1 \sim_{\mathcal{E}} \text{new } \bar{n}_2.\sigma_2 \text{ iff } \text{new } \bar{n}_1 \uplus \bar{m}.\sigma_1 \uplus \theta \sim_{\mathcal{E}} \text{new } \bar{n}_2 \uplus \bar{m}.\sigma_2 \uplus \theta$$

where $\bar{m} \cap (\text{names}(\sigma_1) \cup \text{names}(\sigma_2)) = \emptyset$ and $(\bar{n}_1 \cup \bar{n}_2) \cap \text{names}(\theta) = \emptyset$.

A direct corollary of this proposition is the following.

Corollary 6.1

$$\begin{aligned}
&\text{new } \bar{n}_1.\sigma_1 \sim_{\mathcal{E}} \text{new } \bar{n}_2.\sigma_2 \text{ and } \text{new } \bar{m}_1.\theta_1 \sim_{\mathcal{E}} \text{new } \bar{m}_2.\theta_2 \\
&\quad \text{iff} \\
&\text{new } \bar{n}_1 \uplus \bar{m}_1.\sigma_1 \uplus \theta_1 \sim_{\mathcal{E}} \text{new } \bar{n}_2 \uplus \bar{m}_2.\sigma_2 \uplus \theta_2
\end{aligned}$$

where $\bar{m}_i \cap \text{names}(\sigma_i) = \emptyset$ and $\bar{n}_i \cap \text{names}(\theta_i) = \emptyset$ ($i \in \{1, 2\}$).

The \Rightarrow direction can be seen as a composition result provided the frames do not share any secret name. The \Leftarrow direction states that whenever two frames are statically equivalent then any subframe is also statically equivalent.

When proving results on static equivalence it is often more convenient to use the following equivalent definition of static equivalence.

Definition 6.3 *Two frames $\text{new } \bar{n}_1.\sigma_1$ and $\text{new } \bar{n}_2.\sigma_2$ are statically equivalent, written $\text{new } \bar{n}_1.\sigma_1 \sim_{\mathcal{E}} \text{new } \bar{n}_2.\sigma_2$ iff $\text{Dom}(\sigma_1) = \text{Dom}(\sigma_2)$ and for all terms M, N , such that $(\text{names}(M) \cup \text{names}(N)) \cap (\bar{n}_1 \cup \bar{n}_2) = \emptyset$ we have that $M\sigma_1 =_{\mathcal{E}} N\sigma_1 \Leftrightarrow M\sigma_2 =_{\mathcal{E}} N\sigma_2$.*

We leave it as an exercise to show that these two definitions are equivalent.

Proof : [of Proposition 6.1]

\Rightarrow We will show the contrapositive. Suppose that $\text{new } \bar{n}_1 \uplus \bar{m}. \sigma_1 \uplus \theta \not\sim_{\mathcal{E}} \text{new } \bar{n}_2 \uplus \bar{m}. \sigma_2 \uplus \theta$. Hence there exist two terms M, N such that $\text{names}(M, N) \cap (\bar{n}_1 \cup \bar{n}_2 \cup \bar{m}) = \emptyset$ and $M(\sigma_1 \uplus \theta) =_{\mathcal{E}} N(\sigma_1 \uplus \theta)$ and $M(\sigma_2 \uplus \theta) \neq_{\mathcal{E}} N(\sigma_2 \uplus \theta)$ (or vice-versa). As σ_i are ground $M(\sigma_i \uplus \theta) = (M\theta)\sigma_i$ and $N(\sigma_i \uplus \theta) = (N\theta)\sigma_i$. Because $\text{names}(M, N) \cap (\bar{n}_1 \cup \bar{n}_2) = \emptyset$ and $(\bar{n}_1 \cup \bar{n}_2) \cap \text{names}(\theta) = \emptyset$ we have that $\text{names}(M\theta, N\theta) \cap (\bar{n}_1 \cup \bar{n}_2) = \emptyset$. Hence, $M\theta, N\theta$ gives a valid test to distinguish $\text{new } \bar{n}_1. \sigma_1$ and $\text{new } \bar{n}_2. \sigma_2$.

\Leftarrow We again show the contrapositive. Suppose that $\bar{n}_1. \sigma_1 \not\sim_{\mathcal{E}} \text{new } \bar{n}_2. \sigma_2$. Hence there exist two terms M, N such that $\text{names}(M, N) \cap (\bar{n}_1 \cup \bar{n}_2) = \emptyset$ and $M\sigma_1 =_{\mathcal{E}} N\sigma_1$ and $M\sigma_2 \neq_{\mathcal{E}} N\sigma_2$ (or vice-versa). Let M' and N' be two terms obtained from M and N by applying a bijective renaming of names in $\text{names}(M, N) \cap \bar{m}$ and variables in $\text{Dom}(\theta)$ by fresh names. Hence, $\text{names}(M', N') \cap (\bar{m} \cup \bar{n}_1 \cup \bar{n}_2) = \emptyset$ and $\text{vars}(M', N') \cap \text{Dom}(\theta) = \emptyset$. Because $=_{\mathcal{E}}$ and $\neq_{\mathcal{E}}$ are closed under bijective renaming and $\bar{m} \cap (\text{names}(\sigma_1) \cup \text{names}(\sigma_2)) = \emptyset$ we have that $M'\sigma_1 =_{\mathcal{E}} N'\sigma_1$ and $M'\sigma_2 \neq_{\mathcal{E}} N'\sigma_2$. As $\text{vars}(M', N') \cap \text{Dom}(\theta) = \emptyset$ and σ_i are ground we obtain that $M'(\sigma_1 \uplus \theta) =_{\mathcal{E}} N'(\sigma_1 \uplus \theta)$ and $M'(\sigma_2 \uplus \theta) \neq_{\mathcal{E}} N'(\sigma_2 \uplus \theta)$.

While composition with shared secrets does not hold in general, resistance against offline guessing attacks (in the presence of a passive attacker) composes when the same password is used.

Proposition 6.2

$$\begin{aligned} & \text{new } w. \text{new } \bar{n}_1. \sigma_1 \uplus \{x \mapsto w\} \sim_{\mathcal{E}} \text{new } w. \text{new } w'. \text{new } \bar{n}_1. \sigma_1 \uplus \{x \mapsto w'\} \\ \text{and} \quad & \text{new } w. \text{new } \bar{n}_2. \sigma_2 \uplus \{x \mapsto w\} \sim_{\mathcal{E}} \text{new } w. \text{new } w'. \text{new } \bar{n}_2. \sigma_2 \uplus \{x \mapsto w'\} \\ & \text{implies that} \\ & \text{new } w. \text{new } \bar{n}_1 \uplus \bar{n}_2. \sigma_1 \uplus \sigma_2 \uplus \{x \mapsto w\} \sim_{\mathcal{E}} \text{new } w. \text{new } w'. \text{new } \bar{n}_1 \uplus \bar{n}_2. \sigma_1 \uplus \sigma_2 \uplus \{x \mapsto w'\} \end{aligned}$$

where $w' \notin \text{names}(\sigma_i)$ and $\bar{n}_1 \cap \text{names}(\sigma_2) = \bar{n}_2 \cap \text{names}(\sigma_1) = \emptyset$.

Exercise 6.4 will guide us through the proof of this proposition.

6.1.4 Further readings

This section does not give an exhaustive overview of all papers on the subject. The aim is to give some pointers to the interested reader to papers that are closely related (using similar techniques and notations) to the subjects covered in this chapter.

Static equivalence was first introduced in the applied pi calculus [?] which is a process calculus for cryptographic proptocols. In this context it was used to characterize observational equivalence by a labelled bisimulation, relying on static equivalence. Many properties of static equivalence were already introduced in that paper. Deducibility and static equivalence for equational theories are also discussed in detail in [?]. This paper also presents decidability and complexity results, which we will cover in the next section.

Guessing attacks were first formalized using static equivalence in [?] (other previous formalizations not based on static equivalence were presented in [?, ?] but they seem less natural). Composition results for guessing attacks were shown in [?]: this paper includes Proposition 6.2 showing that offline guessing attacks do compose in the presence of a passive adversary, it

is shown that offline guessing attacks do not compose in general in the presence of an active adversary, but sufficient conditions are given to achieve this composition.

The modelling of anonymity by the use of static and above all observational equivalence has been addressed for electronic voting in [?, ?, ?]. Anonymity in a different context is studied in [?]: there it is shown how to model anonymity in authentication protocols by the use of observational equivalence.

6.2 Procedure for subterm convergent equational theories

In this chapter we show that for a large class of equational theories, namely *subterm convergent* equational theories, static equivalence can be decided in polynomial time. This part follows closely the paper by Abadi and Cortier [?]. Other results and more practical procedures will be discussed in the further readings section.

6.2.1 Preliminaries

Let us first introduce some preliminary definitions and notations.

Notations. Given a set of function symbols \mathcal{F} we write $\text{ar}(\mathcal{F})$ for the maximal arity of $f \in \mathcal{F}$. Sometimes in this section, we will write $==$ when we want to emphasize that we mean syntactic equality. For a term t we define $|t|$ to be its size as follows: $|t| = 1$ if t is a variable, a name or a constant and $f(t_1, \dots, t_n) = 1 + \sum_{1 \leq i \leq n} |t_i|$. We denote by $st(t)$ the set of all subterms of t and by $pos(t)$ the set of its positions. Then, for $p \in pos(t)$ we write $t|_p$ for the subterm of t at position p . $t[u]_p$ is the term obtained by replacing the subterm at position p by u . The notions of size and subterms are naturally lifted to frames: for $\varphi = \text{new } \bar{n}. \{x_1 \mapsto M_1, \dots, x_n \mapsto M_n\}$ we define $|\varphi| = \sum_{1 \leq i \leq n} |M_i|$ and $st(\varphi) = \cup_{1 \leq i \leq n} st(M_i)$.

Rewrite systems and equational theories. A rewrite system \mathcal{R} is a set of rewrite rules $\ell \rightarrow r$ such that $\text{vars}(r) \subseteq \text{vars}(\ell)$. We say that a term t rewrites to u by \mathcal{R} if there exists $\ell \rightarrow r \in \mathcal{R}$ and $p \in pos(t)$ such that $t|_p = \ell\sigma$ for some substitution σ and $u = t[r\sigma]_p$. A rewrite system \mathcal{R} is terminating if there is no infinite chain $t_1 \rightarrow_{\mathcal{R}} t_2 \rightarrow_{\mathcal{R}} \dots$. A rewrite system \mathcal{R} is confluent if for any t_1, t_2 such that $t \rightarrow_{\mathcal{R}}^* t_1$, $t \rightarrow_{\mathcal{R}}^* t_2$ there exists u such that $t_1 \rightarrow_{\mathcal{R}}^* u$ and $t_2 \rightarrow_{\mathcal{R}}^* u$. \mathcal{R} is convergent if it is both confluent and terminating. For a convergent rewrite system \mathcal{R} we denote by $t \downarrow_{\mathcal{R}}$ the unique normal form of t . A rewrite system is subterm convergent if for any rule $\ell \rightarrow r \in \mathcal{R}$ we have that r is either a strict subterm of ℓ or r is a constant. Given an equational theory \mathcal{E} we associate a rewrite system $\mathcal{R}_{\mathcal{E}}$ to it by orienting the equations in \mathcal{E} . For readability we generally write $\rightarrow_{\mathcal{E}}$ instead of $\rightarrow_{\mathcal{R}_{\mathcal{E}}}$. We say that an equational theory is subterm convergent if it can be oriented to a subterm convergent rewrite system.

Example 6.2 The equational theories *enc*, *cipher* and *EKE* encountered before are subterm convergent. However, the equational theories for homomorphic encryption *homo*, extending *enc* by the equations

$$\{\langle x, y \rangle\}_z = \langle \{x\}_z, \{y\}_z \rangle \quad \text{sdec}(\langle x, y \rangle, z) = \langle \text{sdec}(x, z), \text{sdec}(y, z) \rangle$$

and for blind signatures blind defined by

$$\begin{aligned} \text{check}(\text{sign}(x, y), \text{pk}(y)) &= x \\ \text{unblind}(\text{blind}(x, y), y) &= x \\ \text{unblind}(\text{sign}(\text{blind}(x, y), z), y) &= \text{sign}(x, z) \end{aligned}$$

are not subterm convergent.

DAG representation of terms. In order to achieve a polynomial time complexity we will use a DAG representation of terms (which can be exponentially more succinct than its tree representation). A DAG representation of a term is a directed acyclic graph (V, l, E, v_0) where

- V is the set of vertices;
- $l : V \rightarrow \mathcal{F} \cup \mathcal{N} \cup \mathcal{X}$ is a labelling function;
- $E \subseteq (V \times \{1..ar(\mathcal{F})\} \times V)$ is the edge relation;
- v_0 is the root vertex.

For each vertex $v \in V$ such that $l(v) = f \in \mathcal{F}$, v has exactly $ar(f)$ outgoing edges labelled by 1 to $ar(f)$. Vertices labelled by names and variables have no outgoing edge. We denote by $E(v, i)$ the unique vertex v' such that $(v, i, v') \in E$. A DAG $D = (V, l, E, v_0)$ defines a term $t(D)$ as $t(D) = f(t(V, l, E, E(v_0, 1)), \dots, t(V, l, E, E(v_0, ar(f))))$ if $l(v_0) = f \in \mathcal{F}$ and $t(D) = l(v_0)$ if $l(v_0) \in \mathcal{N} \cup \mathcal{X}$. We say that a DAG representation (V, l, E, v_0) is minimal if there are no two distinct vertices $v_1, v_2 \in V$ such that $t(V, l, E, v_1) = t(V, l, E, v_2)$.

The size of a DAG D denoted $|D|$ is the number of its vertices. We define the DAG size of a term t , denoted $|t|_{\text{DAG}}$, as $|st(t)|$ which coincides with $|D|$ when D is a minimal DAG representation of t .

Given a DAG representation D we can compute its minimal DAG representation in $\mathcal{O}(|D|^3)$. We check whether there are two vertices v_1, v_2 (at most $|D|^2$ possibilities) such that $l(v_1) = l(v_2)$ and $E(v_1, i) = E(v_2, i)$ for $1 \leq i \leq ar(l(v_1))$. In that case we delete v_2 from V and replace any occurrence of v_2 by v_1 . We iterate this at most $|D|$ times.

Hence, we can also check whether $t(D_1) = t(D_2)$ in polynomial time in $|D_1| + |D_2|$.

Given a subterm convergent rewrite system \mathcal{R} and a minimal DAG representation $D = (V, l, E, v_0)$ of a term t we can compute the minimal DAG representation of $t \downarrow_{\mathcal{R}}$ in $\mathcal{O}(|D|^4)$. To see this note that each rewrite rule is of the form $C[x_1, \dots, x_n] \rightarrow C'[x_1, \dots, x_n]$ or $C[x_1, \dots, x_n] \rightarrow c$. Starting from the root we check whether the rewrite rule applies in D (at most $|C||D|$ tests). If it applies to some vertex v , i.e. $t(V, l, E, v) = C[x_1, \dots, x_n]\theta$ for some θ , replace v by the vertex representing $C'[x_1, \dots, x_n]\theta$ (if it is not a constant this vertex exists because it is a subterm; otherwise add a vertex labelled by the constant c). Then minimize the DAG in $\mathcal{O}(|D|^3)$. At each step (except for a constant number of cases) we delete one vertex. Hence the procedure stops after at most $|D|$ iterations and we compute the DAG representation of $t \downarrow_{\mathcal{R}}$ in $\mathcal{O}(|D|^4)$.

6.2.2 Deciding $\sim_{\mathcal{E}}$ in polynomial time for subterm convergent equational theories

Let \mathcal{E} be a subterm convergent equational theory defined by the axioms $\cup_{1 \leq i \leq n} \{(\ell_i, r_i)\}$. We define the theory constant $c_{\mathcal{E}} = \max_{1 \leq i \leq n} (|\ell_i|, ar(\mathcal{F}) + 1)$. By convention, we define $c_{\mathcal{E}}$ to be 1 in the case where \mathcal{E} and \mathcal{F} are empty.

Theorem 6.1 *Let φ and φ' be two frames. We can decide whether $\varphi \sim_{\mathcal{E}} \varphi'$ in polynomial time in $|\varphi| + |\varphi'|$.*

The remaining of this section will be dedicated to the proof of this result. The proof can be split into 3 steps:

1. frame saturation,
2. characterization of a frame by a finite set of equalities,
3. decidability of $\sim_{\mathcal{E}}$.

We now detail each of these steps.

6.2.2.1 Frame saturation

For each frame φ we can compute a set of terms $\text{sat}(\varphi)$. This set contains all terms that are deducible from the frame by applying only “small” contexts.

Definition 6.4 *Given a frame $\varphi = \text{new } \bar{n}. \{x_1 \mapsto M_1, \dots, x_k \mapsto M_k\}$ we define $\text{sat}(\varphi)$ to be the smallest set such that*

1. $\{M_1, \dots, M_k\} \subseteq \text{sat}(\varphi)$,
2. if $M_1, \dots, M_n \in \text{sat}(\varphi)$ and $f(M_1, \dots, M_n) \in \text{st}(\varphi)$ then $f(M_1, \dots, M_n) \in \text{sat}(\varphi)$,
3. if $M_1, \dots, M_n \in \text{sat}(\varphi)$ and C is a context such that $C[M_1, \dots, M_n] \rightarrow_{\mathcal{E}} M$ and $|C| \leq c_{\mathcal{E}}$, $\text{names}(C) \cap \bar{n} = \emptyset$ and $M \in \text{st}(\varphi)$ then $M \in \text{sat}(\varphi)$.

Example 6.3 *We again consider the equational theory enc which can be oriented into the following subterm convergent equational theory:*

$$\begin{aligned} \text{sdec}(\{x\}_y, y) &\rightarrow_{\text{enc}} x \\ \pi_i(\langle x_1, x_2 \rangle) &\rightarrow_{\text{enc}} x_i \quad (i \in \{1, 2\}) \end{aligned}$$

We have that $c_{\text{enc}} = 5$. Consider the frame

$$\varphi = \text{new } s_1, s_2, k_1, k_2. \{x_1 \mapsto \{\langle s_1, s_2 \rangle\}_{\langle k_1, k_2 \rangle}, x_2 \mapsto k_1, x_3 \mapsto k_2\}$$

We have that $\text{sat}(\varphi) = \{\{\langle s_1, s_2 \rangle\}_{\langle s_1, s_2 \rangle}, k_1, k_2, \langle k_1, k_2 \rangle, \langle s_1, s_2 \rangle, s_1, s_2\}$. The first three terms result directly from the frame (rule 1). The term $\langle k_1, k_2 \rangle$ can be added by applying the pairing function symbol (rule 2). Term $\langle s_1, s_2 \rangle$ can be added in two different ways using rule 3: either apply the context $\text{sdec}(-, -)$ on $\{\langle s_1, s_2 \rangle\}_{\langle s_1, s_2 \rangle}$ and $\langle k_1, k_2 \rangle$ or $\text{sdec}(-, \langle -, - \rangle)$ on $\{\langle s_1, s_2 \rangle\}_{\langle s_1, s_2 \rangle}$, k_1 and k_2 . s_i can be added by applying $\pi_i(-, -)$ on $\langle s_1, s_2 \rangle$. Note that we are not allowed to directly use the context $C = \pi_i(\text{sdec}(-, \langle -, - \rangle))$ (corresponding to the actual recipe for deducing s_i directly from φ) because $|C| > 5$.

We now show that this set can be computed in polynomial time in $|\varphi|$

Proposition 6.3 *The set $\text{sat}(\varphi)$ can be computed in $\mathcal{O}(|\varphi|^{\max(\text{ar}(\mathcal{F}), c_{\mathcal{E}})+2})$.*

Proof : We initialize the set with the elements $\{M_1, \dots, M_k\}$ (rule 1) and then saturate the set using rules 2 and 3. We notice that $\text{sat}(\varphi) \subseteq \text{st}(\varphi)$. Hence, we have that $\text{sat}(\varphi)$ contains at most $|\varphi|$ elements and the saturation will take at most $|\varphi|$ steps.

- If the step is an application of rule 2 we have to construct at most $|\mathcal{F}||\varphi|^{\text{ar}(\mathcal{F})}$ terms. For each of these terms we check whether it is in $\text{sat}(\varphi)$ which can be done in linear time in $|\varphi|$. Hence, this step can be computed in $\mathcal{O}(|\varphi|^{\text{ar}(\mathcal{F})+1})$.
- If the step is an application of rule 3 we have to compare whether any context of size $\leq c_{\mathcal{E}}$ applied to some M_i s in $\text{sat}(\varphi)$ is an instance of the lhs of a rewrite rule and check whether the resulting term is in $\text{sat}(\varphi)$. This can be computed in $\mathcal{O}(|\varphi|^{c_{\mathcal{E}}+1})$. To see this note that the number of contexts is constant (as $c_{\mathcal{E}}$ is fixed) and that a context has at most $c_{\mathcal{E}}$ holes. Hence, we need to consider at most $\mathcal{O}(|\varphi|^{c_{\mathcal{E}}})$ terms for which we need to check whether they are in $\text{sat}(\varphi)$.

Hence each step can be computed in $\mathcal{O}(|\varphi|^{\max(\text{ar}(\mathcal{F}), c_{\mathcal{E}})+1})$. As there are at most $|\varphi|$ steps we can compute $\text{sat}(\varphi)$ in $\mathcal{O}(|\varphi|^{\max(\text{ar}(\mathcal{F}), c_{\mathcal{E}})+2})$.

Moreover any element of $\text{sat}(\varphi)$ can be deduced using a recipe of small DAG size.

Proposition 6.4 *Let $\varphi = \text{new } \bar{n}.\sigma$. For all $M \in \text{sat}(\varphi)$ there exists a recipe R_M such that $|R_M|_{\text{DAG}} \leq c_{\mathcal{E}} \cdot |\varphi|$, $\text{names}(R_M) \cap \bar{n} = \emptyset$ and $R_M\sigma =_{\mathcal{E}} M$.*

Proof : For each term $M \in \text{sat}(\varphi)$ we can give a recipe R_M according to the rule of Definition 6.4 which added M to $\text{sat}(\varphi)$.

1. Let $R_M = x_i$ if $x_i\sigma = M$ and M is added by rule 1.
2. Let $R_M = f(R_{M_1}, \dots, R_{M_n})$ if M is added by rule 2.
3. Let $R_M = C(R_{M_1}, \dots, R_{M_n})$ if M is added by rule 3.

We can now construct a graph which contains each DAG corresponding to R_M for $M \in \text{sat}(\varphi)$. The graph is constructed as follows.

1. For each recipe R_M constructed by rule 1 add a vertex labelled by x_i .
2. For each recipe $R_M = f(R_{M_1}, \dots, R_{M_n})$ constructed by rule 2 add a vertex labelled by f and connect it to the vertices corresponding to the terms R_{M_1}, \dots, R_{M_n} .
3. For each recipe $R_M = C(R_{M_1}, \dots, R_{M_n})$ constructed by rule 3 construct the minimal DAG corresponding to C and connect it to the vertices corresponding to the terms R_{M_1}, \dots, R_{M_n} .

Steps 1 and 2 add 1 vertex to the graph. As $|C| \leq c_{\mathcal{E}}$ each step 3 adds at most $c_{\mathcal{E}}$ vertices to the graph. Moreover, we know that the graph can be constructed in $|\varphi|$ steps. Hence its size is bounded by $c_{\mathcal{E}} \cdot |\varphi|$.

Example 6.4 *Continuing example 6.3, let*

$$\begin{array}{ll}
 M_1 &= \{\langle s_1, s_2 \rangle\}_{\langle k_1, k_2 \rangle} \\
 M_2 &= k_1 \\
 M_3 &= k_2 \\
 M_4 &= \langle k_1, k_2 \rangle \\
 M_5 &= \langle s_1, s_2 \rangle \\
 M_6 &= s_1 \\
 M_7 &= s_2
 \end{array}$$

denote the elements of $\text{sat}(\varphi)$. We can define the following recipes

$$\begin{array}{ll} R_{M_i} = x_i & (1 \leq i \leq 3) \\ R_{M_4} = \langle x_2, x_3 \rangle \end{array} \quad \begin{array}{ll} R_{M_5} = \text{sdec}(x_1, R_{M_4}) \\ R_{M_6} = \pi_1(R_{M_5}) \\ R_{M_7} = \pi_2(R_{M_6}) \end{array}$$

6.2.2.2 Characterization of a frame by a finite set of equalities

To each frame φ , we can associate a set $\text{Eq}(\varphi)$ which is finite (up to renaming) and characterizes all equalities which hold in φ .

Definition 6.5 Given a frame φ we define $\text{Eq}(\varphi)$ to be the set of equalities

$$C_1[R_{M_1}, \dots, R_{M_k}] = C_2[R_{N_1}, \dots, R_{N_\ell}]$$

such that $|C_1|, |C_2| \leq c_\mathcal{E}$, $M_{1 \leq i \leq k}, N_{1 \leq j \leq \ell} \in \text{sat}(\varphi)$ and $(C_1[R_{M_1}, \dots, R_{M_k}] =_\mathcal{E} C_2[R_{N_1}, \dots, R_{N_\ell}])\varphi$. We write $\varphi' \models \text{Eq}(\varphi)$ iff $(M =_\mathcal{E} N)\varphi'$ for all $(M = N) \in \text{Eq}(\varphi)$.

Example 6.5 Let us continue with Example 6.4. Omitting redundant and trivial equations we have that $\text{Eq}(\varphi) = \{(\{R_{M_5}\}_{R_{M_4}} = R_{M_1}), \langle R_{M_6}, R_{M_7} \rangle = R_{M_5}\}$. Intuitively these equalities state that M_1 is indeed an encryption with the key M_4 and M_1 is the encryption of a pair.

We now show two crucial properties of $\text{Eq}(\varphi)$.

Lemma 6.1 Let $\varphi = \text{new } \bar{n}.\sigma$ and φ' be two frames such that $\varphi' \models \text{Eq}(\varphi)$. For all contexts C_1 and C_2 such that $\text{names}(C_1, C_2) \cap \bar{n} = \emptyset$ and for all terms $M_{1 \leq i \leq k}, N_{1 \leq j \leq \ell} \in \text{sat}(\varphi)$ if $C_1[M_1, \dots, M_k] = C_2[N_1, \dots, N_\ell]$ then $(C_1[R_{M_1}, \dots, R_{M_k}] =_\mathcal{E} C_2[R_{N_1}, \dots, R_{N_\ell}])\varphi'$.

Proof : The proof is done by induction on $|C_1| + |C_2|$.

Base case. $|C_1|, |C_2| \leq c_\mathcal{E}$. We have that

$$C_1[M_1, \dots, M_k] = C_2[N_1, \dots, N_\ell]$$

As $M_i, N_j \in \text{sat}(\varphi)$ we have by Proposition 6.4 that R_{M_i} , resp. R_{N_j} , are recipes for M_i , resp. N_j , in φ . Hence,

$$(C_1[R_{M_1}, \dots, R_{M_k}] =_\mathcal{E} C_2[R_{N_1}, \dots, R_{N_\ell}])\varphi$$

As $|C_1|, |C_2| \leq c_\mathcal{E}$, we have that $(C_1[R_{M_1}, \dots, R_{M_k}] = C_2[R_{N_1}, \dots, R_{N_\ell}]) \in \text{Eq}(\varphi)$ and hence

$$(C_1[R_{M_1}, \dots, R_{M_k}] =_\mathcal{E} C_2[R_{N_1}, \dots, R_{N_\ell}])\varphi'.$$

Inductive case. We consider two cases.

- $C_1 \neq _$ and $C_2 \neq _$.

In that case, we have that $C_1 = f(C_1^1, \dots, C_1^n)$ and $C_2 = f(C_2^1, \dots, C_2^n)$ and for $1 \leq i \leq n$

$$C_1^i[M_1, \dots, M_k] = C_2^i[N_1, \dots, N_\ell].$$

By induction hypothesis, we have that $(C_1^i[R_{M_1}, \dots, R_{M_k}] =_\mathcal{E} C_2^i[R_{N_1}, \dots, R_{N_\ell}])\varphi'$.

Hence, as $=_\mathcal{E}$ is closed under application of function symbols we also have that

$$(C_1[R_{M_1}, \dots, R_{M_k}] =_\mathcal{E} C_2[R_{N_1}, \dots, R_{N_\ell}])\varphi'.$$

- Either $C_1 == _$ or $C_2 == _$. Let us suppose that $C_1 == f(C_1^1, \dots, C_1^n)$ and $C_2 == _$. Let $M, M_1, \dots, M_k \in \text{sat}(\varphi)$ such that $C_1[M_1, \dots, M_k] == M$. Hence we have

$$f(C_1^1[M_1, \dots, M_k], \dots, C_1^n[M_1, \dots, M_k]) == M.$$

Let $N_i = C_1^i[M_1, \dots, M_k]$ for $1 \leq i \leq n$. As $M \in \text{sat}(\varphi)$ and $N_i \in \text{st}(M)$ we have that $N_i \in \text{st}(\text{sat}(\varphi))$. Applying iteratively rule 2 of Definition 6.4 we obtain that $N_i \in \text{sat}(\varphi)$. We can apply the induction hypothesis on $N_i == C_1^i[M_1, \dots, M_k]$ and obtain that $(R_{N_i} =_{\mathcal{E}} C_1^i[R_{M_1}, \dots, R_{M_k}])\varphi'$. Moreover, $M == f(N_1, \dots, N_n)$. Applying the base case (with contexts $_$ and $f(_, \dots, _)$) we obtain that $(R_M =_{\mathcal{E}} f(R_{N_1}, \dots, R_{N_n}))\varphi'$. From $(R_{N_i} =_{\mathcal{E}} C_1^i[R_{M_1}, \dots, R_{M_k}])\varphi'$ and $(R_M =_{\mathcal{E}} f(R_{N_1}, \dots, R_{N_n}))\varphi'$ we conclude that

$$(C_1[R_{M_1}, \dots, R_{M_k}] =_{\mathcal{E}} R_M)\varphi'.$$

Lemma 6.2 *Let $\varphi = \text{new } \bar{n}.\sigma$ and C_1 a context such that $\text{names}(C_1) \cap \bar{n} = \emptyset$. For all $M_1, \dots, M_k \in \text{sat}(\varphi)$ and T such that $C_1[M_1, \dots, M_k] \rightarrow_{\mathcal{E}}^* T$ there exist C_2 and $M'_1, \dots, M'_\ell \in \text{sat}(\varphi)$ such that $\text{names}(C_2) \cap \bar{n} = \emptyset$ and $T == C_2[M'_1, \dots, M'_\ell]$. Moreover, if $\varphi' \models \text{Eq}(\varphi)$ then $(C_1[R_{M_1}, \dots, R_{M_k}] =_{\mathcal{E}} C_2[R_{M'_1}, \dots, R_{M'_\ell}])\varphi'$.*

6.2.2.3 Decidability of $\sim_{\mathcal{E}}$

We now show the key proposition of the decidability proof.

Proposition 6.5

$$\varphi \sim_{\mathcal{E}} \varphi' \Leftrightarrow \varphi \models \text{Eq}(\varphi') \text{ and } \varphi' \models \text{Eq}(\varphi)$$

Proof :

\Rightarrow It follows directly from the definition of static equivalence that $\varphi \sim_{\mathcal{E}} \varphi' \Rightarrow \varphi \models \text{Eq}(\varphi')$ and $\varphi' \models \text{Eq}(\varphi)$.

\Leftarrow We have that $\varphi' \models \text{Eq}(\varphi)$. Let M, N be such that $(M =_{\mathcal{E}} N)\varphi$, i.e., $\varphi =_{\alpha} \text{new } \bar{n}.\sigma$ such that $\text{names}(M, N) \cap \bar{n} = \emptyset$ and $M\sigma =_{\mathcal{E}} N\sigma$. Hence $M\sigma \downarrow_{\mathcal{E}} = N\sigma \downarrow_{\mathcal{E}}$. Let $T = M\sigma \downarrow_{\mathcal{E}}$. By Lemma 6.2 we have that there exist C_M and $M_1, \dots, M_k \in \text{sat}(\varphi)$ such that $\text{names}(C_M) \cap \bar{n} = \emptyset$ and

$$T == C_M[M_1, \dots, M_k] \text{ and } (M =_{\mathcal{E}} C_M[R_{M_1}, \dots, R_{M_k}])\varphi'.$$

Similarly, as $T == N\sigma \downarrow_{\mathcal{E}}$ there exist C_N and $N_1, \dots, N_\ell \in \text{sat}(\varphi)$ such that $\text{names}(C_N) \cap \bar{n} = \emptyset$ and

$$T == C_N[N_1, \dots, N_\ell] \text{ and } (N =_{\mathcal{E}} C_N[R_{N_1}, \dots, R_{N_\ell}])\varphi'.$$

We obtain that $C_M[M_1, \dots, M_k] == C_N[N_1, \dots, N_\ell]$ and by Lemma 6.1 we obtain that

$$(C_M[R_{M_1}, \dots, R_{M_k}] =_{\mathcal{E}} C_N[R_{N_1}, \dots, R_{N_\ell}])\varphi'$$

and hence $(M =_{\mathcal{E}} N)\varphi'$.

Conversely, assuming $\varphi \models \text{Eq}(\varphi')$ and $(M =_{\mathcal{E}} N)\varphi'$ we obtain that $(M =_{\mathcal{E}} N)\varphi$. We conclude that $\varphi \sim_{\mathcal{E}} \varphi'$.

From this and previous propositions follows a polynomial time decision procedure. To decide $\varphi \sim_{\mathcal{E}} \varphi'$ construct $\text{sat}(\varphi)$ and $\text{sat}(\varphi')$. These constructions can be done in polynomial

time (Proposition 6.3). Moreover, each term in $\text{sat}(\varphi) \cup \text{sat}(\varphi')$ has polynomial DAG size (Proposition 6.4). Because of renamings, $\text{Eq}(\varphi)$ and $\text{Eq}(\varphi')$ may be infinite. However, as each of the equalities in $\text{Eq}(\varphi)$ and $\text{Eq}(\varphi')$ is of the form $C_1[R_{M_1}, \dots, R_{M_k}] = C_2[R_{N_1}, \dots, R_{N_\ell}]$ with $|C_i| \leq c_\mathcal{E}$ each equality contains at most $2 \cdot c_\mathcal{E}$ distinct names which can be fixed. There are at most $\mathcal{O}((|\varphi|^{c_\mathcal{E}})^2)$ equalities in $\text{Eq}(\varphi)$ each having a polynomial DAG size. Checking whether two terms in DAG representation are equal can be done in polynomial time as well. Hence, we can check in polynomial time that $\varphi \models \text{Eq}(\varphi')$ and $\varphi' \models \text{Eq}(\varphi)$.

6.2.3 Deciding $\sim_\mathcal{E}$ vs deciding $\vdash_\mathcal{E}$

An interesting question is what is the relation between the decidability problems for $\sim_\mathcal{E}$ and $\vdash_\mathcal{E}$. In general, these problems cannot be reduced to each other: there exist an equational theory, such that $\vdash_\mathcal{E}$ is decidable and $\sim_\mathcal{E}$ is not; there also exists (more surprisingly) an equational theory, such that $\sim_\mathcal{E}$ is decidable and $\vdash_\mathcal{E}$ is not. Indeed for many equational theories deciding $\vdash_\mathcal{E}$ can be reduced to $\sim_\mathcal{E}$. In particular, suppose that $\sim_\mathcal{E}$ is decidable over a signature $\mathcal{F} \uplus \{h\}$ where h is a free symbol. Then we can decide $\vdash_\mathcal{E}$ over the signature \mathcal{F} by the following encoding:

$$\text{new } \bar{n}.\sigma \vdash_\mathcal{E} t \text{ iff } \text{new } \bar{n}.\sigma \uplus \{x \mapsto h(t)\} \sim_\mathcal{E} \text{new } \bar{n}, a.\sigma \uplus \{x \mapsto a\}$$

where $a \notin \text{names}(\sigma)$. In particular this implies that $\vdash_\mathcal{E}$ is decidable in polynomial time for subterm convergent equational theories.

6.2.4 Further readings

While the here described procedure (from [?]) is indeed effective a direct implementation would not be efficient. Procedures for deciding static equivalence for subterm convergent equational theories have been proposed in [?, ?] and have been implemented in the tools YAPA and KISS. Both procedures can also be used to decide the theories *blind* and *homo* presented in Example 6.2. (The procedure presented in [?] and its implementation in the tool KISS is actually the outcome of an M2 internship following this course.) In [?], decidability for *blind* and *homo* (and a more general class of equational theories) was already shown, however, no generic procedure was presented.

In [?] it is shown that static equivalence is also decidable for a class of monoidal theories (which are theories over associative commutative operators including theories for exclusive or and abelian groups). Another important result [?] is that (under some mild assumptions) decision procedures for disjoint equational theories can be combined to obtain a procedure for deciding the joint equational theory.

While decidability of static equivalence has been extensively studied, the current knowledge on decidability of observational equivalence (the generalization to the active case) is very partial. We currently only have approximate procedures for an unboundned number of sessions [?] (implemented in the tool ProVerif), approximate procedures for a bounded number of sessions [?, ?, ?] and a decision procedure for a restricted class of processes [?].

The relationship between deciding deducibility and static equivalence has been first investigated in [?]. There the above encoding of deducibility into static equivalence is proposed and its correctness proven in detail. An equational theory which is decidable for static equivalence but not for deducibility is also presented. In [?] another example of such a theory is given with detailed proofs.

6.3 Exercises

Exercise 6.1 Consider the equation theories enc defined by the equations

$$\text{sdec}(\{x\}_y, y) = x \quad \pi_1(\langle x_1, x_2 \rangle) = x_1 \quad \pi_2(\langle x_1, x_2 \rangle) = x_2$$

and cipher which extends enc by the equation $\{\text{sdec}(x, y)\}_y = x$. Which of the following pairs of frames are statically equivalent? Whenever applicable give the distinguishing test.

$$\begin{array}{ll} \{x \mapsto a\} & \stackrel{?}{\sim}_{\text{enc}} \{x \mapsto b\} \\ \{x \mapsto \{0\}_k\} & \stackrel{?}{\sim}_{\text{enc}} \{x \mapsto \{1\}_k\} \\ \text{new } k. \{x \mapsto \{0\}_k\} & \stackrel{?}{\sim}_{\text{enc}} \text{new } k. \{y \mapsto \{1\}_k\} \\ \text{new } k. \{x \mapsto \{0\}_k\} & \stackrel{?}{\sim}_{\text{enc}} \text{new } k. \{x \mapsto \{1\}_k\} \\ \text{new } n, k. \{x \mapsto \{n\}_k, y \mapsto k\} & \stackrel{?}{\sim}_{\text{enc}} \text{new } n, k, k'. \{x \mapsto \{n\}_k, y \mapsto k'\} \\ \text{new } n, k. \{x \mapsto \{n\}_k, y \mapsto k\} & \stackrel{?}{\sim}_{\text{cipher}} \text{new } n, k, k'. \{x \mapsto \{n\}_k, y \mapsto k'\} \end{array}$$

Exercise 6.2 In Section 6.1.2 we considered a toy voting protocol. However, the anonymity relied on the fact that the administrator waits to have both votes before publishing the result (to model this we could enhance the language with a synchronization construct). Show that as specified above $VP_1 \not\approx_\ell VP_2$.

Exercise 6.3 Proposition 6.1 does not hold any more if we drop one of the condition $\overline{m} \cap (\text{names}(\sigma_1) \cup \text{names}(\sigma_2)) = \emptyset$ and $(\overline{n}_1 \cup \overline{n}_2) \cap \text{names}(\theta) = \emptyset$. Give a counter-example in each case when a condition is omitted.

Exercise 6.4 The aim of this exercise is to guide us through a proof of Proposition 6.2. This proof requires us to prove several properties of static equivalence. Suppose $\text{new } \overline{n}_1. \sigma_1 \sim_\mathcal{E} \text{new } \overline{n}_2. \sigma_2$. Show that

$$\text{new } n. \text{new } \overline{n}_1. \sigma_1 \sim_\mathcal{E} \text{new } n. \text{new } \overline{n}_2. \sigma_2 \text{ where } n \notin (\overline{n}_1 \cup \overline{n}_2), \quad (6.1)$$

$$\text{new } \overline{n}_1. \sigma_1 \{^s/n\} \sim_\mathcal{E} \text{new } \overline{n}_2. \sigma_2 \{^s/n\} \text{ where } n \notin (\overline{n}_1 \cup \overline{n}_2) \text{ and } s \text{ is a fresh name.} \quad (6.2)$$

Let $s \notin (\overline{n}_1 \cup \overline{n}_2)$. Show that

$$\text{new } \overline{n}_1. \sigma_1 \sim_\mathcal{E} \text{new } \overline{n}_2. \sigma_2 \text{ iff } \text{new } s. \text{new } \overline{n}_1. \sigma_1 \uplus \{x \mapsto s\} \sim_\mathcal{E} \text{new } s. \text{new } \overline{n}_2. \sigma_2 \uplus \{x \mapsto s\} \quad (6.3)$$

Use these properties (as well as Proposition 6.1) to show that the following 3 statements are equivalent.

1. $\text{new } w. \text{new } \overline{n}. \sigma \uplus \{x \mapsto w\} \sim_\mathcal{E} \text{new } w. \text{new } w'. \text{new } \overline{n}. \sigma \uplus \{x \mapsto w\}$
2. $\text{new } \overline{n}. \sigma \sim_\mathcal{E} \text{new } w. \text{new } \overline{n}. \sigma$
3. $\text{new } \overline{n}. \sigma \sim_\mathcal{E} \text{new } \overline{n}. \sigma \{^{w'}/w\}$ where w' is a fresh name.

From these results show Proposition 6.2.

Exercise 6.5 Give an example illustrating that the procedure given in Section 6.2 is wrong if we would define the theory constant $c_\mathcal{E}$ to be $\max_{1 \leq i \leq n}(|\ell_i|)$ instead of $\max_{1 \leq i \leq n}(|\ell_i|, \text{ar}(\mathcal{F})+1)$.

Exercise 6.6 We discussed the modelling of guessing attacks by static equivalence. Use the procedure of Section 6.2 to show that

$$\text{new } w.\text{new } n.\{z_1 \mapsto \{n\}_w; x_w \mapsto w\} \not\sim_{\text{enc}} \text{new } w.\text{new } w'.\text{new } n.\{z_1 \mapsto \{n\}_w; x_w \mapsto w'\}$$

and

$$\text{new } w.\text{new } n.\{z_1 \mapsto \{n\}_w; x_w \mapsto w\} \sim_{\text{cipher}} \text{new } w.\text{new } w'.\text{new } n.\{z_1 \mapsto \{n\}_w; x_w \mapsto w'\}$$

Exercise 6.7 The here presented procedure terminates in polynomial time given that terms are implemented by DAGs rather than trees. Give an example of two frames φ_1, φ_2 such that $\varphi_1 \not\sim_{\text{enc}} \varphi_2$ but every distinguishing test is of exponential size in $|\varphi_1| + |\varphi_2|$.

Exercise 6.8 Give an example of two frames φ_1 and φ_2 such that $\varphi_1 \not\sim_{\text{homo}} \varphi_2$ (for the theory of homomorphic encryption introduced in Example 6.2) for which the above procedure fails to distinguish the frames.

Chapter 7

Composition Results

This chapter is related to the work of R. Canetti *et al.* who study universal composability of protocols [?] in a different context (cryptographic model). They consider composition of a protocol with arbitrary other protocols and composition of a protocol with copies of itself. However, in the initial approach [?], they do not allow shared data between protocols, meaning that new encryption and signing keys have to be generated for each component. Note that in the symbolic setting, this kind of composition trivially holds since the definition of security properties always involved an arbitrary context. Indeed, if $\text{new } \bar{n}.P_1$ satisfies a secrecy property, this means that $(\text{new } \bar{n}.P_1) \parallel A$ satisfies the secrecy property for any attacker A . Hence, in particular, we have that the secrecy property holds if we consider a protocol P_2 running in parallel.

Composition theorems that allow joint states between protocols are proposed in further work (see *e.g.* [?, ?, ?]). Such composition theorems are more useful. They allow one to compose different protocols that share some keying material.

“if $\text{new } \bar{n}.P_1$ is secure, then $\text{new } \bar{n}.(P_1 \parallel P_2)$ is secure.”

In this chapter, we prove composability for tagged protocols and secrecy property only. Many parts of this chapter are borrowed from [?] and [?].

7.1 Motivation

When a protocol has been proved secure for an unbounded number of sessions, against a fully active adversary that can intercept, block and send new messages, there is absolutely no guarantee if the protocol is executed in an environment where other protocols are executed, possibly sharing some common keys like public keys or long-term symmetric keys. The interaction with the other protocols may dramatically damage the security of a protocol. Consider for example the two following naive protocols.

$$\begin{array}{ll}
 P_1 : & A \rightarrow B : \text{aenc}(s, \text{pk}(b)) \\
 P_2 : & A \rightarrow B : \text{aenc}(N_a, \text{pk}(b)) \\
 & B \rightarrow A : N_a
 \end{array}$$

In protocol P_1 , the agent A simply sends a secret s encrypted under B 's public key. In protocol P_2 , the agent sends some fresh nonce to B encrypted under B 's public key. The agent B acknowledges A 's message by forwarding A 's nonce. While P_1 executed alone easily

guarantees the secrecy of s , even against active adversaries, the secrecy of s is no more guaranteed when the protocol P_2 is executed. Indeed, an adversary may use the protocol P_2 as an oracle to decrypt any message. More realistic examples illustrating interactions between protocols can be found in *e.g.* [?].

The purpose of this chapter is to investigate sufficient conditions for a protocol to be safely used in an environment where other protocols may be executed as well. We show that whenever a protocol is proved secure when it is executed alone, its security is not compromised by the interactions with any other protocol, provided each protocol is given an identifier (e.g. the protocol's name) that should appear in any encrypted message. Continuing our example, let us consider the two slightly modified protocols.

$$P'_1 : \quad A \rightarrow B : \text{aenc}(\langle 1, s \rangle, \text{pk}(b)) \qquad P'_2 : \quad \begin{array}{l} A \rightarrow B : \text{aenc}(\langle 2, N_a \rangle, \text{pk}(b)) \\ B \rightarrow A : N_a \end{array}$$

The composition theorem (formally stated in Section 7.2) will ensure that P'_1 can be safely executed together with P'_2 , without compromising the secrecy of s .

Being able to share keys between protocols is a very desirable property as it allows to save both memory (for storing keys) and time since generating keys is time consuming in particular in the case of public key encryption. We provide in Section 7.2.3 examples of protocols that share secret materials. For security reasons however, most protocols currently make use of different keys.

The idea of adding an identifier in encrypted messages follows the spirit of the rules proposed in the paper of M. Abadi and R. Needham on prudent engineering practice for cryptographic protocols [?] (Principle 10). The use of unique protocol identifiers is also recommended in [?, ?] and has also been used in the design of fail-stop protocols [?]. In this chapter, we prove that it is sufficient for securely executing several protocols in the same environment. We will see in Section 7.3 how to extend this technique to compose a protocol with copies of itself.

7.2 Parallel Composition under Shared Keys

For sake of simplicity, we consider symmetric encryption only. We consider protocol written with pattern matching (no conditionals) and we do not use explicitly destructors.

7.2.1 Theorem

Even if a protocol is secure for an unbounded number of sessions, its security may collapse if the protocol is executed in an environment where other protocols sharing some common keys are executed. We have seen a first example in the previous section. We need to consider some restrictions.

Hypothesis 1. To avoid a ciphertext from a protocol P_1 to be decrypted in an another protocol P_2 , we introduce the notion of *well-tagged* protocol. This notion is formally defined below and relies on the following notion of *encrypted subterms*.

Let α be a term. We say that a term t is α -tagged if for every $\text{senc}(t_1, t_2) \in st(t)$, we have that $t_1 = \langle \alpha, t'_1 \rangle$ for some term t'_1 . A term is said *well-tagged* if it is α -tagged for some term α .

A protocol P is α -tagged if any term occurring in the role of the protocol is α -tagged. A protocol is said *well-tagged* if it is α -tagged for some term α .

Requiring that a protocol is well-tagged can be very easily achieved in practice: it is sufficient for example to add the name of the protocol in each encrypted term. Moreover, note that (as opposite to [?]) this does not require that the agents check that nested encrypted terms are correctly tagged. For example, let P be the following protocol:

$$P = \text{in}(c, \text{senc}(\langle \alpha, x \rangle, k)). \text{out}(c, \text{senc}(\langle \alpha, x \rangle, k')).$$

The protocol P is α -tagged. Note that the message $\text{senc}(\langle \alpha, \text{senc}(a, k) \rangle, k)$ (which is not α -tagged) would be accepted by P .

Hypothesis 2. Tagging protocols is not sufficient, indeed critical long-term keys should not be revealed in clear. Consider for example the two well-tagged protocols

$$P_3 : \quad A \rightarrow B : \text{senc}(\langle \alpha, s \rangle, k_{ab}) \qquad P_4 : \quad A \rightarrow B : k_{ab}$$

The security of protocol P_3 is again compromised by the execution of P_4 . Thus we will require that long-term keys do not occur in plaintext in the protocol. The set $\text{plaintext}(t)$ of *plaintext* of a term t is the set of names and variables, that is recursively defined as follows:

- $\text{plaintext}(u) = \{u\}$ when u is a variable or a name,
- $\text{plaintext}(\langle u_1, u_2 \rangle) = \text{plaintext}(u_1) \cup \text{plaintext}(u_2)$, and
- $\text{plaintext}(\text{senc}(u_1, u_2)) = \text{plaintext}(u_1)$.

This notation is extended to set of terms and protocols as expected.

Some weird protocols may still reveal critical keys in a hidden way. Consider for example the following (α -tagged) protocol.

$$P = \text{new } k_{ab}.(\text{out}(c, \text{senc}(\langle \alpha, a \rangle, k_{ab})) \parallel \text{in}(c, \text{senc}(\langle \alpha, a \rangle, x)). \text{out}(c, x))$$

While the long-term key k_{ab} does not appear in plaintext, the key k_{ab} is revealed after simply one normal execution of the protocol. This protocol is however not realistic since an unknown value cannot be learned (and sent) if it does not appear previously in plaintext. Thus we will further require (Condition 2 of Theorem 7.1) that a variable occurring in plaintext in a sent message, has to previously occur in plaintext in a received message.

We show that two well-tagged protocols can be safely composed as soon as they use different tags and that critical long-term keys do not appear in plaintext. Here, we assume that names are restricted. In other words, names are not known by the attacker unless they are explicitly given in his initial knowledge. Sometimes, we write $\text{new } \bar{n}.(P \parallel Q)$ to insist on the fact that those names \bar{n} are not known by the attacker or shared between the processes P and Q .

Theorem 7.1 *Let P_1 and P_2 be two well-tagged protocols such that P_1 is α -tagged and P_2 is β -tagged with $\alpha \neq \beta$. Let T_0 (intuitively the initial knowledge of the intruder) be a set of names. Let $\bar{n} = (fn(P_1) \cup fn(P_2)) \setminus T_0$ be the set of critical names. Let m be a term constructed from P_1 such that m is α -tagged and $fv(m) \subseteq bv(P_1)$. Moreover, we assume that*

1. critical names do not appear in plaintext, that is

$$\bar{n} \cap (\text{plaintext}(P_1) \cup \text{plaintext}(P_2)) = \emptyset.$$

2. any $x \in \text{bv}(P_1)$ (resp. $\text{bv}(P_2)$) “occurs first” at a plaintext position.

Then $\text{new } \bar{n}.P_1$ preserves the secrecy of m for the initial knowledge T_0 if, and only if, $\text{new } \bar{n}.(P_1 \parallel P_2)$ preserves the secrecy of m for T_0 .

We require that terms from P_1 and P_2 are tagged with distinct tags for simplicity. The key condition is actually that for any encrypted subterm t_1 of P_1 and for any encrypted subterm t_2 of P_2 , the terms t_1 and t_2 cannot be unified.

Theorem 7.1 is proved by contradiction. Assume that $\text{new } \bar{n}.(P_1 \parallel P_2)$ does not preserve the secrecy of m for T_0 . It means that there exists a scenario sce for $\text{new } \bar{n}.(P_1 \parallel P_2)$ such that the constraint system associated to this scenario, T_0 and m is satisfiable. In this case, the scenario can be turned into a scenario for P_1 such that the constraint system associated to this new scenario, T_0 and m is satisfiable, which means that P_1 does not preserve the secrecy of m for some initial knowledge T_0 , contradiction.

To prove this composition result, we first need to show that messages from two combined protocols do not need to be mixed up to mount an attack. For this purpose, we refine the decision procedure presented in Chapter 4 that allows us to control the form of the execution traces.

7.2.2 Main Steps of the Proof

To prove our composition result, we first refine the decision procedure presented in Chapter 4 for solving constraint systems. The *simplification rules* we consider are the following ones:

$$\begin{array}{ll}
R_1 : & \mathcal{C} \wedge T \vdash^? u \rightsquigarrow \mathcal{C} \quad \text{if } T \cup \{x \mid T' \vdash^? x \in \mathcal{C}, T' \subsetneq T\} \vdash u \\
R_2 : & \mathcal{C} \wedge T \vdash^? u \rightsquigarrow_{\sigma} \mathcal{C}\sigma \wedge T\sigma \vdash^? u\sigma \quad \text{if } \sigma = \text{mgu}(t, u) \text{ where } t \in \text{st}(T), t \neq u, \\
& \text{and } t, u \text{ are neither variables nor pairs} \\
R_3 : & \mathcal{C} \wedge T \vdash^? u \rightsquigarrow_{\sigma} \mathcal{C}\sigma \wedge T\sigma \vdash^? u\sigma \quad \text{if } \sigma = \text{mgu}(t_1, t_2), t_1, t_2 \in \text{st}(T), t_1 \neq t_2, \\
& \text{and } t_1, t_2 \text{ are neither variables nor pairs} \\
R_4 : & \mathcal{C} \wedge T \vdash^? u \rightsquigarrow \perp \quad \text{if } \text{vars}(T, u) = \emptyset \text{ and } T \not\vdash u \\
R_f : & \mathcal{C} \wedge T \vdash^? f(u, v) \rightsquigarrow \mathcal{C} \wedge T \vdash^? u \wedge T \vdash^? v \quad \text{for } f \in \{\langle \rangle, \text{senc}\}
\end{array}$$

Our rules are the same than those presented in Chapter 4 (or more precisely those given in Exercise 4.9) except that we forbid unification of terms headed by $\langle _, _ \rangle$. Correction and termination are still ensured and we show that they still form a complete decision procedure. Intuitively, unification between pairs is useless since pairs can be decomposed in order to perform unification on its components. Then, it is possible to build again the pair if necessary. Note that this is not always possible for encryption since the key used to decrypt or encrypt may be unknown by the attacker. Proving that forbidding unification between pairs still leads to a complete decision procedure required in particular to introduce a new notion of minimality for tree proofs for deduction (see Exercises 7.1 and 7.2). Note that this result is

of independent interest. Indeed, we provide a more efficient decision procedure for solving constraint systems, thus for deciding secrecy for a bounded number of sessions. Of course, the theoretical complexity remains the same (NP).

Theorem 7.2 *Let \mathcal{C} be an unsolved constraint system.*

1. (Correctness) *If $\mathcal{C} \rightsquigarrow_{\sigma}^* \mathcal{C}'$ for some constraint system \mathcal{C}' and some substitution σ and if θ is a solution of \mathcal{C}' then $\sigma\theta$ is a solution of \mathcal{C} .*
2. (Completeness) *If θ is a solution of \mathcal{C} , then there exist a solved constraint system \mathcal{C}' and substitutions σ, θ' such that $\theta = \sigma\theta'$, $\mathcal{C} \rightsquigarrow_{\sigma}^* \mathcal{C}'$ and θ' is a solution of \mathcal{C}' .*
3. (Termination) *There is no infinite chain $\mathcal{C} \rightsquigarrow_{\sigma_1} \mathcal{C}_1 \dots \rightsquigarrow_{\sigma_n} \mathcal{C}_n$.*

Theorem 7.1 is then proved by contradiction in three main steps. First, Theorem 7.2 serves as a key result for proving that if \mathcal{C} is satisfiable, then there exists a solution θ where messages from P_1 and P_2 are not mixed up. This is obtained by observing that the simplification rules enable us to build θ step by step through unification of subterms of P_1 and P_2 . Now, since unification between pairs is forbidden, the rules R_2 and R_3 only involve subterms that convey the same tag, *i.e.* subterms issued from the same protocol. Second, conditions 1 and 2 ensure that for any solution θ of \mathcal{C} , the critical names of \bar{n} do not appear in plaintext in $\mathcal{C}\theta$. Third, thanks to the two previous steps, we prove that β -tagged terms (intuitively messages from P_2) are not useful for deducing α -tagged terms. For this, we establish that $T \vdash u$ implies $\bar{T} \vdash \bar{u}$ where $\bar{\cdot}$ is a function that keep the terms issued from P_1 unchanged and projects the terms issued from P_2 on a public name. The proof is done by induction on the proof tree witnessing $T \vdash u$. It requires in particular the introduction of a new locality lemma for deduction of ground terms (see Exercise 7.3). Then, we deduce that, removing from \mathcal{C} all constraints inherited from P_2 and all β -tagged terms, we obtain a satisfiable constraint \mathcal{C}' that is associated to a scenario of P_1 .

7.2.3 Applications

Security protocols can be analysed using several existing tools, *e.g.* [?, ?]. The security of a protocol P is however guaranteed provided that no other protocols share any of the private data of P . Our result shows that, once the security of an isolated protocol has been established, this protocol can be safely executed in environments that may use some common data provided that a tag is added each time a party performs an encryption.

For security reasons however, most protocols actually make use of different keys. In this section, we provide a simple criteria for safely composing protocols that share keys. This would allow to save both memory (for storing keys) and time since generating keys is time consuming in particular in the case of public key encryption. For example, the SSL protocol should contain the name “ssl2.0” in any of its encrypted messages. This would ensure that no harmful interaction can occur with any other protocols even if they share some data with the SSL protocol, provided that these other protocols are also tagged. In other words, to avoid harmful interaction between protocols, one should simply use a tagged version of them.

There are also situations where running different protocols with common keys already occur. We provide three examples of such cases.

- It is often the case that several versions of a protocol can be used concurrently. In this case and for backward compatibility reasons, the same keys can be used in the different versions of the protocol, which may lead to potentially dangerous interactions.
- When encrypting emails the same public key can be used for two distinct encryption protocols: the PGP protocol (Pretty Good Privacy) and its open source version OpenPGP. Note that the PGP protocol contains also other sub-protocols such as digitally signing message, all sharing the same public-key infrastructure.
- In the slightly different context of APIs, J. Clulow [?] discovered an attack when the VISA PIN verification values (PVV) protocol and the IBM CCA API share the same verification key.

7.3 From One Session to Many

In this section we briefly present a protocol transformation which maps a protocol that is secure for a single session to a protocol that is secure for an unbounded number of sessions. This is another kind of parallel composition. Moreover, this provides an effective strategy to design secure protocols: (i) design a protocol intended to be secure for one protocol session (this can be efficiently verified with existing automated tools); (ii) apply our transformation and obtain a protocol which is secure for an unbounded number of sessions.

7.3.1 Protocol Transformation

Suppose that Π is a protocol between k participants A_1, \dots, A_k . Our transformation adds to Π a preamble in which each participant sends a freshly generated nonce N_i together with his identity to all other participants. This allows each participant to compute a dynamic, session dependent *tag* $\langle A_1, N_1 \rangle, \dots, \langle A_k, N_k \rangle$ that will be used to tag each encryption and signature in Π . Our transformation is surprisingly simple and does not require any cryptographic protection of the preamble. Intuitively, the security relies on the fact that the participant A_i decides on a given tag for a given session which is ensured to be fresh as it contains his own freshly generated nonce N_i . The transformation is computationally light as it does not add any cryptographic application; it may merely increase the size of messages to be encrypted or signed. The transformation applies to a large class of protocols, which may use symmetric and asymmetric encryption, digital signature and hash function.

We may note that, *en passant*, we identify a class of tagged protocols for which security is decidable for an unbounded number of sessions. This directly follows from our main result as it stipulates that verifying security for a single protocol session is sufficient to conclude security for an unbounded number of sessions.

Given an input protocol Π , our transformation will compute a new protocol $\tilde{\Pi}$ which consists of two phases. During the first phase, the protocol participants try to agree on some common, dynamically generated, session identifier τ . For this, each participant sends a freshly generated nonce N_i together with his identity A_i to all other participants. (Note that if broadcast is not practical or if not all identities are known to each participant, the message can be sent to some of the participants who forward the message.) At the end of this preamble, each participant computes a session identifier: $\tau = \langle \langle A_1, N_1 \rangle, \dots, \langle A_k, N_k \rangle \rangle$. Note that an active attacker may interfere with this initialisation phase and may intercept

and replace some of the nonces. Hence, the protocol participants do not necessarily agree on the same session identifier τ after this preamble. In fact, each participant computes his own session identifier, say τ_j . During the second phase, each participant j executes the original protocol in which the dynamically computed identifier is used for tagging each application of a cryptographic primitive. In this phase, when a participant opens an encryption, he will check that the tag is in accordance with the nonces he received during the initialisation phase. In particular he can test the presence of his own nonce.

The transformation, using the informal Alice-Bob notation, is described below and relies on the tagging operation that is formally defined in Definition ??.

$$\Pi = \left\{ \begin{array}{l} A_{i_1} \rightarrow A_{j_1} : m_1 \\ \vdots \\ A_{i_\ell} \rightarrow A_{j_\ell} : m_\ell \end{array} \right. \quad \tilde{\Pi} = \left\{ \begin{array}{ll} \text{Phase 1} & \text{Phase 2} \\ A_1 \rightarrow All : \langle A_1, N_1 \rangle & A_{i_1} \rightarrow A_{j_1} : [m_1]_\tau \\ \vdots & \vdots \\ A_k \rightarrow All : \langle A_k, N_k \rangle & A_{i_\ell} \rightarrow A_{j_\ell} : [m_\ell]_\tau \\ \text{where } \tau = \langle \text{tag}_1, \dots, \text{tag}_k \rangle \text{ with } \text{tag}_i = \langle A_i, N_i \rangle \end{array} \right.$$

Note that, the Alice-Bob notation only represents what happens in a normal execution, *i.e.* with no intervention of the attacker. Of course, in such a situation, the participants agree on the same session identifier τ used in the second phase.

A k -tag is a term $\langle \langle a_1, v_1 \rangle, \dots, \langle a_k, v_k \rangle \rangle$ where each a_i is a name of an agent and each v_i is a term. Let u be a term and tag be a k -tag. The k -tagging of u with tag , denoted $[u]_{\text{tag}}$, is inductively defined as follows:

$$\begin{aligned} [\langle u_1, u_2 \rangle]_{\text{tag}} &= \langle [u_1]_{\text{tag}}, [u_2]_{\text{tag}} \rangle \\ [\text{senc}(u_1, u_2)]_{\text{tag}} &= \text{senc}(\langle \text{tag}, [u_1]_{\text{tag}} \rangle, [u_2]_{\text{tag}}) \\ [u]_{\text{tag}} &= u && \text{otherwise} \end{aligned}$$

7.3.2 Composition Result

We only state the composition result informally.

“if the compiled protocol admits an attack that may involve several honest and dishonest sessions, then there exists an attack which only requires one honest session of each role (and no dishonest sessions). “

The situation is however slightly more complicated than it may seem at first sight since there is an infinite number of honest sessions, which one would need to verify separately. Actually we can avoid this combinatorial explosion thanks to the following well-known result [?]: when verifying secrecy properties it is sufficient to consider one single honest agent (which is allowed to “talk to herself”). Hence we can instantiate all the parameters with the same honest agent name.

Our dynamic tagging is useful to avoid interaction between different sessions of the same role in a protocol execution and allows us for instance to prevent man-in-the-middle attacks. We need also to forbid long-term secrets in plaintext position (even under an encryption). Note that this condition is generally satisfied by protocols and considered as a prudent engineering practice.

7.3.3 Other ways of tagging

We can also consider an alternative, slightly different transformation that does not include the identities in the tag, *i.e.*, the tag is simply the sequence of nonces. In that case we obtain a different result: if a protocol admits an attack then there exists an attack which only requires one (not necessarily honest) session for each role. In this case, we need to additionally check for attacks that involve a session engaged with the attacker. On the example of the Needham-Schroeder protocol the man-in-the-middle attack is not prevented by this weaker tagging scheme. However, the result requires one to also consider one dishonest session for each role, hence including the attack scenario. In both cases, it is important for the tags to be *collaborative*, *i.e.* all participants do contribute by adding a fresh nonce.

7.4 Further Readings

First, the results presented in this chapter are still true in a more general setting [?, ?], *e.g.* asymmetric encryption, authentication properties, ... Some other results have been obtained for a broader class of composition operations [?, ?]. In particular, their results allow sequential composition. This is useful to do refinement. For instance, we may study a protocol where a key is magically shared between two agents. Then, we want to ensure that the security property still hold when a real key establishment protocol is used.

Finally, different kinds of tags have also been considered in [?, ?, ?]. However these tags are *static* and have a different aim. While our dynamic tagging scheme presented in Section 7.3 avoids confusing messages from different sessions, these static tags avoid confusing different messages inside a same session and do not prevent that a same message is reused in two different sessions. Under some additional assumptions (*e.g.* no temporary secret, no ciphertext forwarding), several decidability results [?, ?] have been obtained by showing that it is sufficient to consider one session per role. But those results can not deal with protocols such as the Yahalom protocol or protocols which rely on a temporary secret. In the framework we consider here, the question whether such static tags would be sufficient to obtain decidability is still an open question (see [?]). In a similar way, static tags have also been used by Heather *et al.* [?] to avoid type confusion attacks.

Following the approach presented in this chapter, it has been shown [?] that tagging hashes enable to preserve resistance against guessing attacks under parallel composition. This allows one to safely reuse the same passwords for different applications.

A. Datta *et al.* (*e.g.* [?]) have also studied secure protocol composition in a broader sense: protocols can be composed in parallel, sequentially or protocols may use other protocols as components. However, they do not provide any syntactic conditions for a protocol P to be safely executed in parallel with other protocols. For any protocol P' that might be executed in parallel, they have to prove that the two protocols P and P' satisfy each other invariants. Their approach is thus rather designed for component-based design of protocols.

7.5 Exercises

Exercise 7.1 (★★) Let $T_1 \subseteq T_2 \subseteq \dots \subseteq T_n$. We say that a proof Π of $T_i \vdash u$ is left-minimal if for any $j < i$ such that $T_j \vdash u$, Π is also a proof of $T_j \vdash u$.

We say that a proof Π is simple if

1. any subproof of Π is left-minimal,
2. a composition rule is not directly followed by a decomposition rule,
3. any term of the form $\langle u_1, u_2 \rangle$ obtained by application of a decomposition rule or an axiom rule is followed by a projection rule.

Questions:

1. Let $T_1 = \{a\}$ and $T_2 = \{a, \text{senc}(\langle a, b \rangle, k), k\}$. We have that $T_2 \vdash \langle a, b \rangle$. Give a proof of $T_2 \vdash \langle a, b \rangle$ that is not simple.
2. Give a simple proof of $T_2 \vdash \langle a, b \rangle$.
3. Show that if there exists a proof of $T_i \vdash u$ then there exists a simple proof of it.

Exercise 7.2 (★★) We consider symmetric encryption/decryption and pairing/projection (again no asymmetric encryption). Show that the procedure made up of the rules R_1 , R_2 , R'_3 , and R_f is still complete if we do not perform unification on pairs. This result is stated in Theorem 7.2.

hint: use the notion of simple proof introduced in Exercise 7.1.

Exercise 7.3 (★) Consider the following inference system:

$$\frac{x \quad y}{\langle x, y \rangle} \quad \frac{\langle x, y \rangle}{x} \quad \frac{\langle x, y \rangle}{y} \quad \frac{x \quad y}{\text{senc}(x, y)} \quad \frac{\text{senc}(x, y) \quad y}{x}$$

Let T be a set of ground terms and u be a ground term such that $T \vdash u$. Then, we have that $\text{plaintext}(u) \subseteq \text{plaintext}(T)$.

Exercise 7.4 (★) Apply the transformation described in Section 7.3 on the Needham-Schroeder protocol.

1. Check that the man-in-the-middle attack does not exist anymore.
2. Remove the identity of the agent from the tag. Is this tag sufficient to avoid the man-in-the-middle attack?

Exercise 7.5 (★) In this chapter, we consider static tags to compose different protocols, and dynamic tags to compose different sessions coming from the same protocol. Show that dynamic tags, as described in Section 7.3 are not sufficient to compose different protocols.

Part III

Verification in the Computational Setting

Chapter 8

The Computational Model

8.1 Introduction

Since the beginning of public-key cryptography, with the seminal Diffie-Hellman paper [?], many suitable algorithmic problems for cryptography have been proposed and many cryptographic schemes have been designed, together with more or less heuristic proofs of their security. However, as already explained, most of those schemes have thereafter been broken, even when some kind of proofs existed. The simple fact that a cryptographic algorithm withstood cryptanalytic attacks for several years has often been considered as a kind of validation procedure, but some schemes take a long time before being broken. An example is the Chor-Rivest cryptosystem [?, ?], based on the knapsack problem, which took more than 10 years to be totally broken [?], whereas before this attack it was believed to be strongly secure. As a consequence, the lack of attacks at some time should never be considered as a security validation of the proposal.

A completely different paradigm is provided by the concept of “provable” security in the computational model. A significant line of research has tried to provide proofs in the framework of complexity theory (*a.k.a.* “reductionist” security proofs [?]): the proofs provide reductions from a well-studied computational problem (RSA or the discrete logarithm) to an attack against a cryptographic protocol. Adversaries, and any player, are now modeled by probabilistic Turing machines.

At the beginning, people just tried to define the security notions required by actual cryptographic schemes, and then to design protocols which achieve these notions. The techniques were directly derived from the complexity theory, providing polynomial reductions. However, their aim was essentially theoretical. They were indeed trying to minimize the required assumptions on the primitives (one-way functions or permutations, possibly trap-door, etc) [?, ?, ?, ?] without considering practicality. Therefore, they just needed to design a scheme with polynomial algorithms, and to exhibit polynomial reductions from the basic assumption on the primitive into an attack of the security notion, in an asymptotic way.

However, as just said, such a result has no practical impact on actual security. Indeed, even with a polynomial reduction, one may be able to break the cryptographic protocol within a few hours, whereas the reduction just leads to an algorithm against the underlying problem which requires many years. Therefore, those reductions only prove the security when very huge (and thus maybe unpractical) parameters are in use, under the assumption that no polynomial time algorithm exists to solve the underlying problem.

8.1.1 Exact Security and Practical Security

For more than 10 years now, more efficient reductions have been expected, under the denominations of either “exact security” [?] or “concrete security” [?], which provide more practical security results. The perfect situation is reached when one manages to prove that, from an attack, one can describe an algorithm against the underlying problem, with almost the same success probability within almost the same amount of time. We have then achieved “practical security”.

Unfortunately, in many cases, even just provable security is at the cost of an important loss in terms of efficiency for the cryptographic protocol. Thus some models have been proposed, trying to deal with the security of efficient schemes: some concrete objects are identified with ideal (or black-box) ones.

For example, it is by now usual to identify hash functions with ideal random functions, in the so-called “random-oracle model”, informally introduced by Fiat and Shamir [?], and formalized by Bellare and Rogaway [?]. Similarly, block ciphers are identified with families of truly random permutations in the “ideal cipher model” [?].

From a more algebraic point of view, another kind of idealization was introduced in cryptography: the black-box group [?, ?]. Here, the group operation is defined by a black-box: a new element necessarily comes from the addition (or the subtraction) of two already known elements. It is by now called the “generic model”. It has been more recently extended to bilinear groups [?]. Some works [?, ?] even require several ideal models together to provide some validations.

8.2 Security Proofs and Security Arguments

8.2.1 Computational Assumptions

In both symmetric and asymmetric scenarios, many security notions can not be unconditionally guaranteed (whatever the computational power of the adversary). Therefore, security generally relies on a computational assumption: the existence of one-way functions, or permutations, possibly trapdoor. A one-way function is a function f which anyone can easily compute, but given $y = f(x)$ it is computationally intractable to recover x (or any pre-image of y). A one-way permutation is a bijective one-way function. For encryption, one would like the inversion to be possible for the recipient only: a trapdoor one-way permutation is a one-way permutation for which a secret information (the trapdoor) helps to invert the function on any point.

Given such objects, and thus computational assumptions about the intractability of the inversion without possible trapdoors, we would like that security could be achieved without extra assumptions. The only way to formally prove such a fact is by showing that an attacker against the cryptographic protocol can be used as a sub-part in an algorithm that can break the basic computational assumption.

A partial order therefore exists between computational assumptions (and intractable problems too): if a problem P is more difficult than the problem P' (P' reduces to P , see below) then the assumption of the intractability of the problem P is weaker than the assumption of the intractability of the problem P' . The weaker the required assumption is, the more secure the cryptographic scheme is.

8.2.2 Complexity

In complexity theory, such an algorithm which uses the attacker as a sub-part in a global algorithm is called a reduction. If this reduction is polynomial, we can say that the attack of the cryptographic protocol is at least as hard as inverting the function: if one has a polynomial algorithm to solve the latter problem, one can polynomially solve the former one. In the complexity theory framework, a *polynomial* algorithm is the formalization of *efficiency*.

Therefore, in order to prove the security of a cryptographic protocol, one first needs to make precise the security notion one wants the protocol to achieve: which adversary's goal one wants to be intractable, under which kind of attack. At the beginning of the 1980's, such security notions have been defined for encryption [?] and signature [?, ?], and provably secure schemes have been suggested. However, those proofs had a theoretical impact only, because both the proposed schemes and the reductions were completely unpractical, yet polynomial. The reductions were indeed efficient (*i.e.* polynomial), and thus a polynomial attack against a cryptosystem would have led to a polynomial algorithm that broke the computational assumption. But the latter algorithm, even polynomial, may require hundreds of years to solve a small instance.

For example, let us consider a cryptographic protocol based on integer factoring. Let us assume that one provides a polynomial reduction from the factorization into an attack. But such a reduction may just lead to a factorization algorithm with a complexity in $2^{25}k^{10}$, where k is the bit-size of the integer to factor. This indeed contradicts the assumption that no-polynomial algorithm exists for factoring. However, on a 1024-bit number ($k = 2^{10}$), it provides an algorithm that requires 2^{125} basic operations, which is much more than the complexity of the best current algorithm, such as NFS [?], which needs less than 2^{100} (see Section 9.2). Therefore, such a reduction would just be meaningful for numbers above 4096 bits (since with $k = 2^{12}$, $2^{145} < 2^{149}$, where 2^{149} is the estimate effort for factoring a 4096-bit integer with the best algorithm.) Concrete examples are given later.

8.2.3 Practical Security

Moreover, most of the proposed schemes were unpractical as well. Indeed, the protocols were polynomial in time and memory, but not efficient enough for practical implementation.

For a few years, people have tried to provide both practical schemes, with practical reductions and exact complexity, which prove the security for realistic parameters, under a well-defined assumption: exact reduction in the standard model (which means in the complexity-theoretic framework). For example, under the assumption that a 1024-bit integer cannot be factored with less than 2^{70} basic operations, the cryptographic protocol cannot be broken with less than 2^{60} basic operations. We will see such an example later.

Unfortunately, as already remarked, practical or even just efficient reductions in the standard model can rarely be conjugated with practical schemes. Therefore, one needs to make some hypotheses on the adversary: the attack is generic, independent of the actual implementation of some objects

- hash functions, in the “random-oracle model”;
- symmetric block ciphers, in the “ideal-cipher model”;
- algebraic groups, in the “generic model”.

The “random-oracle model” was the first to be introduced in the cryptographic community [?, ?], and has already been widely accepted. By the way, flaws have been shown in the “generic model” [?] on practical schemes, and the “random-oracle model” is not equivalent to the standard one either. Several gaps have already been exhibited [?, ?, ?]. However, all the counter-examples in the random-oracle model are pathological, counter-intuitive and not natural. Therefore, in the sequel, we focus on security analyses in this model, for real and natural constructions. A security proof in the random-oracle model will at least give a strong argument in favor of the security of the scheme. Furthermore, proofs in the random-oracle model under a weak computational assumption may be of more practical interest than proofs in the standard model under a strong computational assumption.

8.2.4 The Random-Oracle Model

As said above, efficiency rarely meets with provable security. More precisely, none of the most efficient schemes in their category have been proven secure in the standard model. However, some of them admit security validations under ideal assumptions: the random-oracle model is the most widely accepted one.

Many cryptographic schemes use a hash function \mathcal{H} (such as MD5 [?], SHA-1 [?] or more recent functions). This use of hash functions was originally motivated by the wish to sign long messages with a single short signature. In order to achieve *non-repudiation*, a minimal requirement on the hash function is the impossibility for the signer to find two different messages providing the same hash value. This property is called *collision-resistance*.

It was later realized that hash functions were an essential ingredient for the security of, first, signature schemes, and then of most cryptographic schemes. In order to obtain security arguments, while keeping the efficiency of the designs that use hash functions, a few authors suggested using the hypothesis that \mathcal{H} behaves like a random function. First, Fiat and Shamir [?] applied it heuristically to provide a signature scheme “as secure as” factorization. Then, Bellare and Rogaway [?, ?, ?] formalized this concept for cryptography, and namely for signature and public-key encryption.

In this model, the so-called “random-oracle model”, the hash function can be formalized by an oracle which produces a truly random value for each new query. Of course, if the same query is asked twice, identical answers are obtained. This is precisely the context of relativized complexity theory with “oracles,” hence the name.

About this model, no one has ever been able to provide a convincing contradiction to its practical validity, but just theoretical counter-examples on either clearly wrong designs for practical purpose [?], or artificial security notions [?, ?]. Therefore, this model has been strongly accepted by the community, and is considered as a good one, in which security analyses give a good taste of the actual security level. Even if it does not provide a formal proof of security (as in the standard model, without any ideal assumption), it is argued that proofs in this model ensure security of the overall design of the scheme provided that the hash function has no weakness, hence the name “security arguments”.

This model can also be seen as a restriction on the adversary’s capabilities. Indeed, it simply means that the attack is generic without considering any particular instantiation of the hash functions. Therefore, an actual attack would necessarily use a weakness or a specific feature of the hash function. The replacement of the hash function by another one would rule out this attack.

On the other hand, assuming the tamper-resistance of some devices, such as smart cards,

the random-oracle model is equivalent to the standard model, which simply requires the existence of pseudo-random functions [?, ?].

As a consequence, almost all the standards bodies by now require designs provably secure, at least in that model, thanks to the security validation of very efficient protocols.

8.2.5 The General Framework

Before going into more details of this kind of proofs, we would like to insist on the fact that in the current general framework, we model all the players by (interactive) probabilistic polynomial Turing machines, and we give the adversary complete access to the cryptographic primitive, but as a black-box. It can ask any query of its choice, and the box always answers correctly, in constant time. Such a model does not consider timing attacks [?], where the adversary tries to extract the secrets from the computational time. Some other attacks analyze the electrical energy required by a computation to get the secrets [?], or to make the primitive fail on some computation [?, ?]. They are not captured either by this model.

Chapter 9

Provable Security

9.1 Security Notions

In this section we describe more formally what a signature scheme and an encryption scheme are. Moreover, we make precise their goals, and thus the security notions one wants the schemes to achieve. This is the first imperative step towards provable security.

9.1.1 Public-Key Encryption

The aim of a public-key encryption scheme is to allow anybody who knows the public key of Alice to send her a message that she will be the only one able to recover, granted her private key.

9.1.1.1 Definitions

A public-key encryption scheme $S = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ is defined by the three following algorithms:

- The *key generation algorithm* \mathcal{K} . On input 1^k where k is the security parameter, the algorithm \mathcal{K} produces a pair (pk, sk) of matching public and private keys. Algorithm \mathcal{K} is probabilistic.
- The *encryption algorithm* \mathcal{E} . Given a message m and a public key pk , \mathcal{E} produces a ciphertext c of m . This algorithm may be probabilistic. In the latter case, we write $\mathcal{E}_{\text{pk}}(m; r)$ where r is the random input to \mathcal{E} .
- The *decryption algorithm* \mathcal{D} . Given a ciphertext c and the private key sk , $\mathcal{D}_{\text{sk}}(c)$ gives back the plaintext m . This algorithm is necessarily deterministic.

9.1.1.2 Security Notions

The **goals** of the adversary may be various. The first common security notion that one would like for an encryption scheme is *one-wayness* (OW): with just public data, an attacker cannot get back the whole plaintext of a given ciphertext. More formally, this means that for any adversary \mathcal{A} , its success in inverting \mathcal{E} without the private key should be negligible over the probability space $\mathbf{M} \times \Omega$, where \mathbf{M} is the message space and Ω is the space of the random

coins r used for the encryption scheme, and the internal random coins of the adversary:

$$\text{Succ}_S^{\text{ow}}(\mathcal{A}) = \Pr_{m,r}[(\text{pk}, \text{sk}) \leftarrow \mathcal{K}(1^k) : \mathcal{A}(\text{pk}, \mathcal{E}_{\text{pk}}(m; r)) = m].$$

However, many applications require more from an encryption scheme, namely the *semantic security* (IND) [?], *a.k.a. polynomial security/indistinguishability of encryptions*: if the attacker has some information about the plaintext, for example that it is either “yes” or “no” to a crucial query, any adversary should not learn more with the view of the ciphertext. This security notion requires computational impossibility to distinguish between two messages, chosen by the adversary, which one has been encrypted, with a probability significantly better than one half: its advantage $\text{Adv}_S^{\text{ind}}(\mathcal{A})$, formally defined as

$$2 \times \Pr_{b,r} \left[\begin{array}{l} (\text{pk}, \text{sk}) \leftarrow \mathcal{K}(1^k), (m_0, m_1, s) \leftarrow \mathcal{A}_1(\text{pk}), \\ c = \mathcal{E}_{\text{pk}}(m_b; r) : \mathcal{A}_2(m_0, m_1, s, c) = b \end{array} \right] - 1,$$

where the adversary \mathcal{A} is seen as a 2-stage attacker $(\mathcal{A}_1, \mathcal{A}_2)$, should be negligible.

A later notion is *non-malleability* (NM) [?]. To break it, the adversary, given a ciphertext, tries to produce a new ciphertext such that the plaintexts are meaningfully related. This notion is stronger than the above semantic security, but it is equivalent to the latter in the most interesting scenario [?] (the CCA attacks, see below). Therefore, we will just focus on one-wayness and semantic security.

On the other hand, an attacker can play many kinds of attacks, according to the **available information**: since we are considering asymmetric encryption, the adversary can encrypt any plaintext of its choice, granted the public key, hence the *chosen-plaintext attack* (CPA). It may furthermore have access to additional information, modeled by partial or full access to some oracles:

- A validity-checking oracle which, on input a ciphertext c , answers whether it is a valid ciphertext or not. Such a weak oracle, involved in the so-called reaction attacks [?] or *Validity-Checking Attack* (VCA), had been enough to break some famous encryption schemes [?, ?].
- A plaintext-checking oracle which, on input a pair (m, c) , answers whether c encrypts the message m . This attack has been termed the *Plaintext-Checking Attack* (PCA) [?].
- The decryption oracle itself, which on any ciphertext answers the corresponding plaintext. There is of course the natural restriction not to ask the challenge ciphertext to that oracle.

For all these oracles, access may be restricted as soon as the challenge ciphertext is known, the attack is thus said *non-adaptive* since oracle queries cannot depend on the challenge ciphertext, while they depend on previous answers. On the opposite, access can be unlimited and attacks are thus called *adaptive attacks* (w.r.t. the challenge ciphertext). This distinction has been widely used for the chosen-ciphertext attacks, for historical reasons: the *non-adaptive chosen-ciphertext attacks* (CCA1) [?], *a.k.a.* lunchtime attacks, and *adaptive chosen-ciphertext attacks* (CCA2) [?]. The latter scenario which allows adaptively chosen ciphertexts as queries to the decryption oracle is definitely the strongest attack, and will be named the *chosen-ciphertext attack* (CCA).

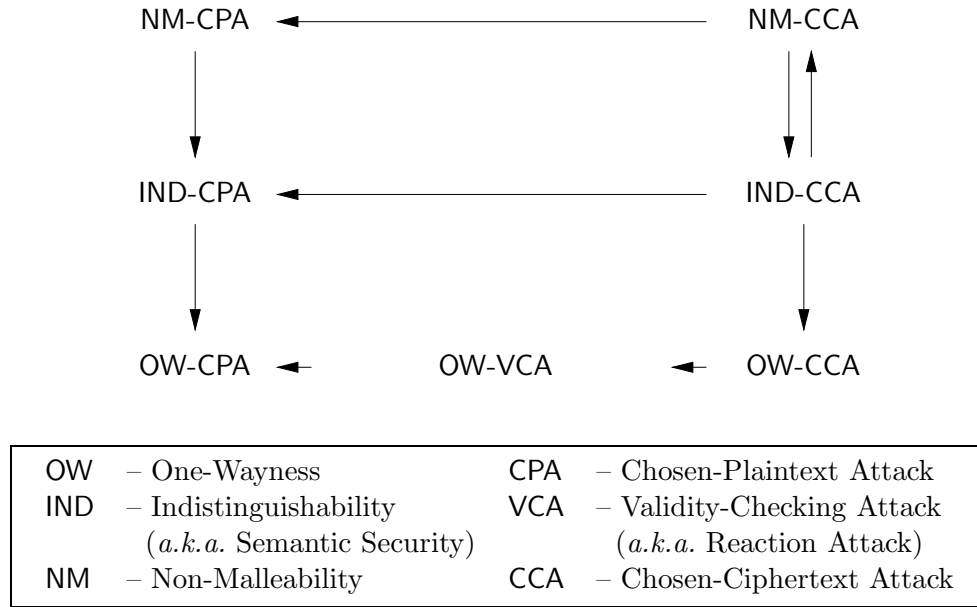


Figure 9.1: Relations between the Security Notions for Asymmetric Encryption

Furthermore, multi-user scenarios can be considered where related messages are encrypted under different keys to be sent to many people (*e.g.* broadcast of encrypted data). This may provide many useful data for an adversary. For example, RSA is well-known to be weak in such a scenario [?, ?], namely with a small encryption exponent, because of the Chinese Remainders Theorem. But once again, semantic security has been shown to be the appropriate security level, since it automatically extends to the multi-user setting: if an encryption scheme is semantically secure in the classical sense, it is also semantically secure in multi-user scenarios, against both passive [?] and active [?] adversaries.

A general study of these security notions and attacks was conducted in [?], we therefore refer the reader to this paper for more details. See also the summary diagram on Figure 9.1. However, we can just review the main scenarios we will consider in the following:

- one-wayness under chosen-plaintext attacks (OW-CPA) – where the adversary wants to recover the whole plaintext from just the ciphertext and the public key. This is the weakest scenario.
- semantic security under adaptive chosen-ciphertext attacks (IND-CCA) – where the adversary just wants to distinguish which plaintext, between two messages of its choice, has been encrypted, while it can ask any query it wants to a decryption oracle (except the challenge ciphertext). This is the strongest scenario one can define for encryption (still in our general framework.) Thus, this is our goal when we design a cryptosystem.

9.1.2 Digital Signature Schemes

Digital signature schemes are the electronic version of handwritten signatures for digital documents: a user's signature on a message m is a string which depends on m , on public and secret data specific to the user and —possibly— on randomly chosen data, in such a way that anyone can check the validity of the signature by using public data only. The user's public

data are called the *public key*, whereas his secret data are called the *private key*. The intuitive security notion would be the impossibility to forge user's signatures without the knowledge of his private key. In this section, we give a more precise definition of signature schemes and of the possible attacks against them (most of those definitions are based on [?]).

9.1.2.1 Definitions

A signature scheme $S = (\mathcal{K}, \mathcal{S}, \mathcal{V})$ is defined by the three following algorithms:

- The *key generation algorithm* \mathcal{K} . On input 1^k , which is a formal notation for a machine with running time polynomial in k (1^k is indeed k in basis 1), the algorithm \mathcal{K} produces a pair (pk, sk) of matching public and private keys. Algorithm \mathcal{K} is probabilistic. The input k is called the security parameter. The sizes of the keys, or of any problem involved in the cryptographic scheme, will depend on it, in order to achieve an appropriate security level (the expected minimal time complexity of any attack).
- The *signing algorithm* \mathcal{S} . Given a message m and a pair of matching public and private keys (pk, sk) , \mathcal{S} produces a signature σ . The signing algorithm might be probabilistic.
- The *verification algorithm* \mathcal{V} . Given a signature σ , a message m and a public key pk , \mathcal{V} tests whether σ is a valid signature of m with respect to pk . In general, the verification algorithm need not be probabilistic.

9.1.2.2 Forgeries and Attacks

In this subsection, we formalize some security notions which capture the main practical situations. On the one hand, the **goals** of the adversary may be various:

- Disclosing the private key of the signer. It is the most serious attack. This attack is termed *total break*.
- Constructing an efficient algorithm which is able to sign messages with good probability of success. This is called *universal forgery*.
- Providing a new message-signature pair. This is called *existential forgery*. The corresponding security level is called *existential unforgeability* (EUF).

In many cases the latter forgery, the *existential forgery*, is not dangerous because the output message is likely to be meaningless. Nevertheless, a signature scheme which is existentially forgeable does not guarantee by itself the identity of the signer. For example, it cannot be used to certify randomly looking elements, such as keys. Furthermore, it cannot formally guarantee the non-repudiation property, since anyone may be able to produce a message with a valid signature.

On the other hand, various **means** can be made available to the adversary, helping it into its forgery. We focus on two specific kinds of attacks against signature schemes: the *no-message attacks* and the *known-message attacks* (KMA). In the former scenario, the attacker only knows the public key of the signer. In the latter, the attacker has access to a list of valid message-signature pairs. According to the way this list was created, we usually distinguish many subclasses, but the strongest is definitely the *adaptive chosen-message attack* (CMA),

where the attacker can ask the signer to sign any message of its choice, in an adaptive way: it can adapt its queries according to previous answers.

When signature generation is not deterministic, there may be several signatures corresponding to a given message. And then, some notions defined above may become ambiguous [?]. First, in known-message attacks, an existential forgery becomes the ability to forge a fresh message/signature pair that has not been obtained during the attack. There is a subtle point here, related to the context where several signatures may correspond to a given message. We actually adopt the stronger rule that the attacker needs to forge the signature of message, whose signature was not queried. The more liberal rule, which makes the attacker successful when it outputs a second signature of a given message different from a previously obtained signature of the same message, is called *malleability*, while the corresponding security level is called *non-malleability* (NM). Similarly, in adaptive chosen-message attacks, the adversary may ask several times the same message, and each new answer gives it some information. A slightly weaker security model, by now called *single-occurrence adaptive chosen-message attack* (SO-CMA), allows the adversary at most one signature query for each message. In other words the adversary cannot submit the same message twice for signature.

When one designs a signature scheme, one wants to computationally rule out at least existential forgeries, or even achieve non-malleability, under adaptive chosen-message attacks. More formally, one wants that the success probability of any adversary \mathcal{A} with a reasonable time is small, where

$$\text{Succ}_S^{\text{euf}}(\mathcal{A}) = \Pr \left[(\text{pk}, \text{sk}) \leftarrow \mathcal{K}(1^k), (m, \sigma) \leftarrow \mathcal{A}^{\mathcal{S}_{\text{sk}}}(\text{pk}) : \mathcal{V}(\text{pk}, m, \sigma) = 1 \right].$$

We remark that since the adversary is allowed to play an adaptive chosen-message attack, the signing algorithm is made available, without any restriction, hence the oracle notation $\mathcal{A}^{\mathcal{S}_{\text{sk}}}$. Of course, in its answer, there is the natural restriction that, at least, the returned message-signature has not been obtained from the signing oracle \mathcal{S}_{sk} itself (non-malleability) or even the output message has not been queried (existential unforgeability).

9.2 The Computational Assumptions

There are two major families in number theory-based public-key cryptography:

1. the schemes based on integer factoring, and on the RSA problem [?];
2. the schemes based on the discrete logarithm problem, and on the Diffie-Hellman problems [?], in any “suitable” group. The first groups in use were cyclic subgroups of \mathbb{Z}_p^* , the multiplicative group of the modular quotient ring $\mathbb{Z}_p = \mathbb{Z}/p\mathbb{Z}$. But many schemes are now converted on cyclic subgroups of elliptic curves, or of the Jacobian of hyper-elliptic curves, with namely the so-called ECDSA [?], the US Digital Signature Standard [?] on elliptic curves. The pairing-based cryptography also works in groups (elliptic curves) where the discrete logarithm problem is difficult, but a bilinear map provides variable intractability levels for the Diffie-Hellman problems, according to the actual curve.

9.2.1 Integer Factoring and the RSA Problem

The most famous intractable problem is factorization of integers: while it is easy to multiply two prime integers p and q to get the product $n = p \cdot q$, it is not simple to decompose n into its prime factors p and q .

Currently, the most efficient algorithm is based on sieving on number fields. The Number Field Sieve (NFS) method [?] has a super-polynomial, but sub-exponential, complexity in $\mathcal{O}(\exp((1.923 + o(1))(\ln n)^{1/3}(\ln \ln n)^{2/3}))$. It has been used to establish the main records, in august 1999 (by factoring a 155-digit integer, 512-bit long [?]) and in december 2009 (by factoring a 232-digit integer, 768-bit long [?]). The factored numbers, called RSA-155 and RSA-768, were taken from the “RSA Challenge List”, which is used as a yardstick for the security of the RSA cryptosystem (see later). The latter is used extensively in hardware and software to protect electronic data traffic such as in the SSL (Security Sockets Layer) Handshake Protocol.

```

RSA-155 =
1094173864157052742180970732204035761200373294544920\
5990913842131476349984288934784717997257891267332497\
625752899781833797076537244027146743531593354333897
= 102639592829741105772054196573991675900\
716567808038066803341933521790711307779
* 106603488380168454820927220360012878679\
207958575989291522270608237193062808643

RSA-232 =
1230186684530117755130494958384962720772853569595334792197\
3224521517264005072636575187452021997864693899564749427740\
6384592519255732630345373154826850791702612214291346167042\
9214311602221240479274737794080665351419597459856902143413
= 3347807169895689878604416984821269081770479498371376856891\
2431388982883793878002287614711652531743087737814467999489
* 3674604366679959042824463379962795263227915816434308764267\
6032283815739666511279233373417143396810270092798736308917

```

Unfortunately, integer multiplication just provides a one-way function, without any possibility to invert the process. No information is known to make factoring easier. However, some algebraic structures are based on the factorization of an integer n , where some computations are difficult without the factorization of n , but easy with it: the finite quotient ring \mathbb{Z}_n which is isomorphic to the product ring $\mathbb{Z}_p \times \mathbb{Z}_q$ if $n = p \cdot q$.

For example, the e -th power of any element x can be easily computed using the *square-and-multiply* method. It consists in using the binary representation of the exponent $e = e_k e_{k-1} \dots e_0$, computing the successive 2 powers of x ($x^{2^0}, x^{2^1}, \dots, x^{2^k}$) and eventually to multiply altogether the ones for which $e_i = 1$. However, to compute e -th roots, it seems that one requires to know an integer d such that $ed = 1 \bmod \varphi(n)$, where $\varphi(n)$ is the totient Euler function which denotes the cardinality of the multiplicative subgroup \mathbb{Z}_n^* of \mathbb{Z}_n . In the particular case where $n = pq$, $\varphi(n) = (p-1)(q-1)$. And therefore, $ed - 1$ is a multiple of $\varphi(n)$ which is equivalent to the knowledge of the factorization of n [?]. In 1978, Rivest, Shamir and Adleman [?] defined the following problem:

The RSA Problem. Let $n = pq$ be the product of two large primes of similar size and e an integer relatively prime to $\varphi(n)$. For a given $y \in \mathbb{Z}_n^*$, compute the modular e -th root x of y (i.e. $x \in \mathbb{Z}_n^*$ such that $x^e = y \bmod n$.)

Year	Required Complexity	modulus bitlength
before 2000	64	768
before 2010	80	1024
before 2020	112	2048
before 2030	128	3072
	192	7680
	256	15360

Figure 9.2: Bitlength of RSA Moduli

The Euler function can be easily computed from the factorization of n , since for any $n = \prod p_i^{v_i}$,

$$\varphi(n) = n \times \prod \left(1 - \frac{1}{p_i}\right).$$

Therefore, with the factorization of n (the trapdoor), the RSA problem can be easily solved. But nobody knows whether the factorization is required, and how to do without it either:

The RSA Assumption. For any product of two primes, $n = pq$, large enough, the RSA problem is intractable (presumably as hard as the factorization of n).

More precisely, the intractability level is assumed as depicted on Figure 9.2. Until 2000, a security level of 2^{64} was considered enough since the computational power that an adversary could collect within a reasonable time was bounded by 2^{64} . Since 2010, 2^{80} is no longer considered enough, and thus 2048-bit RSA moduli are recommended. However, this is under the assumption that the RSA problem can be efficiently reduced to an attack, without any loss (practical security).

9.2.2 The Discrete Logarithm and the Diffie-Hellman Problems

The setting is quite general: one is given

- a cyclic group \mathcal{G} of prime order q (such as the finite group $(\mathbb{Z}_q, +)$, a subgroup of (\mathbb{Z}_p^*, \times) for $q|p-1$, of an elliptic curve, etc);
- a generator \mathbf{g} (*i.e.* $\mathcal{G} = \langle \mathbf{g} \rangle$).

We note in bold (such as \mathbf{g}) any element of the group \mathcal{G} , to distinguish it from a scalar $x \in \mathbb{Z}_q$. But such a \mathbf{g} could be an element in \mathbb{Z}_p^* or a point of an elliptic curve, according to the setting. Above, we talked about a “suitable” group \mathcal{G} . In such a group, some of the following problems have to be hard to solve (using the additive notation).

- the **Discrete Logarithm** problem (**DL**): given $\mathbf{y} \in \mathcal{G}$, compute $x \in \mathbb{Z}_q$ such that $\mathbf{y} = x \cdot \mathbf{g} = \mathbf{g} + \dots + \mathbf{g}$ (x times), then one writes $x = \log_{\mathbf{g}} \mathbf{y}$.
- the **Computational Diffie-Hellman** problem (**CDH**): given two elements in the group \mathcal{G} , $\mathbf{a} = a \cdot \mathbf{g}$ and $\mathbf{b} = b \cdot \mathbf{g}$, compute $\mathbf{c} = ab \cdot \mathbf{g}$. Then one writes $\mathbf{c} = \mathbf{DH}(\mathbf{a}, \mathbf{b})$.
- the **Decisional Diffie-Hellman** Problem (**DDH**): given three elements in the group \mathcal{G} , $\mathbf{a} = a \cdot \mathbf{g}$, $\mathbf{b} = b \cdot \mathbf{g}$ and $\mathbf{c} = c \cdot \mathbf{g}$, decide whether $\mathbf{c} = \mathbf{DH}(\mathbf{a}, \mathbf{b})$ (or equivalently, whether $c = ab \bmod q$).




Figure 9.3: Proof by Reduction

It is clear that they are sorted from the strongest problem to the weakest one. Furthermore, one may remark that they all are “random self-reducible”, which means that any instance can be reduced to a uniformly distributed instance: for example, given a specific element \mathbf{y} for which one wants to compute the discrete logarithm x in basis \mathbf{g} , one can choose a random $z \in \mathbb{Z}_q$, and compute $\mathbf{z} = z \cdot \mathbf{y}$. The element \mathbf{z} is therefore uniformly distributed in the group, and the discrete logarithm $\alpha = \log_{\mathbf{g}} \mathbf{z}$ leads to $x = \alpha/z \bmod q$. As a consequence, there are only average complexity cases. Thus, the ability to solve a problem for a non-negligible fraction of instances in polynomial time is equivalent to solve any instance in expected polynomial time.

A variant of the Diffie-Hellman problem has been defined by Tatsuaki Okamoto and David Pointcheval [?], the so-called *Gap Diffie-Hellman Problem (GDH)*, where one wants to solve the **CDH** problem with an access to a **DDH** oracle. One may easily remark the following properties about above problems: $\mathbf{DL} \geq \mathbf{CDH} \geq \{\mathbf{DDH}, \mathbf{GDH}\}$, where $A \geq B$ means that the problem A is at least as hard as the problem B . However, in practice, no one knows how to solve any of them without breaking the **DL** problem itself. On pairing-friendly elliptic curves, the **DDH** can be easy to decide.

Currently, the most efficient algorithms to solve the latter problem depend on the underlying group. For generic groups (for which no specific algebraic property can be used), algorithms have a complexity in the square root of q , the order of the generator \mathbf{g} [?, ?]. For example, on well-chosen elliptic curves only these algorithms can be used.

However, for subgroups of \mathbb{Z}_p^* , some better techniques can be applied. The best algorithm is based on sieving on number fields, as for the factorization. The General Number Field Sieve method [?] has a super-polynomial, but sub-exponential, complexity in $\mathcal{O}(\exp((1.923 + o(1))(\ln p)^{1/3}(\ln \ln p)^{2/3}))$.

For signature applications, one only requires groups where the **DL** problem is hard, whereas encryption needs trapdoor problems and therefore requires groups where some of the **DH**’s problems are also hard to solve.

9.3 Proof Methodology

The actual security proof consists of a reduction of the underlying computational problem to an attack against the cryptographic scheme: if there exist an adversary \mathcal{A} able to win the security game, we can use this adversary \mathcal{A} as a subroutine to solve the computational problem, as shown of Figure 9.3. More precisely, a proof consists in building a simulator. In order to be able to evaluate the success probability of the simulator in solving the computational problem, we have to show that the view of the adversary remains unchanged from the one its

has during a real attack.

Until the early 2000's, a proof consisted in exhibiting the simulator, and then to directly analyze the success probability. This analysis was intricate and error-prone. Victor Shoup introduced in [?, ?, ?] a game-based approach, also revisited by Bellare and Rogaway [?], that we will extensively use in these notes.

In this technique, we define a sequence $\mathbf{G}_1, \mathbf{G}_2$, etc., of modified attack games starting from the actual security game \mathbf{G}_0 . Each of the games operates on the same underlying probability space: the public and private keys of the cryptographic scheme, the coin tosses of the adversary \mathcal{A} and the various oracles. Only the rules defining how the view is computed differ from game to game: we modify the behavior of the oracles and of the challenger. The view of the adversary (similar to the trace of an execution in the symbolic model) can be seen as a random variable following a distribution probability \mathcal{D}_i in the game \mathbf{G}_i .

Then, several situations can appear:

- the distribution remains perfectly unchanged, then the distance between the two distributions is 0;
- the distribution remains statistically unchanged, then the distance between the two distributions is negligible;
- the distributions are computationally indistinguishable, then a decisional problem is involved;

In these three cases, the probability of any event is almost the same in the two games: the difference is bounded by the distance between the two distributions.

We can also modify the distributions, but in specific cases only: unless an event is raised, the two games run identically. As a consequence, the distributions of outputs of the two games are related by the event. We can indeed apply the following Shoup lemma:

Lemma 9.1 *Let E_1, E_2 and F_1, F_2 be events defined on a probability space*

$$\Pr[E_1 \mid \neg F_1] = \Pr[E_2 \mid \neg F_2] \text{ and } \Pr[F_1] = \Pr[F_2] = \varepsilon \Rightarrow |\Pr[E_1] - \Pr[E_2]| \leq \varepsilon.$$

Proof : The proof follows from easy computations:

$$\begin{aligned} |\Pr[E_1] - \Pr[E_2]| &= |\Pr[E_1 \mid F_1] \cdot \Pr[F_1] + \Pr[E_1 \mid \neg F_1] \cdot \Pr[\neg F_1] \\ &\quad - \Pr[E_2 \mid F_2] \cdot \Pr[F_2] - \Pr[E_2 \mid \neg F_2] \cdot \Pr[\neg F_2]| \\ &= |(\Pr[E_1 \mid F_1] - \Pr[E_2 \mid F_2]) \cdot \varepsilon \\ &\quad + (\Pr[E_1 \mid \neg F_1] - \Pr[E_2 \mid \neg F_2]) \cdot (1 - \varepsilon)| \\ &= |(\Pr[E_1 \mid F_1] - \Pr[E_2 \mid F_2]) \cdot \varepsilon| \leq \varepsilon. \end{aligned}$$

9.4 Exercises

Exercise 9.1 (Amplification of RSA) *Let us consider the RSA problem, with modulus n and exponent e , prime to $\varphi(n)$. We say that the algorithm \mathcal{A} is an (ε, t) -adversary against RSA with parameters (n, e) if, within time t , its success probability is greater than ε :*

$$\text{Succ}(\mathcal{A}) = \Pr_{x \in \mathbb{Z}_n^*} [\mathcal{A}(n, e, x^e \bmod n) = x] \geq \varepsilon.$$

Show that if there exists an (ε, t) -adversary \mathcal{A} against RSA with parameters (n, e) , then for any $0 < \eta < 1$, there exists an (η, t') -adversary \mathcal{B} against RSA with the same parameters (n, e) , within a reasonable time bound t' .

Exercise 9.2 (Amplification of Diffie-Hellman) Let us consider the Diffie-Hellman problem in a cyclic group \mathcal{G} of prime order q , with a generator g : given $X = g^x$ and $Y = g^y$, one has to compute $Z = g^{xy}$.

We say that the algorithm \mathcal{A} is an (ε, t) -adversary against DH with parameters (\mathcal{G}, g) if, within time t , its success probability is greater than ε :

$$\text{Succ}(\mathcal{A}) = \Pr_{x, y \in \mathbb{Z}_q^*} [\mathcal{A}(g^x, g^y) = g^{xy}] \geq \varepsilon.$$

1. Can we use the same amplification method as above to build an (η, t') -adversary for any $0 < \eta < 1$?

From an instance $(A = g^a, B = g^b)$, let us derive two instances: $(X_i = A^{\alpha_i} g^{u_i}, Y_i = B^{\beta_i} g^{v_i})$, for $i = 0, 1$, with random exponents $\alpha_i, \beta_i, u_i, v_i$. We then run twice our adversary \mathcal{A} on each derived instance.

2. Show that we can detect if \mathcal{A} succeeded on the two instances with negligible error
3. Show that we can detect if \mathcal{A} failed for at least one instance, with negligible error
4. Describe the amplified algorithm \mathcal{B} , and estimate the time and the success probability
5. Show that u_i and v_i are important in the randomization: zero values could lead to a wrong algorithm \mathcal{B} .

Chapter 10

Public-Key Encryption Schemes

10.1 Introduction

10.1.1 The RSA Encryption Scheme

In their seminal paper [?], Rivest, Shamir and Adleman proposed both signature and public-key encryption schemes, thanks to the “trapdoor one-way permutation” property of the RSA function: the generation algorithm produces a large composite number $N = pq$, a public key e , and a private key d such that $e \cdot d = 1 \bmod \varphi(N)$. The encryption of a message m , encoded as an element in \mathbb{Z}_N^* , is simply $c = m^e \bmod N$. This ciphertext can be easily decrypted thanks to the knowledge of d , $m = c^d \bmod N$. Clearly, this encryption is OW-CPA, relative to the RSA problem. The determinism makes a plaintext-checking oracle useless. Indeed, the encryption of a message m , under a public key \mathbf{pk} is always the same, and thus it is easy to check whether a ciphertext c really encrypts m , by re-encrypting it. Therefore the RSA-encryption scheme is OW-PCA relative to the RSA problem as well.

Because of this determinism, it cannot be semantically secure: given the encryption c of either m_0 or m_1 , the adversary simply computes $c' = m_0^e \bmod N$ and checks whether $c' = c$. Furthermore, with a small exponent e (e.g. $e = 3$), any security vanishes under a multi-user attack: given $c_1 = m^3 \bmod N_1$, $c_2 = m^3 \bmod N_2$ and $c_3 = m^3 \bmod N_3$, one can easily compute $m^3 \bmod N_1 N_2 N_3$ thanks to the Chinese Remainders Theorem, which is exactly m^3 in \mathbb{Z} and therefore leads to an easy recovery of m .

10.1.2 The El Gamal Encryption Scheme

In 1985, El Gamal [?] also designed both signature and public-key encryption schemes. The latter is based on the Diffie-Hellman key exchange protocol [?]: given a cyclic group \mathcal{G} of order prime q and a generator \mathbf{g} , the generation algorithm produces a random element $x \in \mathbb{Z}_q^*$ as private key, and a public key $\mathbf{y} = x \cdot \mathbf{g}$. The encryption of a message m , encoded as an element \mathbf{m} in \mathcal{G} , is a pair $(\mathbf{c} = a \cdot \mathbf{g}, \mathbf{d} = a \cdot \mathbf{y} + \mathbf{m})$, for a random $a \in \mathbb{Z}_q$. This ciphertext can be easily decrypted thanks to the knowledge of x , since

$$a \cdot \mathbf{y} = ax \cdot \mathbf{g} = x \cdot \mathbf{c},$$

and thus $\mathbf{m} = \mathbf{d} - x \cdot \mathbf{c}$. This encryption scheme is well-known to be OW-CPA relative to the Computational Diffie-Hellman problem. It is also semantically secure (against chosen-plaintext attacks) relative to the Decisional Diffie-Hellman problem [?].

As we have seen above, the expected security level is IND-CCA, whereas the RSA encryption just reaches OW-CPA under the RSA assumption, and the El Gamal encryption achieves IND-CPA under the DDH assumption. Can we achieve IND-CCA for practical encryption schemes?

10.2 The Cramer-Shoup Encryption Scheme

As we will see later, the scheme using hash functions, modeled as random oracles in the security analyses, are definitely the most efficient schemes, but let us first start with a nice variation of the ElGamal encryption scheme, proposed by Cramer and Shoup [?], that is both rather efficient and IND-CCA secure under the DDH assumption.

10.2.1 Description

For the description of this scheme, we will use the multiplicative notation, as in the original paper: We thus work in a cyclic group \mathcal{G} of order prime q , with two independent generators g_1 and g_2 . We will also need a hash function \mathcal{H} , that will be assumed to be second-preimage resistant. Then the encryption scheme $\text{CS} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ can be described as follows:

- $\mathcal{K}(1^k)$: produces a random element $z \in \mathbb{Z}_q^*$, as in the above ElGamal scheme, but also four additional scalars $x_1, x_2, y_1, y_2 \in \mathbb{Z}_q^*$ as private key. The public key is defined by the values $h = g_1^z$ and $c = g_1^{x_1} g_2^{x_2}$, $d = g_1^{y_1} g_2^{y_2}$. The public key pk is therefore (c, d, h) and the private key sk is (x_1, x_2, y_1, y_2, z) .
- $\mathcal{E}_{\text{pk}}(m; r)$: given a message $m \in \mathcal{G}$ and a random scalar $r \in \mathbb{Z}_q$, one computes

$$u_1 = g_1^r, u_2 = g_2^r, e = m \times h^r, v = (cd^\alpha)^r \quad \text{where } \alpha = \mathcal{H}(u_1, u_2, e).$$

- $\mathcal{D}_{\text{sk}}(u_1 \| u_2 \| e \| v)$: thanks to the private key, the decryption algorithm \mathcal{D}_{sk} first checks the validity of the ciphertext:

$$v \stackrel{?}{=} u_1^{x_1 + \alpha y_1} u_2^{x_2 + \alpha y_2} \quad \text{where } \alpha = \mathcal{H}(u_1, u_2, e).$$

If the equality holds, the plaintext is $m = e/u_1^z$, otherwise it returns “Reject.”

10.2.2 Security Analysis

About this construction, one can prove:

Theorem 10.1 *Let \mathcal{A} be a CCA-adversary against the semantic security of the above encryption scheme CS. Assume that \mathcal{A} has advantage ε and running time τ and makes q_d queries to the decryption oracle. Then*

$$\text{Adv}^{\text{ddh}}(T) \geq \frac{\varepsilon}{2} - \frac{1}{2} \times \text{Succ}^{\mathcal{H}}(T) - \frac{3q_d}{2q},$$

where $T = t + (10 + 4q_d)T_{\text{exp}}$, and T_{exp} the time for one exponentiation.

Proof : In the following we use starred letters (r^* , u_1^* , u_2^* , e^* , v^* , α^* , and c^*) to refer to the challenge ciphertext, whereas unstarred letters (r , u_1 , u_2 , e , v , α and c) refer to the ciphertext asked to the decryption oracle.

Game \mathbf{G}_0 : A pair of keys $(\mathbf{pk}, \mathbf{sk})$ is generated using $\mathcal{K}(1^k)$. Adversary A_1 is fed with \mathbf{pk} , and outputs a pair of messages (m_0, m_1) . Next a challenge ciphertext is produced by flipping a coin b and producing a ciphertext $c^* = u_1^* \| u_2^* \| e^* \| v^*$ of m_b . This ciphertext comes from a random $r^* \xleftarrow{R} \mathbb{Z}_q$ and $u_1^* = g_1^{r^*}$, $u_2^* = g_2^{r^*}$, $e^* = m_b \times h^{r^*}$, $v^* = (cd^{\alpha^*})^{r^*}$ where $\alpha^* = \mathcal{H}(u_1^*, u_2^*, e^*)$. On input c^* , A_2 outputs bit b' . In both stages, the adversary is given additional access to the decryption oracle $\mathcal{D}_{\mathbf{sk}}$. The only requirement is that the challenge ciphertext c^* cannot be queried from the decryption oracle.

We denote by S_0 the event $b' = b$ and use a similar notation S_i in any \mathbf{G}_i below. By definition, we have

$$\Pr[S_0] = \frac{1}{2} + \frac{\varepsilon}{2}. \quad (10.1)$$

\mathcal{K} Oracle	<p>One generates four random scalars $x_1, x_2, y_1, y_2 \in \mathbb{Z}_q^*$.</p> <p>► Rule GenSK⁽⁰⁾</p> <p style="padding-left: 20px;">One chooses a random scalar $z \in \mathbb{Z}_q$, and sets $h = g_1^z$.</p> <p>The private consists of all the random scalars. The public key is defined by h and $c = g_1^{x_1} g_2^{x_2}$, $d = g_1^{y_1} g_2^{y_2}$.</p>
\mathcal{D} Oracle	<p>Query $\mathcal{D}_{\mathbf{sk}}(u_1 \ u_2 \ e \ v)$:</p> <p>► Rule DecM⁽⁰⁾</p> <p style="padding-left: 20px;">One computes $m = e/h^z$.</p> <p>► Rule CheckV⁽⁰⁾</p> <p style="padding-left: 20px;">One computes $\alpha = \mathcal{H}(u_1, u_2, e)$, and $v' = u_1^{x_1 + \alpha y_1} u_2^{x_2 + \alpha y_2}$. If $v \neq v'$ then one returns “Reject.”</p> <p>One returns m.</p>
Challenger	<p>For two messages (m_0, m_1), flip a coin b and set $m^* = m_b$.</p> <p>► Rule Chal–Output⁽⁰⁾</p> <p style="padding-left: 20px;">Choose randomly r^*, then set $u_1^* = g_1^{r^*}$, $u_2^* = g_2^{r^*}$, $e^* = m^* \times h^{r^*}$, $v^* = (cd^{\alpha^*})^{r^*}$ where $\alpha^* = \mathcal{H}(u_1^*, u_2^*, e^*)$.</p> <p>Then, output $c^* = u_1^* \ u_2^* \ e^* \ v^*$.</p>

Figure 10.1: Formal Simulation of the IND-CCA Game against the CS Construction

Game \mathbf{G}_1 : In the previous game, we define an event \mathbf{CBad}_0 that is raised during a decryption querying $(u_1 \| u_2 \| e \| v)$ if $v = u_1^{x_1 + \alpha y_1} u_2^{x_2 + \alpha y_2}$, but $u_1 = g_1^{r_1}$ and $u_2 = g_2^{r_2}$ with $r_1 \neq r_2$. In this

new game, when this event is raised, one stops the game and outputs a random bit b' . Then, clearly,

$$\begin{aligned}
\Pr[\text{CBad}_1] &= \Pr[\text{CBad}_0] & \Pr[S_1 | \neg \text{CBad}_1] &= \Pr[S_0 | \neg \text{CBad}_0] & \Pr[S_1 | \text{CBad}_1] &= \frac{1}{2} \\
|\Pr[S_1] - \Pr[S_0]| &= |\Pr[S_1 | \neg \text{CBad}_1] \Pr[\neg \text{CBad}_1] + \Pr[S_1 | \text{CBad}_1] \Pr[\text{CBad}_1] \\
&\quad - \Pr[S_0 | \neg \text{CBad}_0] \Pr[\neg \text{CBad}_0] - \Pr[S_0 | \text{CBad}_0] \Pr[\text{CBad}_0]| \\
&= |\Pr[S_1 | \text{CBad}_1] - \Pr[S_0 | \text{CBad}_0]| \times \Pr[\text{CBad}_0] \\
&= \left| \frac{1}{2} - \Pr[S_0 | \text{CBad}_0] \right| \times \Pr[\text{CBad}_0]
\end{aligned}$$

As a consequence,

$$|\Pr[S_1] - \Pr[S_0]| \leq \frac{1}{2} \times \Pr[\text{CBad}_1] \quad (10.2)$$

Let us now evaluate $\Pr[\text{CBad}_1]$. This event is raised if, knowing the public parameters $c = g_1^{x_1} g_2^{x_2}$ and $d = g_1^{y_1} g_2^{y_2}$ and the challenge ciphertext $u_1^* = g_1^{r_1}$, $u_2^* = g_2^{r_2}$, and $v^* = (cd^{\alpha^*})^{r^*}$ where $\alpha^* = \mathcal{H}(u_1^*, u_2^*, e^*)$, the adversary generates $u_1 = g_1^{r_1}$, $u_2 = g_2^{r_2}$, and v such that $r_1 \neq r_2$ and $v = u_1^{x_1 + \alpha y_1} u_2^{x_2 + \alpha y_2}$.

If we denote (formally), $g_2 = g_1^\beta$, $c = g_1^\gamma$, and $d = g_1^\delta$, together with $v = g_1^\mu$ and $chv = g_1^{\mu^*}$ we have

$$\begin{aligned}
x_1 + \beta x_2 &= \gamma \\
y_1 + \beta y_2 &= \delta \\
r^*(x_1 + \alpha y_1) + \beta r^*(x_2 + \alpha y_2) &= \mu^* \\
r_1(x_1 + \alpha y_1) + \beta r_2(x_2 + \alpha y_2) &= \mu
\end{aligned}$$

By construction, the third equation is a linear combination of the two first one (μ^* can thus be uniquely determined). However, the fourth is linearly independent ($r_1 \neq r_2$), μ is unpredictable, and thus v is so too. As a consequence, even a powerful adversary has no chance to raise CBad_1 but by chance:

$$\Pr[\text{CBad}_0] = \Pr[\text{CBad}_1] \leq \frac{q_d}{q} \quad (10.3)$$

Game G_2 : In this game we modify the key generation, with

► **Rule GenSK⁽²⁾**

| One chooses two random scalars $z_1, z_2 \in \mathbb{Z}_q$, and sets $h = g_1^{z_1} g_2^{z_2}$.

and thus the decryption procedure becomes:

► **Rule DecM⁽²⁾**

| One computes $m = e / u_1^{z_1} u_2^{z_2}$.

Since we still abort when an acceptable ciphertext satisfies $r_1 \neq r_2$, where $u_1 = g_1^{r_1}$ and $u_2 = g_2^{r_2}$, then necessarily, a decryption only happens when $r = r_1 = r_2$: $u_1^{z_1} u_2^{z_2} = (g_1^{z_1} g_2^{z_2})^r = h^r$. If we formally denote $h = g_1^z$, which means that $z = z_1 + \beta z_2$, then $u_1^{z_1} u_2^{z_2} = (g_1^r)^{z_1 + \beta z_2} = u_1^z$, and thus the decryption is identical to the previous game.

$$\Pr[S_2] = \Pr[S_1] \quad (10.4)$$

Game G_3 : We now want to involve a **DDH** instance. Let us be given a tuple $(g_1, g_2, U = g_1^{r^{\star 1}}, V = g_2^{r^{\star 2}})$, with $r^{\star 1} = r^{\star 2}$. We use it for the challenge ciphertext generation:

► **Rule Chal–Output**⁽³⁾

$$\left| \begin{array}{l} \text{Set } u_1^* = U, u_2^* = V, e^* = m^* \times U^{z_1} V^{z_2}, v^* = U^{x_1 + \alpha^* y_1} V^{x_2 + \alpha^* y_2} \\ \text{where } \alpha^* = \mathcal{H}(u_1^*, u_2^*, e^*). \end{array} \right.$$

Since we assume that $r^* = r^{\star 1} = r^{\star 2}$, $u_1^* = g_1^{r^*}$, $u_2^* = g_2^{r^*}$, $e^* = m^* \times (g_1^{z_1} g_2^{z_2})^{r^*} = h^{r^*}$, and $v^* = g_1^{r^* x_1 + \alpha^* r^* y_1} g_2^{r^* x_2 + \alpha^* r^* y_2} = c^{r^*} d^{\alpha^* r^*}$ where $\alpha^* = \mathcal{H}(u_1^*, u_2^*, e^*)$. Hence, the challenge generation is identical to the previous game:

$$\Pr[S_3] = \Pr[S_2] \quad (10.5)$$

Game G_4 : Now, we would like to replace the Diffie-Hellman tuple by a random tuple, and use the **DDH** assumption. However, our current simulation is not polynomial: in order to detect the event **CBad** and then abort, one needs to be able to compute discrete logarithms (or at least to make the **DDH** decision). We thus forget this event, and do not abort anymore, even in case of wrong acceptable ciphertexts. Since v^* still leads to a linear combination of the exponents of the public key, we can make exactly the same analysis as in game G_1 :

$$|\Pr[S_4] - \Pr[S_3]| \leq \frac{q_d}{2q} \quad (10.6)$$

Game G_5 : We are now given a random input tuple $(g_1, g_2, U = g_1^{r^{\star 1}}, V = g_2^{r^{\star 2}})$. By simply outputting the boolean $b' = b$, we have a distinguisher against the **DDH** problem, within time T , which takes into account the time complexity t of the adversary, and the simulations of the key generation (6 exponentiations), the decryption oracle ($4q_d$ exponentiations) and the challenger (4 exponentiations):

$$|\Pr[S_5] - \Pr[S_4]| \leq \text{Adv}^{\text{ddh}}(t + (10 + 4q_d)T_{\text{exp}}) \quad (10.7)$$

Game G_6 : In order to be sure that no additional information is revealed about the secret key, we again abort the simulation when the event **CBad** is detected. Now, v^* does no longer leads to a linear combination of the exponents of the public key. We can thus take the first part of the analysis of the game G_1 :

$$|\Pr[S_6] - \Pr[S_5]| \leq \frac{1}{2} \times \Pr[\text{CBad}_6] \quad (10.8)$$

The analysis of the event **CBad** is a bit more intricate: the adversary knows $c = g_1^{x_1} g_2^{x_2}$ and $d = g_1^{y_1} g_2^{y_2}$ and the challenge ciphertext $u_1^* = g_1^{r^{\star 1}}$, $u_2^* = g_2^{r^{\star 2}}$, and $v^* = g_1^{r^{\star 1} x_1 + \alpha^* r^{\star 1} y_1} g_2^{r^{\star 2} x_2 + \alpha^* r^{\star 2} y_2}$ where $\alpha^* = \mathcal{H}(u_1^*, u_2^*, e^*)$ and wants to generate $u_1 = g_1^{r_1}$, $u_2 = g_2^{r_2}$, and v such that $r_1 \neq r_2$ (to learn something) and $v = u_1^{x_1 + \alpha y_1} u_2^{x_2 + \alpha y_2}$ (to be accepted).

If we denote (formally), $g_2 = g_1^\beta$, $c = g_1^\gamma$, and $d = g_1^\delta$, together with $v = g_1^\mu$ and $chv = g_1^{\mu^*}$ we have

$$\begin{aligned} x_1 + \beta x_2 &= \gamma \\ y_1 + \beta y_2 &= \delta \\ r^{\star 1}(x_1 + \alpha^* y_1) + \beta r^{\star 2}(x_2 + \alpha^* y_2) &= \mu^* \\ r_1(x_1 + \alpha y_1) + \beta r_2(x_2 + \alpha y_2) &= \mu \end{aligned}$$

A powerful adversary is able to compute $\alpha, \beta, \gamma, \delta, r^*_1, r^*_2$ and μ^* . Its goal is to generate $r_1 \neq r_2$ and μ that satisfies the appropriate equation with the variables “ x_1, x_2, y_1, y_2 ” used by the simulator. However, for any $r_1 \neq r_2$, we can show that the following determinant is

$$\begin{vmatrix} 1 & \beta & 0 & 0 \\ 0 & 0 & 1 & \beta \\ r^*_1 & \beta r^*_2 & r^*_1 \alpha^* & \beta r^*_2 \alpha^* \\ r_1 & \beta r_2 & r_1 \alpha & \beta r_2 \alpha \end{vmatrix} = \beta^2 \times (r^*_2 - r^*_1) \times (r_2 - r_1) \times (\alpha^* - \alpha).$$

But for a decryption query (u_1, u_2, e, α) , three cases appear:

- $(u_1, u_2, e) = (u_1^*, chu_2, e^*)$, then necessarily $chv \neq v$ otherwise the query is exactly the challenge ciphertext, and this is not allowed. Then, we know that the verification check will not pass, the decryption will be rejected;
- $(u_1, u_2, e) \neq (u_1^*, chu_2, e^*)$, but $\alpha^* = \alpha$, which means that $\mathcal{H}(u_1, u_2, e) = \mathcal{H}(u_1^*, chu_2, e^*)$, and then the adversary has found a second pre-image for \mathcal{H} ;
- $(u_1, u_2, e) \neq (u_1^*, chu_2, e^*)$, and $\alpha^* \neq \alpha$. If $r^*_1 \neq r^*_2$, then the above determinant is non-zero, and thus the value μ is unpredictable.

As a consequence, even a powerful adversary has no chance to raise CBad_6 but by chance:

$$\Pr[\text{CBad}_5] = \Pr[\text{CBad}_6] \leq \frac{q_d}{q} + \text{Succ}^{\mathcal{H}}(t + (10 + 4q_d)T_{\text{exp}}) \quad (10.9)$$

Furthermore, in this last game the challenge ciphertext contains $e^* = m^* \times U^{z_1} V^{z_2}$, and the public key $h = g_1^{z_1} g_2^{z_2}$ just reveals a linear relation between z_1 and z_2 , and any decryption queries for valid ciphertexts where $r_1 = r_2$. On the opposite, since $U = g_1^{r^*_1}$ and $V = g_2^{r^*_2}$ with $r^*_1 \neq r^*_2$, even a powerful adversary has no information about the value of $U^{z_1} V^{z_2}$, hence b' is independent of b :

$$\Pr[\text{S}_6] = \frac{1}{2} \quad (10.10)$$

As a conclusion, one can see that

$$\begin{aligned} \Pr[\text{S}_0] &= \frac{1}{2} + \frac{\varepsilon}{2} \\ |\Pr[\text{S}_1] - \Pr[\text{S}_0]| &\leq \frac{q_d}{2q} \\ |\Pr[\text{S}_2] - \Pr[\text{S}_1]| &= 0 \\ |\Pr[\text{S}_3] - \Pr[\text{S}_2]| &= 0 \\ |\Pr[\text{S}_4] - \Pr[\text{S}_3]| &\leq \frac{q_d}{2q} \\ |\Pr[\text{S}_5] - \Pr[\text{S}_4]| &\leq \text{Adv}^{\text{ddh}}(T) \\ |\Pr[\text{S}_6] - \Pr[\text{S}_5]| &\leq \frac{q_d}{2q} + \frac{1}{2} \times \text{Succ}^{\mathcal{H}}(T) \\ \Pr[\text{S}_6] &= \frac{1}{2} \end{aligned}$$

where $T = t + (10 + 4q_d)T_{\text{exp}}$, and thus

$$\frac{\varepsilon}{2} = |\Pr[\text{S}_6] - \Pr[\text{S}_0]| \leq \frac{3q_d}{2q} + \text{Adv}^{\text{ddh}}(T) + \frac{1}{2} \times \text{Succ}^{\mathcal{H}}(T)$$

10.3 A Generic Construction

In [?], Bellare and Rogaway proposed the first generic construction which applies to any trapdoor one-way permutation f onto X .

10.3.1 Description

We need two hash functions \mathcal{G} and \mathcal{H} :

$$\mathcal{G} : X \longrightarrow \{0, 1\}^n \quad \text{and} \quad \mathcal{H} : \{0, 1\}^* \longrightarrow \{0, 1\}^{k_1},$$

where n is the bit-length of the plaintexts, and k_1 a security parameter. Then the encryption scheme $\text{BR} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ can be described as follows:

- $\mathcal{K}(1^k)$: specifies an instance of the function f , and of its inverse f^{-1} . The public key pk is therefore f and the private key sk is f^{-1} .
- $\mathcal{E}_{\text{pk}}(m; r)$: given a message $m \in \{0, 1\}^n$, and a random value $r \xleftarrow{R} X$, the encryption algorithm \mathcal{E}_{pk} computes

$$a = f(r), \quad b = m \oplus \mathcal{G}(r) \quad \text{and} \quad c = \mathcal{H}(m, r),$$

and outputs the ciphertext $y = a \| b \| c$.

- $\mathcal{D}_{\text{sk}}(a \| b \| c)$: thanks to the private key, the decryption algorithm \mathcal{D}_{sk} extracts

$$r = f^{-1}(a), \quad \text{and next} \quad m = b \oplus \mathcal{G}(r).$$

If $c = \mathcal{H}(m, r)$, the algorithm returns m , otherwise it returns “Reject.”

10.3.2 Security Analysis

About this construction, one can prove:

Theorem 10.2 *Let \mathcal{A} be a CCA-adversary against the semantic security of the above encryption scheme BR . Assume that \mathcal{A} has advantage ε and running time τ and makes q_d , q_g and q_h queries to the decryption oracle, and the hash functions \mathcal{G} and \mathcal{H} , respectively. Then*

$$\begin{aligned} \text{Succ}_f^{\text{ow}}(\tau') &\geq \frac{\varepsilon}{2} - \frac{2q_d}{2^{k_1}}, \\ \text{with } \tau' &\leq \tau + (q_g + q_h) \cdot T_f, \end{aligned}$$

where T_f denotes the time complexity for evaluating f .

Proof : In the following we use starred letters (r^* , a^* , b^* , c^* and y^*) to refer to the challenge ciphertext, whereas unstarred letters (r , a , b , c and y) refer to the ciphertext asked to the decryption oracle.

Game \mathbf{G}_0 : A pair of keys (pk, sk) is generated using $\mathcal{K}(1^k)$. Adversary A_1 is fed with pk , the description of f , and outputs a pair of messages (m_0, m_1) . Next a challenge ciphertext is produced by flipping a coin b and producing a ciphertext $y^* = a^* \| b^* \| c^*$ of m_b . This ciphertext

comes from a random $r^* \xleftarrow{R} X$ and $a^* = f(r^*)$, $b^* = m_b \oplus \mathcal{G}(r^*)$ and $c^* = \mathcal{H}(m_b, r^*)$. On input y^* , A_2 outputs bit b' . In both stages, the adversary is given additional access to the decryption oracle \mathcal{D}_{sk} . The only requirement is that the challenge ciphertext y^* cannot be queried from the decryption oracle.

We denote by S_0 the event $b' = b$ and use a similar notation S_i in any \mathbf{G}_i below. By definition, we have

$$\Pr[S_0] = \frac{1}{2} + \frac{\varepsilon}{2}. \quad (10.11)$$

Game \mathbf{G}_1 : In this game, one makes the classical simulation of the random oracles, with random answers for any new query, as shown on Figure 10.2. This game is clearly identical to the previous one.

\mathcal{G}, \mathcal{H} Oracles	<p>Query $\mathcal{G}(r)$: if a record (r, g) appears in G-List, the answer is g. Otherwise the answer g is chosen randomly: $g \in \{0, 1\}^n$ and the record (r, g) is added in G-List.</p> <hr/> <p>Query $\mathcal{H}(m, r)$: if a record (m, r, h) appears in H-List, the answer is h. Otherwise the answer h is chosen randomly: $h \in \{0, 1\}^{k_1}$ and the record (m, r, h) is added in H-List.</p>
\mathcal{D} Oracle	<p>Query $\mathcal{D}_{\text{sk}}(a b c)$: one applies the following rules:</p> <p style="padding-left: 20px;">► Rule Decrypt–R⁽¹⁾ Compute $r = f^{-1}(a)$;</p> <p style="padding-left: 20px;">Then, compute $m = b \oplus \mathcal{G}(r)$, and finally,</p> <p style="padding-left: 20px;">► Rule Decrypt–H⁽¹⁾ If $c = \mathcal{H}(m, r)$, one returns m, otherwise one returns “Reject.”</p>
Challenger	<p>For two messages (m_0, m_1), flip a coin b and set $m^* = m_b$.</p> <p style="padding-left: 20px;">► Rule Chal–Hash⁽¹⁾ Choose randomly r^*, then set $a^* = f(r^*)$, $g^* = \mathcal{G}(r^*)$, $b^* = m^* \oplus g^*$, $c^* = \mathcal{H}(m^*, r^*)$.</p> <p style="padding-left: 20px;">Then, output $y^* = a^* b^* c^*$.</p>

Figure 10.2: Formal Simulation of the IND-CCA Game against the BR Construction

Game \mathbf{G}_2 : In this game, one randomly chooses $h^+ \xleftarrow{R} \{0, 1\}^{k_1}$, and uses it instead of $\mathcal{H}(m^*, r^*)$.

► **Rule Chal–Hash⁽²⁾**

The value $h^+ \xleftarrow{R} \{0, 1\}^{k_1}$ has been chosen ahead of time, choose randomly r^* , then set $a^* = f(r^*)$, $g^* = \mathcal{G}(r^*)$, $b^* = m^* \oplus g^*$, and $c^* = h^+$.

The two games \mathbf{G}_2 and \mathbf{G}_1 are perfectly indistinguishable unless (m^*, r^*) is asked for \mathcal{H} , either by the adversary or the decryption oracle. But the latter case is not possible, otherwise the decryption query would be the challenge ciphertext. More generally, we denote by AskR_2 the event that r^* has been asked to \mathcal{G} or to \mathcal{H} , by the adversary. We have:

$$|\Pr[S_2] - \Pr[S_1]| \leq \Pr[\text{AskR}_2]. \quad (10.12)$$

Game \mathbf{G}_3 : We start modifying the simulation of the decryption oracle, by rejecting any ciphertext $(a||b||c)$ for which the corresponding (m, r) has not been queried to \mathcal{H} :

► **Rule Decrypt–H⁽³⁾**

Look up in H-List for (m, r, c) . If such a triple does not exist, then output “Reject”, otherwise output m .

Such a simulation differs from the previous one if the value c has been correctly guessed, by chance:

$$|\Pr[S_3] - \Pr[S_2]| \leq \frac{q_d}{2^{k_1}} \quad |\Pr[\text{AskR}_3] - \Pr[\text{AskR}_2]| \leq \frac{q_d}{2^{k_1}}. \quad (10.13)$$

Game \mathbf{G}_4 : In this game, one randomly chooses $r^+ \xleftarrow{R} X$ and $g^+ \xleftarrow{R} \{0, 1\}^n$, and uses r^+ instead of r^* , as well as g^+ instead of $\mathcal{G}(r^*)$.

► **Rule Chal–Hash⁽⁴⁾**

The three values $r^+ \xleftarrow{R} X$, $g^+ \xleftarrow{R} \{0, 1\}^n$ and $h^+ \xleftarrow{R} \{0, 1\}^{k_1}$ have been chosen ahead of time, then set $a^* = f(r^+)$, $b^* = m^* \oplus g^+$, $c^* = h^+$.

The two games \mathbf{G}_4 and \mathbf{G}_3 are perfectly indistinguishable unless r^* is asked for \mathcal{G} , either by the adversary or the decryption oracle. The former case has already been cancelled in the previous game, in AskR_3 . The latter case does not make any difference since either $\mathcal{H}(m, r^*)$ has been queried by the adversary, which falls in AskR_3 , or the ciphertext is rejected in both games. We have:

$$\Pr[S_4] = \Pr[S_3] \quad \Pr[\text{AskR}_4] = \Pr[\text{AskR}_3]. \quad (10.14)$$

In this game, m^* is masked by g^+ , a random value which never appears anywhere else. Thus, the input to \mathcal{A}_2 follows a distribution that does not depend on b . Accordingly:

$$\Pr[S_4] = \frac{1}{2}. \quad (10.15)$$

Game \mathbf{G}_5 : Finally, one randomly chooses $a^+ \xleftarrow{R} X$, which implicitly defines a random r^+ in X . Actually, a^+ is the given random challenge for which one is looking for the pre-image r^+ .

► **Rule Chal–Hash⁽⁵⁾**

The three values $a^+ \xleftarrow{R} X$, $g^+ \xleftarrow{R} \{0, 1\}^n$ and $h^+ \xleftarrow{R} \{0, 1\}^{k_1}$ have been chosen/given ahead of time, then set $a^* = a^+$, $b^* = m^* \oplus g^+$, $c^* = h^+$.




Figure 10.3: Optimal Asymmetric Encryption Padding

The two games \mathbf{G}_5 and \mathbf{G}_4 are perfectly indistinguishable, thanks to the permutation property of f .

One may now note that the event AskR_5 leads to the pre-image of a^+ by f in the queries asked to \mathcal{G} and \mathcal{H} , by the adversary. By checking all of them, one gets it:

$$\Pr[\text{AskR}_5] \leq \text{Succ}_f^{\text{ow}}(\tau + (q_g + q_h)T_f). \quad (10.16)$$

10.4 OAEP: the Optimal Asymmetric Encryption Padding

10.4.1 Description

The problem with the above generic construction is the high over-head. When one encrypts with a trapdoor one-way permutation onto X , one could hope the ciphertext to be an element in X , without anything else. In 1994, Bellare and Rogaway proposed such a more compact generic conversion [?], in the random-oracle model, the “Optimal Asymmetric Encryption Padding” (OAEP, see Figure 10.3), obtained from a trapdoor one-way permutation f onto $\{0, 1\}^k$, whose inverse is denoted by f^{-1} . We need two hash functions \mathcal{G} and \mathcal{H} :

$$\mathcal{G} : \{0, 1\}^{k_0} \longrightarrow \{0, 1\}^{k-k_0} \quad \text{and} \quad \mathcal{H} : \{0, 1\}^{k-k_0} \longrightarrow \{0, 1\}^{k_0},$$

for some k_0 . We also need n and k_1 which satisfy $k = n + k_0 + k_1$. Then the encryption scheme $\text{OAEP} = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ can be described as follows:

- $\mathcal{K}(1^k)$: specifies an instance of the function f , and of its inverse f^{-1} . The public key pk is therefore f and the private key sk is f^{-1} .
- $\mathcal{E}_{\text{pk}}(m; r)$: given a message $m \in \{0, 1\}^n$, and a random value $r \xleftarrow{R} \{0, 1\}^{k_0}$, the encryption algorithm \mathcal{E}_{pk} computes

$$s = (m \| 0^{k_1}) \oplus \mathcal{G}(r) \quad \text{and} \quad t = r \oplus \mathcal{H}(s),$$

and outputs the ciphertext $c = f(s, t)$.

- $\mathcal{D}_{\text{sk}}(c)$: thanks to the private key, the decryption algorithm \mathcal{D}_{sk} extracts

$$(s, t) = f^{-1}(c), \quad \text{and next} \quad r = t \oplus \mathcal{H}(s) \quad \text{and} \quad M = s \oplus \mathcal{G}(r).$$

If $[M]_{k_1} = 0^{k_1}$, the algorithm returns $[M]^n$, otherwise it returns “Reject.”

In the above description, $[M]_{k_1}$ denotes the k_1 least significant bits of M , while $[M]^n$ denotes the n most significant bits of M .

10.4.2 About the Security

Paper [?] includes a proof that, provided f is a one-way trapdoor permutation, the resulting OAEP encryption scheme is both semantically secure and weakly plaintext-aware. This implies the semantic security against indifferent chosen-ciphertext attacks, also called security against lunchtime attacks (IND-CCA1). Indeed, the *Weak Plaintext-Awareness* means that the adversary cannot produce a new valid ciphertext, until it has seen any valid one, without knowing (awareness) the plaintext. This is more formally defined by the existence of a plaintext-extractor which, on input a ciphertext and the list of the query-answers of the random oracles, outputs the corresponding plaintext. This plaintext-extractor is thus enough for simulating the decryption oracle, but in the first step of the attack only. We briefly comment on the intuition behind (weak) plaintext-awareness. When the plaintext-extractor receives a ciphertext c , then:

- either s has been queried to \mathcal{H} and r has been queried to \mathcal{G} , in which case the extractor finds the cleartext by inspecting the two query lists G-List and H-List,
- or else the decryption of (s, t) remains highly random and there is little chance to meet the redundancy 0^{k_1} : the plaintext extractor can safely declare the ciphertext invalid.

The argument collapses when the plaintext-extractor receives additional valid ciphertexts, since this puts additional implicit constraints on \mathcal{G} and \mathcal{H} . These constraints cannot be seen by inspecting the query lists. Hence the requirement of a stronger notion of *plaintext-awareness*. In [?], Bellare, Desai, Pointcheval, and Rogaway defined such a stronger notion which extends the previous *awareness* of the plaintext even after having seen valid ciphertexts. But such a plaintext-awareness notion had never been studied for OAEP, while it was still widely admitted.

10.4.2.1 Shoup's Counter-Example

In his papers [?, ?], Shoup showed that it was quite unlikely to extend the results of [?] to obtain adaptive chosen-ciphertext security, under the sole one-wayness of the permutation. His counter-example made use of the ad hoc notion of an *XOR-malleable* trapdoor one-way permutation: for such permutation f_0 , one can compute $f_0(x \oplus a)$ from $f_0(x)$ and a , with non-negligible probability.

Let f_0 be such an XOR-malleable permutation. Define f by $f(s||t) = s||f_0(t)$. Clearly, f is also a trapdoor one-way permutation. However, it leads to a malleable encryption scheme as we now show. Start with a challenge ciphertext $y = f(s||t) = s||u$, where $s||t$ is the output of the OAEP transformation on the redundant message $m||0^{k_1}$ and the random string r (see Figure 10.4),

$$s = \mathcal{G}(r) \oplus (m||0^{k_1}), \quad t = \mathcal{H}(s) \oplus r \quad \text{and} \quad u = f_0(t).$$

Since f is the identity on its leftmost part, we know s , and can define $\Delta = \delta||0^{k_1}$, for any random string δ , and $s' = s \oplus \Delta$. We then set $t' = \mathcal{H}(s') \oplus r = t \oplus (\mathcal{H}(s) \oplus \mathcal{H}(s'))$. The XOR-malleability of f_0 allows one to obtain $u' = f_0(t')$ from $u = f_0(t)$ and $\mathcal{H}(s) \oplus \mathcal{H}(s')$,




Figure 10.4: Shoup's attack.

with significant probability. Finally, $y' = s' \| u'$ is a valid ciphertext of $m' = m \oplus \delta$, built from $r' = r$, since:

$$t' = f_0^{-1}(u') = t \oplus (\mathcal{H}(s) \oplus \mathcal{H}(s')) = \mathcal{H}(s') \oplus r, \quad r' = \mathcal{H}(s') \oplus t' = r$$

and

$$s' \oplus \mathcal{G}(r') = \Delta \oplus s \oplus \mathcal{G}(r) = \Delta \oplus (m \| 0^{k_1}) = (m \oplus \delta) \| 0^{k_1}.$$

Note that the above definitely contradicts adaptive chosen-ciphertext security: asking the decryption of y' after having received the ciphertext y , an adversary obtains m' and easily recovers the actual cleartext m from m' and δ . Also note that Shoup's counter-example exactly stems from where the intuition developed at the end of the previous section failed: a valid ciphertext y' was created without querying the oracle at the corresponding random seed r' , using in place the implicit constraint on \mathcal{G} coming from the received valid ciphertext y .

Using methods from relativized complexity theory, Shoup [?, ?] built a non-standard model of computation, where there exists an XOR-malleable trapdoor one-way permutation. As a consequence, it is very unlikely that one can prove the IND-CCA security of the OAEP construction, under the sole one-wayness of the underlying permutation. Indeed, all methods of proof currently known still apply in relativized models of computation.

10.4.3 The Actual Security of OAEP

Shoup [?, ?] furthermore provided a specific proof for RSA with public exponent 3. However, there is little hope of extending this proof for higher exponents. Hopefully, Fujisaki, Okamoto, Pointcheval, and Stern provided a general security analysis, but under a stronger assumption about the underlying permutation [?, ?]. Indeed, they prove that the scheme is IND-CCA in the random-oracle model [?], relative to the *partial-domain* one-wayness of permutation f .

10.4.3.1 Partial-Domain One-Wayness

Let us first introduce this new computational assumption. Let f be a permutation $f : \{0, 1\}^k \rightarrow \{0, 1\}^k$, which can also be written as

$$f : \{0, 1\}^{n+k_1} \times \{0, 1\}^{k_0} \rightarrow \{0, 1\}^{n+k_1} \times \{0, 1\}^{k_0},$$

with $k = n + k_0 + k_1$. In the original description of OAEP from [?], it is only required that f is a trapdoor one-way permutation. However, in the following, we consider two additional related problems, namely partial-domain one-wayness and set partial-domain one-wayness:

- Permutation f is (τ, ε) -one-way if any adversary \mathcal{A} whose running time is bounded by τ has success probability $\text{Succ}_f^{\text{ow}}(\mathcal{A})$ upper-bounded by ε , where

$$\text{Succ}_f^{\text{ow}}(\mathcal{A}) = \Pr_{s,t}[\mathcal{A}(f(s, t)) = (s, t)].$$

- Permutation f is (τ, ε) -partial-domain one-way if any adversary \mathcal{A} whose running time is bounded by τ has success probability $\text{Succ}_f^{\text{pd-ow}}(\mathcal{A})$ upper-bounded by ε , where

$$\text{Succ}_f^{\text{pd-ow}}(\mathcal{A}) = \Pr_{s,t}[\mathcal{A}(f(s, t)) = s].$$

- Permutation f is $(\ell, \tau, \varepsilon)$ -set partial-domain one-way if any adversary \mathcal{A} , outputting a set of ℓ elements within time bound τ , has success probability $\text{Succ}_f^{\text{s-pd-ow}}(\mathcal{A})$ upper-bounded by ε , where

$$\text{Succ}_f^{\text{s-pd-ow}}(\mathcal{A}) = \Pr_{s,t}[s \in \mathcal{A}(f(s, t))].$$

We denote by $\text{Succ}_f^{\text{ow}}(\tau)$ (resp. $\text{Succ}_f^{\text{pd-ow}}(\tau)$ and $\text{Succ}_f^{\text{s-pd-ow}}(\ell, \tau)$) the maximal success probability $\text{Succ}_f^{\text{ow}}(\mathcal{A})$ (resp. $\text{Succ}_f^{\text{pd-ow}}(\mathcal{A})$ and $\text{Succ}_f^{\text{s-pd-ow}}(\mathcal{A})$). The maximum ranges over all adversaries whose running time is bounded by τ . In the third case, there is an obvious additional restriction on this range from the fact that \mathcal{A} outputs sets with ℓ elements. It is clear that for any τ and $\ell \geq 1$,

$$\text{Succ}_f^{\text{s-pd-ow}}(\ell, \tau) \geq \text{Succ}_f^{\text{pd-ow}}(\tau) \geq \text{Succ}_f^{\text{ow}}(\tau).$$

Note that, by randomly selecting an element in the set returned by an adversary to the set partial-domain one-wayness, one breaks partial-domain one-wayness with probability $\text{Succ}_f^{\text{s-pd-ow}}(\mathcal{A})/\ell$. This provides the following inequality

$$\text{Succ}_f^{\text{pd-ow}}(\tau) \geq \text{Succ}_f^{\text{s-pd-ow}}(\ell, \tau)/\ell.$$

However, for specific choices of f , more efficient reductions may exist. Also, in some cases, all three problems are polynomially equivalent. This is the case for the RSA permutation [?], hence the global security result for RSA-OAEP.

10.4.4 Intuition behind the Proof of Security

In the following we use starred letters (r^* , s^* , t^* and y^*) to refer to the challenge ciphertext, whereas unstarred letters (r , s , t and y) refer to the ciphertext asked to the decryption oracle.

Referring to our description of the intuition behind the original OAEP proof of security, given above, we can carry a more subtle analysis by distinguishing the case where s has not been queried from oracle \mathcal{H} from the case where r has not been queried from \mathcal{G} . If s is not queried, then $\mathcal{H}(s)$ is random and uniformly distributed and r is necessarily defined as $t \oplus \mathcal{H}(s)$. This holds even if s matches with the string s^* coming from the valid ciphertext y^* . There is a minute probability that $t \oplus \mathcal{H}(s)$ is queried from \mathcal{G} or equals r^* . Thus, $\mathcal{G}(r)$ is random: there is little chance that the redundancy 0^{k_1} is met and the extractor can safely reject.

We claim that r cannot match with r^* , unless s^* is queried from \mathcal{H} . This is because $r^* = t^* \oplus \mathcal{H}(s^*)$ equals $r = t \oplus \mathcal{H}(s)$ with minute probability. Thus, if r is not queried, then

$\mathcal{G}(r)$ is random and we similarly infer that the extractor can safely reject. The argument fails only if s^* is queried.

Thus rejecting when it cannot combine elements of the lists G-List and H-List so as to build a pre-image of y , the plaintext-extractor is only wrong with minute probability, unless s^* has been queried by the adversary. This seems to show that OAEP leads to an IND-CCA encryption scheme if it is difficult to invert f “partially”, which means: given $y^* = f(s^* \| t^*)$, find s^* .

Chosen-ciphertext security is actually addressed, by turning the intuition explained above into a formal argument, involving a restricted variant of plaintext-awareness (where the list C of ciphertexts is limited to only one ciphertext, the challenge ciphertext y^*):

Theorem 10.3 *Let \mathcal{A} be a CCA-adversary against the semantic security of the encryption scheme OAEP. Assume that \mathcal{A} has advantage ε and running time τ and makes q_d , q_g and q_h queries to the decryption oracle, and the hash functions \mathcal{G} and \mathcal{H} , respectively. Then*

$$\begin{aligned} \text{Succ}_f^{\text{s-pd-ow}}(q_h, \tau') &\geq \frac{\varepsilon}{2} - \left(\frac{2(q_d + 2)(q_d + 2q_g)}{2^{k_0}} + \frac{3q_d}{2^{k_1}} \right), \\ \text{with } \tau' &\leq \tau + q_g \cdot q_h \cdot (T_f + \mathcal{O}(1)), \end{aligned}$$

where T_f denotes the time complexity for evaluating f .

Unfortunately, because of the additional reduction of the basic RSA to the partial-domain RSA problem, the global reduction is very expensive, and is thus meaningful for huge moduli only, more than 4096-bit long. Indeed, the RSA inverter we can build, thanks to the full reduction, has a complexity at least greater than $q_h \cdot (q_h + 2q_g) \times \mathcal{O}(k^3)$. As already remarked, the adversary can ask up to 2^{60} queries to the hash functions, and thus this overhead in the inversion is at least 2^{151} . However, current factoring algorithms can factor up to 4096 bit-long integers within this number of basic operations (see [?] for complexity estimates of the most efficient factoring algorithms).

Anyway, the formal proof shows that the global design of OAEP is sound, and that it is still probably safe to use it in practice (*e.g.* in PKCS #1 v2.0, while being very careful during the implementation [?]).

Chapter 11

Digital Signature Schemes

11.1 Introduction

Until 1996, no practical **DL**-based cryptographic scheme has ever been formally studied, but heuristically only. And surprisingly, at the Eurocrypt '96 conference, two opposite studies were conducted on the El Gamal signature scheme [?], the first **DL**-based signature scheme designed in 1985 and depicted on Figure 11.1.

Initialization $\rightarrow (p, g)$
g a generator of \mathbb{Z}_p^* , where p is a large prime $\rightarrow (p, g)$
K: Key Generation $\rightarrow (y, x)$
private key $x \in \mathbb{Z}_{p-1}^*$ public key $y = g^x \bmod p$ $\rightarrow (y, x)$
S: Signature of m $\rightarrow (r, s)$
K is randomly chosen in \mathbb{Z}_{p-1}^* $r = g^K \bmod p$ $s = (m - xr)/K \bmod p - 1$ $\rightarrow (r, s)$ is a signature of m
V: Verification of (m, r, s)
check whether $g^m \stackrel{?}{=} y^r r^s \bmod p$ $\rightarrow \text{Yes/No}$

Figure 11.1: The El Gamal Signature Scheme.

Whereas existential forgeries were known for that scheme, it was believed to prevent universal forgeries. The first analysis, from Daniel Bleichenbacher [?], showed such a universal forgery when the generator g is not properly chosen. The second one, from David Pointcheval and Jacques Stern [?], proved the security against existential forgeries under adaptive chosen-message attacks of a slight variant with a randomly chosen generator g . The latter variant simply replaces the message m by $\mathcal{H}(m, r)$ in the computation, while one uses a hash function \mathcal{H} that is assumed to behave like a random oracle. It is amazing to remark that the Bleichenbacher's attack also applies on Pointcheval-Stern's variant. Therefore, depending on

the initialization, the variant could be a very strong signature scheme or become a very weak one!

As a consequence, a proof has to be performed in details, with precise assumptions and achievements. Furthermore, the conclusions have to be strictly followed by developers, otherwise the concrete implementation of a secure scheme can be very weak.

11.2 Some Schemes

The first *secure* signature scheme was proposed by Goldwasser *et al.* [?] in 1984. It used the notion of claw-free permutations. A pair of permutations (f, g) is said *claw-free* if it is computationally impossible to find a *claw* (x, y) , which satisfies $f(x) = g(y)$. Their proposal provided polynomial algorithms with a polynomial reduction between the research of a claw and an existential forgery under an adaptive chosen-message attack. However, the scheme was totally unpractical. What about practical schemes?

11.2.1 The RSA Signature Scheme

Two years after the Diffie-Hellman paper [?], Rivest, Shamir and Adleman [?] proposed the first signature scheme based on the “trapdoor one-way permutation paradigm”, using the RSA function: the generation algorithm produces a large composite number $N = pq$, a public key e , and a private key d such that $e \cdot d = 1 \bmod \varphi(N)$. The signature of a message m , encoded as an element in \mathbb{Z}_N^* , is its e -th root, $\sigma = m^{1/e} = m^d \bmod N$. The verification algorithm simply checks whether $m = \sigma^e \bmod N$.

However, the RSA scheme is not secure by itself since it is subject to existential forgery: it is easy to create a valid message-signature pair, without any help of the signer, first randomly choosing a certificate σ and getting the signed message m from the public verification relation, $m = \sigma^e \bmod N$.

11.2.2 The Schnorr Signature Scheme

In 1986 a new paradigm for signature schemes was introduced. It is derived from fair zero-knowledge identification protocols involving a prover and a verifier [?], and uses hash functions in order to create a kind of virtual verifier. The first application was derived from the Fiat-Shamir [?] zero-knowledge identification protocol, based on the hardness of extracting square roots, with a brief outline of its security. Another famous identification scheme [?], together with the signature scheme [?], has been proposed later by Schnorr, based on that paradigm: the generation algorithm produces two large primes p and q , such that $q \geq 2^k$, where k is the security parameter, and $q \mid p - 1$, as well as an element g in \mathbb{Z}_p^* of order q . It also creates a pair of keys, the private key $x \in \mathbb{Z}_q^*$ and the public key $y = g^{-x} \bmod p$. The signature of a message m is a triple (r, e, s) , where $r = g^K \bmod p$, with a random $K \in \mathbb{Z}_q$, the “challenge” $e = \mathcal{H}(m, r)$ and $s = K + ex \bmod q$. The latter satisfies $r = g^s y^e \bmod p$ with $e = \mathcal{H}(m, r)$, which is checked by the verification algorithm.

The security results for that paradigm have been considered as folklore for a long time but without any formal validation.




Figure 11.2: The Oracle Replay Technique

Initialization (security parameter k) $\rightarrow (\mathcal{G}, g, \mathcal{H})$
\mathbf{g} a generator of any cyclic group $(\mathcal{G}, +)$ of order q , with $2^{k-1} \leq q < 2^k$
\mathcal{H} a hash function: $\{0, 1\}^* \rightarrow \mathbb{Z}_q$
$\rightarrow (\mathcal{G}, g, \mathcal{H})$
\mathcal{K}: Key Generation $\rightarrow (\mathbf{y}, x)$
private key $x \in \mathbb{Z}_q^*$
public key $\mathbf{y} = -x \cdot \mathbf{g}$
$\rightarrow (\mathbf{y}, x)$
\mathcal{S}: Signature of m $\rightarrow (\mathbf{r}, h, s)$
K is randomly chosen in \mathbb{Z}_q^*
$\mathbf{r} = K \cdot \mathbf{g}$ $h = \mathcal{H}(m, \mathbf{r})$ $s = K + xh \bmod q$
$\rightarrow (\mathbf{r}, h, s)$ is a signature of m
\mathcal{V}: Verification of (m, \mathbf{r}, s)
check whether $h \stackrel{?}{=} \mathcal{H}(m, \mathbf{r})$
and $\mathbf{r} \stackrel{?}{=} s \cdot \mathbf{g} + h \cdot \mathbf{y}$
$\rightarrow \text{Yes/No}$

Figure 11.3: The Schnorr Signature Scheme.

11.3 DL-Based Signatures

In [?, ?], David Pointcheval and Jacques Stern formally proved the above paradigm when \mathcal{H} is assumed to behave like a random oracle. The proof is based on the by now classical *oracle replay technique*: by a polynomial replay of the attack with different random oracles (the \mathcal{Q}_i 's are the queries and the ρ_i 's are the answers), one make the attacker forge signatures that are suitably related. This generic technique is depicted on Figure 11.2, where the signature of a message m is a triple (σ_1, h, σ_2) , with $h = \mathcal{H}(m, \sigma_1)$ which depends on the message and the first part of the signature, both bound not to change for the computation of σ_2 , which really relies on the knowledge of the private key. If the probability of fraud is high enough, then with good probability, the adversary is able to answer to many distinct outputs from the \mathcal{H} function, on the input (m, σ_1) .

To be more concrete, let us consider the Schnorr signature scheme, which is presented on Figure 11.3, in any “suitable” cyclic group \mathcal{G} of prime order q , where at least the Discrete Logarithm problem is hard. We expect to obtain two signatures $(\mathbf{r} = \sigma_1, h, s = \sigma_2)$ and $(\mathbf{r}' = \sigma'_1, h', s' = \sigma'_2)$ of an identical message m such that $\sigma_1 = \sigma'_1$, but $h \neq h'$. Thereafter, we

can easily extract the discrete logarithm of the public key:

$$\left. \begin{aligned} \mathbf{r} &= s \cdot \mathbf{g} + h \cdot \mathbf{y} \\ \mathbf{r} &= s' \cdot \mathbf{g} + h' \cdot \mathbf{y} \end{aligned} \right\} \Rightarrow (s - s') \cdot \mathbf{g} = (h' - h) \cdot \mathbf{y},$$

which leads to $\log_{\mathbf{g}} \mathbf{y} = (s - s') \cdot (h' - h)^{-1} \bmod q$.

11.3.1 General Tools

First, let us recall the “Splitting Lemma” which will be the main probabilistic tool for the “Forking Lemma”. It translates the fact that when a subset A is “large” in a product space $X \times Y$, it has many “large” sections.

Lemma 11.1 (The Splitting Lemma) *Let $A \subset X \times Y$ such that $\Pr[(x, y) \in A] \geq \varepsilon$. For any $\alpha < \varepsilon$, define*

$$B = \left\{ (x, y) \in X \times Y \mid \Pr_{y' \in Y}[(x, y') \in A] \geq \varepsilon - \alpha \right\},$$

then the following statements hold:

- (i) $\Pr[B] \geq \alpha$
- (ii) $\forall (x, y) \in B, \Pr_{y' \in Y}[(x, y') \in A] \geq \varepsilon - \alpha$.
- (iii) $\Pr[B \mid A] \geq \alpha/\varepsilon$.

Proof : In order to prove statement (i), we argue by contradiction, using the notation \bar{B} for the complement of B in $X \times Y$. Assume that $\Pr[B] < \alpha$. Then

$$\varepsilon \leq \Pr[B] \cdot \Pr[A \mid B] + \Pr[\bar{B}] \cdot \Pr[A \mid \bar{B}] < \alpha \cdot 1 + 1 \cdot (\varepsilon - \alpha) = \varepsilon.$$

This implies a contradiction, hence the result.

Statement (ii) is a straightforward consequence of the definition.

We finally turn to the last assertion, using Bayes’ law:

$$\begin{aligned} \Pr[B \mid A] &= 1 - \Pr[\bar{B} \mid A] \\ &= 1 - \Pr[A \mid \bar{B}] \cdot \Pr[\bar{B}] / \Pr[A] \geq 1 - (\varepsilon - \alpha) / \varepsilon = \alpha / \varepsilon. \end{aligned}$$

11.3.2 No-Message Attacks

The following *Forking Lemma* just states that the above oracle replay technique will often success with any good adversary.

Theorem 11.1 (The Forking Lemma) *Let $(\mathcal{K}, \mathcal{S}, \mathcal{V})$ be a digital signature scheme with security parameter k , with a signature as above, of the form $(m, \sigma_1, h, \sigma_2)$, where $h = \mathcal{H}(m, \sigma_1)$ and σ_2 depends on σ_1 and h only. Let \mathcal{A} be a probabilistic polynomial time Turing machine whose input only consists of public data and which can ask q_h queries to the random oracle, with $q_h > 0$. We assume that, within the time bound T , \mathcal{A} produces, with probability $\varepsilon \geq 7q_h/2^k$, a valid signature $(m, \sigma_1, h, \sigma_2)$. Then, within time $T' \leq 16q_h T / \varepsilon$, and with probability $\varepsilon' \geq 1/9$, a replay of this machine outputs two valid signatures $(m, \sigma_1, h, \sigma_2)$ and $(m, \sigma_1, h', \sigma'_2)$ such that $h \neq h'$.*

Proof : We are given an adversary \mathcal{A} , which is a probabilistic polynomial time Turing machine with random tape ω . During the attack, this machine asks a polynomial number of questions to the random oracle \mathcal{H} . We may assume that these questions are distinct: for instance, \mathcal{A} can store questions and answers in a table. Let $\mathcal{Q}_1, \dots, \mathcal{Q}_{q_h}$ be the q_h distinct questions and let $\rho = (\rho_1, \dots, \rho_{q_h})$ be the list of the q_h answers of \mathcal{H} . It is clear that a random choice of \mathcal{H} exactly corresponds to a random choice of ρ . Then, for a random choice of (ω, \mathcal{H}) , with probability ε , \mathcal{A} outputs a valid signature $(m, \sigma_1, h, \sigma_2)$. Since \mathcal{H} is a random oracle, it is easy to see that the probability for h to be equal to $\mathcal{H}(m, \sigma_1)$ is less than $1/2^k$, unless it has been asked during the attack. So, it is likely that the question (m, σ_1) is actually asked during a successful attack. Accordingly, we define $Ind_{\mathcal{H}}(\omega)$ to be the index of this question: $(m, \sigma_1) = \mathcal{Q}_{Ind_{\mathcal{H}}(\omega)}$ (we let $Ind_{\mathcal{H}}(\omega) = \infty$ if the question is never asked). We then define the sets

$$\begin{aligned} \mathbf{S} &= \{(\omega, \mathcal{H}) \mid \mathcal{A}^{\mathcal{H}}(\omega) \text{ succeeds} \ \& \ Ind_{\mathcal{H}}(\omega) \neq \infty\}, \\ \text{and } \mathbf{S}_i &= \{(\omega, \mathcal{H}) \mid \mathcal{A}^{\mathcal{H}}(\omega) \text{ succeeds} \ \& \ Ind_{\mathcal{H}}(\omega) = i\} \quad \text{for } i \in \{1, \dots, q_h\}. \end{aligned}$$

We thus call \mathbf{S} the set of the successful pairs (ω, \mathcal{H}) .

One should note that the set $\{\mathbf{S}_i \mid i \in \{1, \dots, q_h\}\}$ is a partition of \mathbf{S} . With those definitions, we find a lower bound for the probability of success, $\nu = \Pr[\mathbf{S}] \geq \varepsilon - 1/2^k$. Since we did the assumption that $\varepsilon \geq 7q_h/2^k \geq 7/2^k$, then $\nu \geq 6\varepsilon/7$. Let I be the set consisting of the most likely indices i ,

$$I = \{i \mid \Pr[\mathbf{S}_i \mid \mathbf{S}] \geq 1/2q_h\}.$$

The following lemma claims that, in case of success, the index lies in I with probability at least $1/2$.

Lemma 11.2

$$\Pr[Ind_{\mathcal{H}}(\omega) \in I \mid \mathbf{S}] \geq \frac{1}{2}.$$

Proof : By definition of the sets \mathbf{S}_i , $\Pr[Ind_{\mathcal{H}}(\omega) \in I \mid \mathbf{S}] = \sum_{i \in I} \Pr[\mathbf{S}_i \mid \mathbf{S}]$. This probability is equal to $1 - \sum_{i \notin I} \Pr[\mathbf{S}_i \mid \mathbf{S}]$. Since the complement of I contains fewer than q_h elements, this probability is at least $1 - q_h \times 1/2q_h \geq 1/2$.

We now run the attacker $2/\varepsilon$ times with random ω and random \mathcal{H} . Since $\nu = \Pr[\mathbf{S}] \geq 6\varepsilon/7$, with probability greater than $1 - (1 - 6\varepsilon/7)^{2/\varepsilon}$, we get at least one pair (ω, \mathcal{H}) in \mathbf{S} . It is easily seen that this probability is lower bounded by $1 - e^{-12/7} \geq 4/5$.

We now apply the Splitting-lemma (Lemma 11.1, with $\varepsilon = \nu/2q_h$ and $\alpha = \varepsilon/2$) for each integer $i \in I$: we denote by $\mathcal{H}_{|i}$ the restriction of \mathcal{H} to queries of index strictly less than i . Since $\Pr[\mathbf{S}_i] \geq \nu/2q_h$, there exists a subset Ω_i of executions such that,

$$\begin{aligned} \text{for any } (\omega, \mathcal{H}) \in \Omega_i, \Pr_{\mathcal{H}'}[(\omega, \mathcal{H}') \in \mathbf{S}_i \mid \mathcal{H}'_{|i} = \mathcal{H}_{|i}] &\geq \frac{\nu}{4q_h} \\ \Pr[\Omega_i \mid \mathbf{S}_i] &\geq \frac{1}{2}. \end{aligned}$$

Since all the subsets \mathbf{S}_i are disjoint,

$$\Pr_{\omega, \mathcal{H}}[(\exists i \in I) (\omega, \mathcal{H}) \in \Omega_i \cap \mathbf{S}_i \mid \mathbf{S}]$$

$$\begin{aligned}
&= \Pr \left[\bigcup_{i \in I} (\Omega_i \cap \mathbf{S}_i) \mid \mathbf{S} \right] = \sum_{i \in I} \Pr[\Omega_i \cap \mathbf{S}_i \mid \mathbf{S}] \\
&= \sum_{i \in I} \Pr[\Omega_i \mid \mathbf{S}_i] \cdot \Pr[\mathbf{S}_i \mid \mathbf{S}] \geq \left(\sum_{i \in I} \Pr[\mathbf{S}_i \mid \mathbf{S}] \right) / 2 \geq \frac{1}{4}.
\end{aligned}$$

We let β denote the index $\text{Ind}_{\mathcal{H}}(\omega)$ corresponding to the successful pair. With probability at least $1/4$, $\beta \in I$ and $(\omega, \mathcal{H}) \in \mathbf{S}_\beta \cap \Omega_\beta$. Consequently, with probability greater than $4/5 \times 1/5 = 1/5$, the $2/\varepsilon$ attacks have provided a successful pair (ω, \mathcal{H}) , with $\beta = \text{Ind}_{\mathcal{H}}(\omega) \in I$ and $(\omega, \mathcal{H}) \in \mathbf{S}_\beta$. Furthermore, if we replay the attack, with fixed ω but randomly chosen oracle \mathcal{H}' such that $\mathcal{H}'_\beta = \mathcal{H}_\beta$, we know that $\Pr_{\mathcal{H}'}[(\omega, \mathcal{H}') \in \mathbf{S}_\beta \mid \mathcal{H}'_\beta = \mathcal{H}_\beta] \geq \nu/4q_h$. Then

$$\begin{aligned}
&\Pr_{\mathcal{H}'}[(\omega, \mathcal{H}') \in \mathbf{S}_\beta \text{ and } \rho_\beta \neq \rho'_\beta \mid \mathcal{H}'_\beta = \mathcal{H}_\beta] \\
&\geq \Pr_{\mathcal{H}'}[(\omega, \mathcal{H}') \in \mathbf{S}_\beta \mid \mathcal{H}'_\beta = \mathcal{H}_\beta] - \Pr_{\mathcal{H}'}[\rho'_\beta = \rho_\beta] \geq \nu/4q_h - 1/2^k,
\end{aligned}$$

where $\rho_\beta = \mathcal{H}(\mathcal{Q}_\beta)$ and $\rho'_\beta = \mathcal{H}'(\mathcal{Q}_\beta)$. Using again the assumption that $\varepsilon \geq 7q_h/2^k$, the above probability is lower-bounded by $\varepsilon/14q_h$. We thus replay the attack $14q_h/\varepsilon$ times with a new random oracle \mathcal{H}' such that $\mathcal{H}'_\beta = \mathcal{H}_\beta$, and get another success with probability greater than

$$1 - (1 - \varepsilon/14q_h)^{14q_h/\varepsilon} \geq 1 - e^{-1} \geq 3/5.$$

Finally, after less than $2/\varepsilon + 14q_h/\varepsilon$ repetitions of the attack, with probability greater than $1/5 \times 3/5 \geq 1/9$, we have obtained two signatures $(m, \sigma_1, h, \sigma_2)$ and $(m', \sigma'_1, h', \sigma'_2)$, both valid w.r.t. their specific random oracle \mathcal{H} or \mathcal{H}' , and with the particular relations

$$\mathcal{Q}_\beta = (m, \sigma_1) = (m', \sigma'_1) \text{ and } h = \mathcal{H}(\mathcal{Q}_\beta) \neq \mathcal{H}'(\mathcal{Q}_\beta) = h'.$$

One may have noticed that the mechanics of our reduction depend on some parameters related to the attacker \mathcal{A} , namely, its probability of success ε and the number q_h of queries to the random oracle. This induces a lack of uniformity. A uniform version, in expected polynomial time is also possible.

Theorem 11.2 (The Forking Lemma – The Uniform Case) *Let $(\mathcal{K}, \mathcal{S}, \mathcal{V})$ be a digital signature scheme with security parameter k , with a signature as above, of the form $(m, \sigma_1, h, \sigma_2)$, where $h = \mathcal{H}(m, \sigma_1)$ and σ_2 depends on σ_1 and h only. Let \mathcal{A} be a probabilistic polynomial time Turing machine whose input only consists of public data and which can ask q_h queries to the random oracle, with $q_h > 0$. We assume that, within the time bound T , \mathcal{A} produces, with probability $\varepsilon \geq 7q_h/2^k$, a valid signature $(m, \sigma_1, h, \sigma_2)$. Then there is another machine which has control over \mathcal{A} and produces two valid signatures $(m, \sigma_1, h, \sigma_2)$ and $(m, \sigma_1, h', \sigma'_2)$ such that $h \neq h'$, in expected time $T' \leq 84480Tq_h/\varepsilon$.*

Proof : Now, we try to design a machine \mathbf{M} which succeeds in expected polynomial time:

1. \mathbf{M} initializes $j = 0$;
2. \mathbf{M} runs \mathcal{A} until it outputs a successful pair $(\omega, \mathcal{H}) \in \mathbf{S}$ and denotes by N_j the number of calls to \mathcal{A} to obtain this success, and by β the index $\text{Ind}_{\mathcal{H}}(\omega)$;

3. **M** replays, at most $140N_j\alpha^j$ times, \mathcal{A} with fixed ω and random \mathcal{H}' such that $\mathcal{H}'_{|\beta} = \mathcal{H}_{|\beta}$, where $\alpha = 8/7$;
4. **M** increments j and returns to 2, until it gets a successful forking.

For any execution of **M**, we denote by J the last value of j and by N the total number of calls to \mathcal{A} . We want to compute the expectation of N . Since $\nu = \Pr[\mathbf{S}]$, and $N_j \geq 1$, then $\Pr[N_j \geq 1/5\nu] \geq 3/4$. We define $\ell = \lceil \log_\alpha q_h \rceil$, so that, $140N_j\alpha^j \geq 28q_h/\varepsilon$ for any $j \geq \ell$, whenever $N_j \geq 1/5\nu$. Therefore, for any $j \geq \ell$, when we have a first success in **S**, with probability greater than $1/4$, the index $\beta = \text{Ind}_{\mathcal{H}}(\omega)$ is in the set I and $(\omega, \mathcal{H}) \in \mathbf{S}_\beta \cap \Omega_\beta$. Furthermore, with probability greater than $3/4$, $N_j \geq 1/5\nu$. Therefore, with the same conditions as before, that is $\varepsilon \geq 7q_h/2^k$, the probability of getting a successful fork after at most $28q_h/\varepsilon$ iterations at step 3 is greater than $6/7$.

For any $t \geq \ell$, the probability for J to be greater or equal to t is less than $(1 - 1/4 \times 3/4 \times 6/7)^{t-\ell}$, which is less than $\gamma^{t-\ell}$, with $\gamma = 6/7$. Furthermore,

$$E[N \mid J = t] \leq \sum_{j=0}^{j=t} (E[N_j] + 140E[N_j]\alpha^j) \leq \frac{141}{\nu} \times \sum_{j=0}^{j=t} \alpha^j \leq \frac{141}{\nu} \times \frac{\alpha^{t+1}}{\alpha - 1}.$$

So, the expectation of N is $E[N] = \sum_t E[N \mid J = t] \cdot \Pr[J = t]$ and then it can be shown to be less than $84480q_h/\varepsilon$. Hence the theorem.

11.3.3 Chosen-Message Attacks

However, this just covers the no-message attacks, without any oracle access. Since we can simulate any zero-knowledge protocol, even without having to restart the simulation because of the honest verifier (*i.e.* the challenge is randomly chosen by the random oracle \mathcal{H}) one can easily simulate the signer without the private key:

- one first chooses random $h, s \in \mathbb{Z}_q$;
- one computes $\mathbf{r} = s \cdot \mathbf{g} + h \cdot \mathbf{y}$ and defines $\mathcal{H}(m, \mathbf{r})$ to be equal to h , which is a uniformly distributed value;
- one can output (\mathbf{r}, h, s) as a valid signature of the message m .

This furthermore simulates the oracle \mathcal{H} , by defining $\mathcal{H}(m, \mathbf{r})$ to be equal to h . This simulation is almost perfect since \mathcal{H} is supposed to output a random value to any new query, and h is indeed a random value. Nevertheless, if the query $\mathcal{H}(m, \mathbf{r})$ has already been asked, $\mathcal{H}(m, \mathbf{r})$ is already defined, and thus the definition $\mathcal{H}(m, \mathbf{r}) \leftarrow h$ is impossible. But such a situation is very rare, which allows us to claim the following result, which stands for the Schnorr signature scheme but also for any signature derived from a three-round honest verifier zero-knowledge interactive proof of knowledge:

Theorem 11.3 *Let \mathcal{A} be a probabilistic polynomial time Turing machine whose input only consists of public data. We denote respectively by q_h and q_s the number of queries that \mathcal{A} can ask to the random oracle and the number of queries that \mathcal{A} can ask to the signer. Assume that, within a time bound T , \mathcal{A} produces, with probability $\varepsilon \geq 10(q_s + 1)(q_s + q_h)/2^k$, a valid signature $(m, \sigma_1, h, \sigma_2)$. If the triples (σ_1, h, σ_2) can be simulated without knowing the secret*

\mathcal{K}: Key Generation $\rightarrow (f, f^{-1})$
public key $f : X \rightarrow X$, a trapdoor one-way permutation onto X private key f^{-1} $\rightarrow (f, f^{-1})$
\mathcal{S}: Signature of $m \rightarrow \sigma$
$r = \mathcal{H}(m)$ and $\sigma = f^{-1}(r)$ $\rightarrow \sigma$ is the signature of m
\mathcal{V}: Verification of (m, σ)
check whether $f(\sigma) \stackrel{?}{=} \mathcal{H}(m)$ \rightarrow Yes/No

Figure 11.4: The FDH Signature.

key, with an indistinguishable distribution probability, then, a replay of the attacker \mathcal{A} , where interactions with the signer are simulated, outputs two valid signatures $(m, \sigma_1, h, \sigma_2)$ and $(m, \sigma_1, h', \sigma_2')$ such that $h \neq h'$, within time $T' \leq 23q_h T/\varepsilon$ and with probability $\varepsilon' \geq 1/9$.

A uniform version of this lemma can also be found in [?]. From a more practical point of view, these results state that if an adversary manages to perform an existential forgery under an adaptive chosen-message attack within an expected time T , after q_h queries to the random oracle and q_s queries to the signing oracle, then the discrete logarithm problem can be solved within an expected time less than $Cq_h T$, for some constant C . This result has thereafter been extended to the transformation of any identification scheme secure against passive adversaries into a signature scheme [?].

Brickell, Pointcheval, Vaudenay, and Yung also extended the *forking lemma* technique [?, ?] to many variants of El Gamal [?] and DSA [?], such as the Korean Standard KCDSA [?]. However, the original El Gamal and DSA schemes were not covered by this study, and are certainly not provably secure, even if no attack has ever been found against DSA.

11.4 RSA-Based Signatures

Unfortunately, with the above signatures based on the discrete logarithm, as any construction using the Fiat-Shamir paradigm, we do not really achieve our goal, because the reduction is costly, since q_h can be huge, as much as 2^{60} in practice. This security proof is meaningful for very large groups only.

In 1996, Bellare and Rogaway [?] proposed other candidates, based on the RSA assumption. The first scheme is the by-now classical hash-and-decrypt paradigm (*a.k.a.* the Full-Domain Hash paradigm): as for the basic RSA signature, the generation algorithm produces a large composite number $N = pq$, a public key e , and a private key d such that $e \cdot d = 1 \bmod \varphi(N)$. In order to sign a message m , one first hashes it using a full-domain hash function $\mathcal{H} : \{0, 1\}^* \rightarrow \mathbb{Z}_N^*$, and computes the e -th root, $\sigma = \mathcal{H}(m)^d \bmod N$. The verification algorithm simply checks whether the following equality holds, $\mathcal{H}(m) = \sigma^e \bmod N$.

More generally, the Full-Domain Hash signature can be defined as described on figure 11.4, for any trapdoor one-way permutation f .

For this scheme, Bellare and Rogaway proved, in the random-oracle model:

Theorem 11.4 *Let \mathcal{A} be an adversary which can produce, with success probability ε , an*

existential forgery under a chosen-message attack within a time t , after q_h and q_s queries to the hash function and the signing oracle respectively. Then the permutation f can be inverted with probability ε' within time t' where

$$\varepsilon' \geq \frac{\varepsilon}{q_s + q_h + 1} \quad \text{and} \quad t' \leq t + (q_s + q_h)T_f,$$

with T_f the time for an evaluation of f .

11.4.1 Basic Proof of the FDH Signature

In this proof, we incrementally define a sequence of games starting at the real game \mathbf{G}_0 and ending up at \mathbf{G}_5 . We make a very detailed sequence of games in this proof, since this is the first one. Some steps will be skipped in the other proofs. The goal of this proof is to reduce the inversion of the permutation f on an element y (find x such that $y = f(x)$) to an attack. We are thus given such a random challenge y .

Game \mathbf{G}_0 : This is the real attack game, in the random-oracle model, which includes the verification step. This means that the attack game consists in giving the public key to the adversary, and a full access to the signing oracle. When it outputs its forgery, one furthermore checks whether it is actually valid or not. Note that if the adversary asks q_s queries to the signing oracle and q_h queries to the hash oracle, at most $q_s + q_h + 1$ queries are asked to the hash oracle during this game, since each signing query may make such a new query, and the last verification step too. We are interested in the following event: S_0 which occurs if the verification step succeeds (and the signature is new).

$$\text{Succ}_{\text{fdh}}^{\text{euf}}(\mathcal{A}) = \Pr[S_0]. \quad (11.1)$$

Game \mathbf{G}_1 : In this game, we simulate the oracles, the hash oracle \mathcal{H} and the signing oracle \mathcal{S} , and the last verification step, as shown on Figure 11.5. From this simulation, we easily see that the game is perfectly indistinguishable from the real attack.

$$\Pr[S_1] = \Pr[S_0]. \quad (11.2)$$

Game \mathbf{G}_2 : Since the verification process is included in the attack game, the output message is necessarily asked to the hash oracle. Let us guess the index c of this (first) query. If the guess failed, we abort the game. Therefore, only a correct guess (event **GoodGuess**) may lead to a success.

$$\begin{aligned} \Pr[S_2] &= \Pr[S_1 \wedge \text{GoodGuess}] = \Pr[S_1 \mid \text{GoodGuess}] \times \Pr[\text{GoodGuess}] \\ &\geq \Pr[S_1] \times \frac{1}{q_h + q_s + 1}. \end{aligned} \quad (11.3)$$

Game \mathbf{G}_3 : We can now simulate the hash oracle, incorporating the challenge y , for which we want to extract the pre-image x by f :

► **Rule $\mathcal{H}^{(3)}$**

┌ If this is the c -th query, set $r \leftarrow y$; otherwise, choose a random
└ element $r \in X$. The record (q, \perp, r) is added to H-List.

\mathcal{H} oracle	<p>For a hash-query $\mathcal{H}(q)$, such that a record (q, \star, r) appears in H-List, the answer is r. Otherwise the answer r is defined according to the following rule:</p> <p style="margin-left: 20px;">►Rule $\mathcal{H}^{(1)}$</p> <p style="margin-left: 40px;">Choose a random element $r \in X$. The record (q, \perp, r) is added to H-List.</p> <p>Note: the second component of the elements of this list will be explained later.</p>
\mathcal{S} oracle	<p>For a sign-query $\mathcal{S}(m)$, one first asks for $r = \mathcal{H}(m)$ to the \mathcal{H}-oracle, and then the signature σ is defined according to the following rule:</p> <p style="margin-left: 20px;">►Rule $\mathcal{S}^{(1)}$</p> <p style="margin-left: 40px;">Computes $\sigma = f^{-1}(r)$.</p>
\mathcal{V} oracle	<p>The game ends with the verification of the output (m, σ) from the adversary. One first asks for $r = \mathcal{H}(m)$, and checks whether $r = f(\sigma)$.</p>

Figure 11.5: Simulation of the Attack Game against FDH

Because of the random choice for the challenge y , this rule lets the game indistinguishable from the previous one.

$$\Pr[\mathbf{S}_3] = \Pr[\mathbf{S}_2]. \quad (11.4)$$

Game \mathbf{G}_4 : We now modify the simulation of the hash oracle for other queries, which may be used in signing queries:

►Rule $\mathcal{H}^{(4)}$

If this is the c -th query, set $r \leftarrow y$ and $s \leftarrow \perp$; otherwise, choose a random element $s \in X$, and compute $r = f(s)$. The record (q, s, r) is added to H-List.

Because of the permutation property of f , and the random choice for s , this rule lets the game indistinguishable from the previous one.

$$\Pr[\mathbf{S}_4] = \Pr[\mathbf{S}_3]. \quad (11.5)$$

Game \mathbf{G}_5 : By now, excepted for the c -th hash query, which will be involved in the forgery (and thus not asked to the signing oracle), the pre-image is known. One can thus simulate the signing oracle without quering f^{-1} :

►Rule $\mathcal{S}^{(5)}$

Lookup for (m, s, r) in H-List, and set $\sigma = s$.

Since the message corresponding to the c -th query cannot be asked to the signing oracle, otherwise it would not be a valid forgery, this rule lets the game indistinguishable from the previous one.

$$\Pr[\mathbf{S}_5] = \Pr[\mathbf{S}_4]. \quad (11.6)$$

Note that now, the simulation can easily be performed, without any specific computational power or oracle access. Just a few more evaluations of f are done to simulate the hash oracle, and the forgery leads to the pre-image of y :

$$\Pr[S_5] = \text{Succ}_f^{\text{ow}}(t + (q_h + q_s)T_f). \quad (11.7)$$

As a consequence, using equations (11.1), (11.2), (11.3), (11.4), (11.5), (11.6) and (11.7)

$$\begin{aligned} \text{Succ}_f^{\text{ow}}(t + (q_h + q_s)T_f) &= \Pr[S_5] = \Pr[S_3] = \Pr[S_4] = \Pr[S_2] \\ &\geq \frac{1}{q_h + q_s + 1} \times \Pr[S_1] \geq \frac{1}{q_h + q_s + 1} \times \Pr[S_0]. \end{aligned}$$

And thus,

$$\text{Succ}_{\text{fdh}}^{\text{euf}}(\mathcal{A}) \leq (q_h + q_s + 1) \times \text{Succ}_f^{\text{ow}}(t + (q_h + q_s)T_f).$$

□

11.4.2 Improved Security Result

This reduction has been thereafter improved [?], thanks to the random self-reducibility of the RSA function. The following result applies as soon as the one-way permutation has some homomorphic property on the group X :

$$f(x \otimes y) = f(x) \otimes f(y).$$

Theorem 11.5 *Let \mathcal{A} be an adversary which can produce, with success probability ε , an existential forgery under a chosen-message attack within a time t , after q_h and q_s queries to the hash function and the signing oracle respectively. Then the permutation f can be inverted with probability ε' within time t' where*

$$\varepsilon' \geq \frac{\varepsilon}{q_s} \times \exp(-2) \quad \text{and} \quad t' \leq t + (q_s + q_h)T_f,$$

with T_f the time for an evaluation of f .

This proof can be performed as the previous one, and thus starts at the real game \mathbf{G}_0 , then we can use the same simulation as in the game \mathbf{G}_1 . The sole formal difference in the simulation will be the **H-List** which elements have one more field, and are thus initially of the form (q, \perp, \perp, r) . Things differ much after that, using a real value p between 0 and 1, which will be made precise later. The idea here, is to make any forgery useful for inverting the permutation f , not only a specific (guessed) one. On the other hand, one must still be able to simulate the signing oracle. The probability p will separate the two situations:

Game \mathbf{G}_2 : A random coin decides whether we introduce the challenge y in the hash answer, or an element with a known pre-image:

► **Rule $\mathcal{H}^{(2)}$**

One chooses a random $s \in X$. With probability p , one sets $r \leftarrow y \otimes f(s)$ and $t \leftarrow 1$; otherwise, $r \leftarrow f(s)$ and $t \leftarrow 0$. The record (q, t, s, r) is added to **H-List**.

Because of the homomorphic property on the group X of the permutation f , this rule lets the game indistinguishable from the previous one. Note again that elements in **H-List** contain one more field t than in the previous proof. One may see that $r = y^t \otimes f(s)$.

Game G_3 : For a proportion $1 - p$ of the signature queries, one can simulate the signing oracle without having to invert the permutation f :

► **Rule $S^{(3)}$**

 | Lookup for (m, t, s, r) in **H-List**, if $t = 1$ then halt the game, otherwise set $\sigma = s$.

This rule lets the game indistinguishable, unless one signing query fails ($t = 1$), which happens with probability p , for each signature:

$$\Pr[S_3] = (1 - p)^{q_s} \times \Pr[S_2]. \quad (11.8)$$

Note that now, the simulation can easily be performed, without any specific computational power or oracle access. Just a few more exponentiations are done to simulate the hash oracle, and the forgery (m, σ) leads to the pre-image of y , if ($t = 1$). The latter case holds with probability p . Indeed, (m, t, s, r) can be found in the **H-List**, and then $r = y^t \otimes f(s) = y \otimes f(s) = f(\sigma)$, which easily leads to the pre-image of y by f :

$$\text{Succ}_f^{\text{ow}}(t + (q_h + q_s)T_f) = p \times \Pr[S_3]. \quad (11.9)$$

Using equations (11.1), (11.2), (11.8) and (11.9)

$$\begin{aligned} \text{Succ}_f^{\text{ow}}(t + (q_h + q_s)T_f) &= p \times \Pr[S_3] = p \times (1 - p)^{q_s} \times \Pr[S_2] \\ &= p \times (1 - p)^{q_s} \times \Pr[S_1] = p \times (1 - p)^{q_s} \times \Pr[S_0]. \end{aligned}$$

And thus,

$$\text{Succ}_{\text{fdh}}^{\text{euf}}(\mathcal{A}) \leq \frac{1}{p(1 - p)^{q_s}} \times \text{Succ}_f^{\text{ow}}(t + (q_h + q_s)T_f).$$

Therefore, the success probability of our inversion algorithm is $p(1 - p)^{q_s}\varepsilon$, if ε is the success probability of the adversary. If $q_s > 0$, the latter expression is optimal for $p = 1/(q_s + 1)$. And for this parameter, and a huge value q_s , the success probability is approximately ε/eq_s . It is anyway larger than ε/e^2q_s (where $e = \exp(1) \approx 2.17\dots$).

As far as time complexity is concerned, each random oracle simulation (which can be launched by a signing simulation) requires a modular exponentiation to the power e , hence the result. \square

This is a great improvement since the success probability does not depend anymore on q_h . Furthermore, q_s can be limited by the user, whereas q_h cannot. In practice, one only assumes $q_h \leq 2^{60}$, but q_s can be limited below 2^{30} .

11.4.3 PSS: The Probabilistic Signature Scheme

However, one would like to get more, suppressing any coefficient. In their paper [?], Bellare and Rogaway proposed such a better candidate, the Probabilistic Signature Scheme (PSS, see Figure 11.6): the key generation is still the same, but the signature process involves three hash functions

$$\begin{aligned} \mathcal{F} : \{0, 1\}^{k_2} &\rightarrow \{0, 1\}^{k_0}, & \mathcal{G} : \{0, 1\}^{k_2} &\rightarrow \{0, 1\}^{k_1}, \\ \mathcal{H} : \{0, 1\}^* &\rightarrow \{0, 1\}^{k_2}, \end{aligned}$$




Figure 11.6: Probabilistic Signature Scheme

where $k = k_0 + k_1 + k_2 + 1$ satisfies $\{0, 1\}^{k-1} \subset X \subset \{0, 1\}^k$. We remind that f is a trapdoor one-way permutation onto X , with an homomorphic relationship. For each message m to be signed, one chooses a random string $r \in \{0, 1\}^{k_1}$. One first computes $w = \mathcal{H}(m, r)$, $s = \mathcal{G}(w) \oplus r$ and $t = \mathcal{F}(w)$. Then one concatenates $y = 0\|w\|s\|t$, where $a\|b$ denotes the concatenation of the bit strings a and b . Finally, one computes the pre-image by f , $\sigma = f^{-1}(y)$. The verification algorithm first computes $y = f(\sigma)$, and parses it as $y = b\|w\|s\|t$. Then, one can get $r = s \oplus \mathcal{G}(w)$, and checks whether $b = 0$, $w = \mathcal{H}(m, r)$ and $t = \mathcal{F}(w)$.

About this PSS construction, Bellare and Rogaway proved the security in the random-oracle model.

Theorem 11.6 *Let \mathcal{A} be a CMA-adversary against f -PSS which produces an existential forgery within a time t , after q_f , q_g , q_h and q_s queries to the hash functions \mathcal{F} , \mathcal{G} and \mathcal{H} and the signing oracle respectively. Then its success probability is upper-bounded by*

$$\text{Succ}_f^{\text{ow}}(t + (q_s + q_h)k_2 \cdot T_f) + \frac{1}{2^{k_2}} + (q_s + q_h) \cdot \left(\frac{q_s}{2^{k_1}} + \frac{q_f + q_g + q_h + q_s + 1}{2^{k_2}} \right),$$

with T_f the time for an evaluation of f .

The important point in this security result is the very tight link between success probabilities, but also the almost linear time of the reduction. Thanks to this exact and efficient security result, RSA-PSS has become the new PKCS #1 v2.1 standard for signature [?]. Another variant has been proposed with message-recovery: PSS-R which allows one to include a large part of the message inside the signature. This makes a signed-message shorter than the size of the signature plus the size of the message, since the latter is inside the former one.

Chapter 12

Automating Game-Based Proofs

Part IV

Links between the two Settings

Chapter 13

Soundness of Static Equivalence

In this chapter we will show that under some assumptions on the cryptographic primitives static equivalence is sound with respect to a computational model, i.e. whenever two frames φ_1 and φ_2 are statically equivalent, the distributions corresponding to the implementations of these frames are computationally indistinguishable.

13.1 Abstract and computational algebras

13.1.1 Abstract algebras.

As previously we will use an abstract algebra to represent messages of a protocol symbolically by terms. However, we will suppose that the algebra is sorted. In addition to the sets \mathcal{F} , \mathcal{N} and \mathcal{X} for function symbols, names and variables resp., we consider a set \mathcal{S} of sorts. The arity of a function symbol $f \in \mathcal{F}$ will be written accordingly as $\text{ar}(f) = s_1 \times \dots \times s_n \rightarrow s$ where $s_i, s \in \mathcal{S}$ are the sorts of the arguments and the result of f . Similarly, We suppose that terms and substitutions are always well-sorted.

In the following, to avoid notational clutter, we suppose that a frame $\varphi = \text{new } \bar{a}.\sigma$ binds all names in σ and we generally omit to precise \bar{a} . Note that this is not a loss of generality as for each name n supposed to be free we can simply extend the frame by $\{x_n \mapsto n\}$ (cf Property (6.3) shown in Exercise 6.4).

13.1.2 Abstract and concrete models.

We now give terms and frames a concrete semantics, parameterized by an implementation of the primitives. Provided a set of sorts \mathcal{S} and a set of symbols \mathcal{F} as above, a $(\mathcal{S}, \mathcal{F})$ -computational algebra A consists of

- a non-empty set of bit-strings $\llbracket s \rrbracket_A \subseteq \{0, 1\}^*$ for each sort $s \in \mathcal{S}$; moreover, if s_2 is a subsort of s_1 we require that $\llbracket s_2 \rrbracket_A \subseteq \llbracket s_1 \rrbracket_A$;
- a computable function $\llbracket f \rrbracket_A : \llbracket s_1 \rrbracket_A \times \dots \times \llbracket s_k \rrbracket_A \rightarrow \llbracket s \rrbracket_A$ for each $f \in \mathcal{F}$ with $\text{ar}(f) = s_1 \times \dots \times s_k \rightarrow s$;
- an effective procedure to draw random elements from $\llbracket s \rrbracket_A$; we denote such a drawing by $x \xleftarrow{R} \llbracket s \rrbracket_A$;

Assume a fixed $(\mathcal{S}, \mathcal{F})$ -computational algebra A . We associate to each frame $\varphi = \{x_1 \mapsto t_1, \dots, x_n \mapsto t_n\}$ a distribution $\psi = \llbracket \varphi \rrbracket_A$, of which the drawings $\widehat{\psi} \stackrel{R}{\leftarrow} \psi$ are computed as follows:

1. for each name a of sort s appearing in t_1, \dots, t_n , draw a value $\widehat{a} \stackrel{R}{\leftarrow} \llbracket s \rrbracket_A$;
2. for each x_i ($1 \leq i \leq n$) of sort s_i , compute $\widehat{t}_i \in \llbracket s_i \rrbracket_A$ recursively on the structure of terms: $\widehat{f(t'_1, \dots, t'_m)} = \llbracket f \rrbracket_A(\widehat{t'_1}, \dots, \widehat{t'_m})$;
3. return the value $\widehat{\psi} = \{x_1 \mapsto \widehat{t}_1, \dots, x_n \mapsto \widehat{t}_n\}$.

Such values $\phi = \{x_1 = e_1, \dots, x_n = e_n\}$ with $e_i \in \llbracket s_i \rrbracket_A$ are called *concrete frames*. In the following we will focus on *asymptotic* notions of security and consider families of computational algebra (A_η) indexed by a complexity parameter $\eta \geq 0$. (This parameter η might be thought of as the size of keys and other secret values.) The *concrete semantics* of a frame φ is a family of distributions over concrete frames $(\llbracket \varphi \rrbracket_{A_\eta})$. We only consider families of computational algebras (A_η) such that each required operation on algebras is feasible by a (uniform, probabilistic) polynomial-time algorithm in the complexity parameter η . This ensures that the concrete semantics of terms and frames is efficiently computable (in the same sense). We also suppose that computational algebras implement the equalities dictated by the equational theory, i.e. for all terms t_1, t_2 such that $t_1 =_{\mathcal{E}} t_2$ we suppose that $\mathbb{P}[e_1, e_2 \stackrel{R}{\leftarrow} \llbracket t_1, t_2 \rrbracket_{A_\eta} : e_1 = e_2] = 1$.

Families of distributions over concrete frames benefit from the usual notion of cryptographic indistinguishability. Intuitively, two families of distributions (ψ_η) and (ψ'_η) are *indistinguishable*, written $(\psi_\eta) \approx (\psi'_\eta)$, iff no probabilistic polynomial-time adversary \mathcal{A} can guess whether he is given a sample from ψ_η or ψ'_η with a probability significantly greater than $\frac{1}{2}$.

Definition 13.1 *Let (ψ_η) and (ψ'_η) be two families of distributions. (ψ_η) and (ψ'_η) are indistinguishable, written $(\psi_\eta) \approx (\psi'_\eta)$, iff for all probabilistic polynomial-time adversary \mathcal{A} the advantage of \mathcal{A} ,*

$$\text{Adv}_{\mathcal{A}}^{\text{IND}}(\psi_\eta, \psi'_\eta) = \left| \mathbb{P}[\widehat{\psi} \stackrel{R}{\leftarrow} \psi_\eta : \mathcal{A}(\eta, \widehat{\psi}) = 1] - \mathbb{P}[\widehat{\psi} \stackrel{R}{\leftarrow} \psi'_\eta : \mathcal{A}(\eta, \widehat{\psi}) = 1] \right|$$

is a negligible function of η . A function is negligible in η if for any integer $n > 0$ there exists η_0 such that $f(\eta) < \eta^{-n}$ for all $\eta > \eta_0$.

We can now define what it means for an abstract algebra to be sound w.r.t. static equivalence.

Definition 13.2 *A family of computational algebras (A_η) is $\sim_{\mathcal{E}}$ -sound iff*

$$\varphi_1 \sim_{\mathcal{E}} \varphi_2 \implies (\llbracket \varphi_1 \rrbracket_{A_\eta}) \approx (\llbracket \varphi_2 \rrbracket_{A_\eta}).$$

13.2 Soundness for symmetric encryption

We will now show a soundness result for a theory of symmetric encryption and pairs. We will first introduce a bit more precisely these particular symbolic and computational algebras.

Symbolic model. Our symbolic model consists of the set of sorts $\mathcal{S} = \{\text{Data}\}$, an infinite number of names for sort Data called keys and the function symbols:

senc, sdec	$: \text{Data} \times \text{Data} \rightarrow \text{Data}$	encrypt, decrypt
pair	$: \text{Data} \times \text{Data} \rightarrow \text{Data}$	pair constructor
π_1, π_2	$: \text{Data} \rightarrow \text{Data}$	projections
samekey	$: \text{Data} \times \text{Data} \rightarrow \text{Data}$	key equalities test
$\text{tenc}, \text{tpair}$	$: \text{Data} \rightarrow \text{Data}$	type testers
$0, 1$	$: \text{Data}$	constants

We consider the equational theory **sym** generated by:

$$\begin{array}{ll}
 \text{sdec}(\text{senc}(x, y), y) &= x & \pi_1(\text{pair}(x, y)) &= x \\
 \pi_2(\text{pair}(x, y)) &= y & \text{samekey}(\text{senc}(x, y), \text{senc}(z, y)) &= 1 \\
 \text{tenc}(\text{senc}(x, y)) &= 1 & \text{tpair}(\text{pair}(x, y)) &= 1
 \end{array}$$

As previously $\text{senc}(t, k)$ is also written $\{t\}_k$ and $\text{pair}(t, t')$ is also written $\langle t, t' \rangle$.

Well-formed frames. A soundness result generally does not hold for arbitrary frames and we need to make some restrictions. In particular we need to forbid *key cycles*. (We will discuss in more detail the necessity of this condition in Exercise 13.2). We say that a term k is used at a key position in a term t if there exists a subterm $\text{senc}(t', k)$ of t . Else k is used at a plaintext position. We write $K(\varphi)$ the set of all terms k appearing in key positions in φ . Moreover we denote by $\text{NDK}(\varphi) = \{k \mid k \in K(\varphi), \varphi \not\models k\}$ the set of non-deducible keys in $K(\varphi)$.

Definition 13.3 We define the relation encrypts_φ for a frame φ on the terms in $\text{NDK}(\varphi)$ as follows. $k_1 \text{ encrypts}_\varphi k_2$ if there exists $\{t\}_{k_1} \in \text{st}(\varphi)$ such that $k_2 \in \text{st}(t)$. We say that a frame φ is *acyclic* iff the relation encrypts_φ is acyclic. Otherwise we say that the frame contains a *key-cycle*.

Example 13.1 The following frames contain a key-cycle.

$$\begin{array}{l}
 \{x \mapsto \{k\}_k\} \\
 \{x \mapsto \{k_1\}_{k_2}, y \mapsto \{k_2\}_{k_1}\} \\
 \{x \mapsto \{\{k_3\}_{k_1}\}_{k_2}, y \mapsto \{\{k_2\}_{k_3}\}_{k_1}\}
 \end{array}$$

The following frames are acyclic.

$$\begin{array}{l}
 \{x \mapsto \{k_1\}_{k_2}, y \mapsto \{k_2\}_{k_1}, z \mapsto k_1\} \\
 \{x \mapsto \{k_1\}_{k_2}, y \mapsto \{k_3\}_{k_1}\} \\
 \{x \mapsto \{\{k_3\}_{k_1}\}_{k_2}, y \mapsto \{k_3\}_{k_2}\}
 \end{array}$$

Definition 13.4 A frame $\varphi = \{x_1 \mapsto t_1, \dots, x_n \mapsto t_n\}$ is well-formed if

- $K(\varphi)$ only contains names;
- φ is acyclic;
- the terms t_i do not contain destructors, i.e. they only use symbols senc , pair , 0 and 1 .

Concrete model. We recall the standard definition for symmetric encryption schemes. A symmetric encryption scheme \mathcal{SE} is defined by three algorithms \mathcal{KG} , \mathcal{E} and \mathcal{D} . The key generation algorithm takes as input the security parameter η and outputs a key k . The encryption algorithm \mathcal{E} is randomized, it takes as input a bit-string s and a key k and returns the encryption of s using k . The decryption algorithm \mathcal{D} takes as input a bit-string c representing a ciphertext and a key k and outputs the corresponding plaintext. Given $k \leftarrow \mathcal{KG}(\eta)$, we have that for any bit-string s , if $c \leftarrow \mathcal{E}(k, s)$ then it is required that $\mathcal{D}(c) = s$.

The family of computational algebras (A_η) giving the concrete semantics depends on a symmetric encryption scheme $\mathcal{SE} = (\mathcal{KG}, \mathcal{E}, \mathcal{D})$. The concrete domain $\llbracket \text{Data} \rrbracket_{A_\eta}$ contains all the possible bit-strings and is equipped with the distribution induced by \mathcal{KG} . Interpretation for constants 0 and 1 are respectively bit-strings 0^η and 1^η . The `senc` and `sdec` function are respectively interpreted using algorithm \mathcal{E} and \mathcal{D} . We assume the existence in the concrete model of a concatenation operation which is used to interpret the `pair` symbol. The corresponding left and right projections implement π_1 and π_2 . Finally, as we are only interested in well-formed frames, we do not provide any computational interpretation for `tenc`, `tpair` and `samekey`.

Semantic security. We use schemes that satisfy a variant of semantic security, namely length-concealing semantic security. The definition that we recall below uses a left-right encryption oracle $LR_{\mathcal{SE}}^b$. This oracle first generates a key k using \mathcal{KG} . Then it answers queries of the form (bs_0, bs_1) , where bs_0 and bs_1 are bit-strings. The oracle returns ciphertext $\mathcal{E}(bs_b, k)$. The goal of the adversary \mathcal{A} is to guess the value of bit b . His advantage is defined as:

$$\text{Adv}_{\mathcal{SE}, \mathcal{A}}^{\text{cpa}}(\eta) = \left| \mathbb{P}[\mathcal{A}^{LR_{\mathcal{SE}}^1}(\eta) = 1] - \mathbb{P}[\mathcal{A}^{LR_{\mathcal{SE}}^0}(\eta) = 1] \right|$$

Encryption scheme \mathcal{SE} is IND-CPA secure if the advantage of any adversary \mathcal{A} is negligible in η . The difference with the standard notion of semantic security is that we require that the scheme hides the length of the plaintext (and therefore we do not restrict bs_0 and bs_1 to have equal length). By abuse of notation we call the resulting scheme also IND-CPA secure.

Soundness result. We can now state the soundness result for symmetric encryption.

Proposition 13.1 *Let φ_1, φ_2 be two well-formed frames and A_η a computational algebra based on an IND-CPA secure encryption scheme. Then $\varphi_1 \sim_{\text{sym}} \varphi_2$ implies that $\llbracket \varphi_1 \rrbracket_{A_\eta} \approx \llbracket \varphi_2 \rrbracket_{A_\eta}$.*

13.3 Proof of Proposition 13.1

The proof will be done in two steps:

1. static equivalence implies pattern equivalence;
2. soundness of pattern equivalence.

13.3.1 Pattern equivalence.

We introduce another notion of equivalence on frames which will be more convenient to manipulate to show the soundness result.

Definition 13.5 Let $\varphi = \{x_1 \mapsto t_1, \dots, x_n \mapsto t_n\}$ be a frame containing only constructors and names in key positions. We define $\text{pat}(\varphi)$ to be the sequence $(\text{pat}(t_1, K), \dots, \text{pat}(t_n, K))$ where $K = \{k \mid k \in \text{names}(\varphi) \text{ and } \varphi \vdash_{\text{sym}} k\}$ and

$$\begin{aligned} \text{pat}(\langle t_0, t_1 \rangle, K) &= \langle \text{pat}(t_0, K), \text{pat}(t_1, K) \rangle \\ \text{pat}(\{t\}_k, K) &= \{\text{pat}(t, K)\}_k && \text{if } k \in K \\ \text{pat}(\{t\}_k, K) &= \{\square\}_k && \text{if } k \notin K \\ \text{pat}(t, K) &= t && \text{else} \end{aligned}$$

where \square is a special constant.

Let φ_1 and φ_2 be two frames containing only constructors and names in key positions. We say that φ_1 and φ_2 are pattern-equivalent, denoted $\varphi_1 \equiv_{\text{pat}} \varphi_2$ if there exists φ_2^α such that $\varphi_2^\alpha =_\alpha \varphi_2$ and $\text{pat}(\varphi_1) = \text{pat}(\varphi_2^\alpha)$.

Exercise 13.1 Let $\varphi_i = \{x \mapsto t_i\}$. For each of the following values for t_i determine whether $\varphi_1 \equiv_{\text{pat}} \varphi_2$.

t_1	t_2
0	0
0	1
$\{0\}_k$	$\{1\}_k$
$\langle k, \{0\}_k \rangle$	$\langle k, \{1\}_k \rangle$
$\langle k, \{\langle \{0\}_{k'}, 0 \rangle\}_k \rangle$	$\langle k, \{\langle \{1\}_{k'}, 0 \rangle\}_k \rangle$
$\{0\}_k$	$\{k\}_k$
$\{\langle 0, 0 \rangle\}_k$	$\{0\}_k$
$\langle \{0\}_k, \{0\}_k \rangle$	$\langle \{0\}_k, \{1\}_k \rangle$
$\langle \{0\}_k, \{0\}_k \rangle$	$\langle \{0\}_k, \{1\}_{k'} \rangle$

We now show that static equivalence implies pattern equivalence.

Lemma 13.1

$$\varphi \sim_{\text{sym}} \varphi' \implies \varphi \equiv_{\text{pat}} \varphi'$$

Proof: We suppose that φ and φ' contain only one term, i.e., $\varphi = \{x \mapsto t\}$ and $\varphi' = \{x \mapsto t'\}$. This hypothesis is not a loss of generality in the theory **sym** as sequences of terms can be grouped into a single term using the pair operator.

Let K and K' be the sets of deducible keys from respectively φ and φ' , let u be $\text{pat}(t, K)$ and u' be $\text{pat}(t', K')$. We suppose that $u \neq u'\rho$ for any bijective renaming ρ and show that $\varphi \not\sim_{\text{sym}} \varphi'$.

Let p be the minimal position, such that

- $p \in \text{pos}(u) \cap \text{pos}(u')$;
- $\text{root}(u|_p) \neq \text{root}(u'|_p)$;
- either $\text{root}(u|_p)$ or $\text{root}(u'|_p)$ is not a name.

We have to consider two cases according to whether such a p exists.

If such a position p exists and $\text{root}(u|_p) \neq \square$, we denote by M_p the term such that $M_p \varphi =_{\text{sym}} u|_p$. Note that as p is minimal, p cannot be a key position. It directly follows from the definition of patterns that such a term M_p exists. We have one of the following cases:

- $root(u|_p) = \text{pair}$
Then we have that $\text{tpair}(M_p\varphi) =_{\text{sym}} 1$ while $\text{tpair}(M_p\varphi') \neq_{\text{sym}} 1$.
- $root(u|_p) = \text{senc}$
Then we have that $\text{tenc}(M_p\varphi) =_{\text{sym}} 1$ while $\text{tenc}(M_p\varphi') \neq_{\text{sym}} 1$.
- $root(u|_p) = \square$
Then we have that $root(u|_{p'}) = \text{senc}$ where $p = p' \cdot 1$. Moreover, $root(u|_{p'}) = root(u'|_{p'})$, as p is minimal, and $root(u'|_p) \neq \square$. Hence, by definition of patterns the key at position $p' \cdot 2$ and the term at position p' are deducible from φ' . We denote by M_k and $M_{p'}$ the terms such that $M_k\varphi' =_{\text{sym}} u'|_{p' \cdot 2}$ and $M_{p'}\varphi' =_{\text{sym}} u'|_{p'}$. Then, we have that $\text{samekey}(M_{p'}, \text{senc}(0, M_k))\varphi' =_{\text{sym}} 1$, while this test does not hold under φ .
- $root(u|_p) = 0$
Then we have that $M_p\varphi =_{\text{sym}} 0$ while $M_p\varphi' \neq_{\text{sym}} 0$.
- $root(u|_p) = 1$
Then we have that $M_p\varphi =_{\text{sym}} 1$ while $M_p\varphi' \neq_{\text{sym}} 1$.
- $root(u|_p)$ is a name
By definition of p , $root(u'|_p)$ is not a name. Hence this case is symmetric to one of the above cases.

If no such position p exists, then the patterns only differ by names. As, there exists no bijective renaming there must be a name which is used twice in one of the frames, but two different names are used at the same positions in the other frame. Let p_1 and p_2 be the positions of these two names. We consider the following cases.

- p_1 and p_2 are both plain text positions
From the definition of patterns it follows that they are both deducible using terms which we denote by M_{p_1} and M_{p_2} respectively. The test $M_{p_1} = M_{p_2}$ distinguishes φ and φ' .
- p_1 and p_2 are both key positions
Let $p_1 = p'_1 \cdot 2$ and $p_2 = p'_2 \cdot 2$. Then, the test $\text{samekey}(M_{p'_1}, M_{p'_2}) =_{\text{sym}} 1$ distinguishes φ and φ' , where $M_{p'_1}$, respectively $M_{p'_2}$, denotes the term such that $M_{p'_1}\varphi =_{\text{sym}} t|_{p'_1}$, respectively $M_{p'_2}\varphi =_{\text{sym}} t|_{p'_2}$.
- p_1 is a plain text position and p_2 is a key position
We use the same notations as in the two previous cases. The test $\text{samekey}(M_{p'_2}, \text{senc}(0, M_{p_1})) =_{\text{sym}} 1$ distinguishes φ and φ' .
- p_1 is a key position and p_2 is a plain text position
This case is symmetric to the previous one.

13.3.2 Soundness of \equiv_{pat} .

We now show a soundness result for pattern equivalence.

Lemma 13.2 *Let φ_0, φ_1 be two well-formed frames such that $\varphi_0 \equiv_{\text{pat}} \varphi_1$. Let A_η be a computational algebra based on an IND-CPA secure encryption scheme. We have that $\llbracket \varphi_0 \rrbracket_{A_\eta} \approx \llbracket \varphi_1 \rrbracket_{A_\eta}$.*

Proof : The proof is based on a hybrid argument. Let $\mathcal{SE} = (\mathcal{KG}, \mathcal{E}, \mathcal{D})$ be the IND-CPA secure encryption scheme. We suppose that the set of deducible names in φ_0 and φ_1 are the same (in the sense that the same recipe yields the same name for both frames) and $\text{NDK}(\varphi_0) \cap \text{NDK}(\varphi_1) = \emptyset$. Moreover, we denote by k_1, \dots, k_n the set of names in $\text{NDK}(\varphi_0) \cup \text{NDK}(\varphi_1)$ such that $k_i (\text{encrypts}_{\varphi_0} \cup \text{encrypts}_{\varphi_1}) k_j$ implies $i > j$. These hypotheses can be made without loss of generality as they can be achieved by α -conversion.

Let \mathcal{A} be an adversary against the soundness of \equiv_p . We build two adversaries \mathcal{B}_b against IND-CPA for $b \in \{0, 1\}$. Adversary \mathcal{B}_b randomly samples an integer i between 1 and n . k_i will be the challenge key. \mathcal{B}_b also generates a value for each key using the key generation algorithm \mathcal{KG} . \mathcal{B}_b then executes \mathcal{A} and uses an algorithm S to compute a sequence of bitstrings given as input to \mathcal{A} . (For readability, in the following we will omit the complexity parameter η which is implicitly the first argument of each adversary.)

Adversary \mathcal{B}_b^{LRSE} :

```

i  $\leftarrow [1, n]$ 
for  $k \in \text{names}(\varphi_1, \varphi_2)$  do  $\llbracket k \rrbracket \leftarrow \mathcal{KG}(\eta)$ 
d  $\leftarrow \mathcal{A}(S)$ 
return d

```

S uses the values $\llbracket k \rrbracket$ generated by \mathcal{B}_b and the sequences (t_j^1, \dots, t_j^ℓ) where $\varphi_j = \{x_1 \mapsto t_j^1, \dots, x_\ell \mapsto t_j^\ell\}$ ($j \in \{0, 1\}$). S computes a bitstring bs^k for each pair (t_0^k, t_1^k) ($1 \leq k \leq \ell$) as follows and eventually returns the sequence (bs^1, \dots, bs^ℓ) . In the following we denote by K_j the set of deducible names in φ_j . On a pair (t_0, t_1) , S works as follows:

- If t_0 is a pair $\langle t'_0, t''_0 \rangle$ then because of \equiv_{pat} , t_1 is a pair $\langle t'_1, t''_1 \rangle$ and $\text{pat}(\langle t'_0, t''_0 \rangle, K_0) = \text{pat}(\langle t'_1, t''_1 \rangle, K_1)$. S is recursively applied on (t'_0, t'_1) yielding bs' and on (t''_0, t''_1) yielding bs'' . S returns $\llbracket \text{pair} \rrbracket_{A_\eta}(bs', bs'')$.
- If t_0 is an encryption $\{t'_0\}_k$ and $\varphi_1 \vdash k$ then t_1 is an encryption $\{t'_1\}_k$ and $\text{pat}(t'_0, K_0) \equiv_{\text{pat}} \text{pat}(t'_1, K_1)$. Algorithm S is applied recursively on (t'_0, t'_1) yielding bs' . S returns $\mathcal{E}(bs', \llbracket k \rrbracket)$.
- If t_0 is an encryption $\{t'_0\}_{k_{j_0}}$ and $\varphi_0 \not\vdash k_{j_0}$, then t_1 is an encryption $\{t'_1\}_{k_{j_1}}$ and $\varphi_1 \not\vdash k_{j_1}$. S is applied recursively on (t'_0, t'_1) yielding a bit-string bs_b .

$$S \text{ returns } \begin{cases} \mathcal{E}(bs_b, \llbracket k_{j_b} \rrbracket) & \text{if } j_b < i \\ LR_{SE}(0, bs_b) & \text{if } j_b = i \\ \mathcal{E}(0, \llbracket k_{j_b} \rrbracket) & \text{if } j_b > i \end{cases}$$

Note that in the case where $j_b = i$, \mathcal{B}_b uses his left-right encryption oracle so S outputs $LR(0, bs_b)$ which is the encryption of either bs_b or 0 under the challenge key.

- If t_0 is a key name k then t_1 is the same key k and S outputs the generated value $\llbracket k \rrbracket$.
- If t_0 is a constant 0 or 1 then t_1 is the same constant and S outputs either $\llbracket 0 \rrbracket_{A_\eta}$ or $\llbracket 1 \rrbracket_{A_\eta}$.

We sum up in the following array how the algorithm S proceeds with encryption using a non-deducible key k_j (0 means that 0 is encrypted, bs_b indicates that the part corresponding

to t_b is encrypted).

	k_1	\dots	k_{i-1}	k_i	k_{i+1}	\dots	k_n
\mathcal{B}_b	bs_b	\dots	bs_b	$bs_b/0$	0	\dots	0

For key k_i , bs_b is used if the challenge bit of \mathcal{B}_b is 1, else 0 is used. We note that S never needs to use the value $\llbracket k_i \rrbracket$ and hence produces a coherent sequence of bitstrings for \mathcal{A} .

When the challenge bit of \mathcal{B}_0 is 1 and i equals n , \mathcal{A} is executed normally with challenge bit 0 and hence

$$\mathbb{P}[i = n : \mathcal{B}_0^{LR^1_{S\mathcal{E}}} = 1] = \mathbb{P}[\psi \xleftarrow{R} \llbracket \varphi_0 \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1]$$

When the challenge bit of \mathcal{B}_b is 1, algorithm S encrypts 0 for keys k_{i+1} to k_n . The same thing happens when the challenge bit of \mathcal{B}_b is 0 and the value $i + 1$ is generated instead of i . Hence

$$\mathbb{P}[i = x : \mathcal{B}_b^{LR^1_{S\mathcal{E}}} = 1] = \mathbb{P}[i = x + 1 : \mathcal{B}_b^{LR^0_{S\mathcal{E}}} = 1]$$

When the challenge bit of \mathcal{B}_b is 0 and i equals 1, then \mathcal{A} receives an interpretation for $\text{pat}(\varphi_b)$ (where $\llbracket \square \rrbracket_{A_\eta}$ is defined to be 0). As $\text{pat}(\varphi_0) =_\alpha \text{pat}(\varphi_1)$ we have that

$$\mathbb{P}[i = 1 : \mathcal{B}_0^{LR^0_{S\mathcal{E}}} = 1] = \mathbb{P}[i = 1 : \mathcal{B}_1^{LR^0_{S\mathcal{E}}} = 1]$$

Finally, when the challenge bit of \mathcal{B}_b is 1 and i equals n , \mathcal{A} is executed normally with challenge bit 0 and hence

$$\mathbb{P}[i = n : \mathcal{B}_1^{LR^1_{S\mathcal{E}}} = 1] = \mathbb{P}[\psi \xleftarrow{R} \llbracket \varphi_1 \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1]$$

From these equalities we get that

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{IND}}(\llbracket \varphi_0 \rrbracket_{A_\eta}, \llbracket \varphi_1 \rrbracket_{A_\eta}) &= \left| \mathbb{P}[\psi \xleftarrow{R} \llbracket \varphi_0 \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1] - \mathbb{P}[\psi \xleftarrow{R} \llbracket \varphi_1 \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1] \right| \\ &= \left| \mathbb{P}[i = n : \mathcal{B}_0^{LR^1_{S\mathcal{E}}} = 1] - \mathbb{P}[i = n : \mathcal{B}_1^{LR^1_{S\mathcal{E}}} = 1] \right| \\ &= \left| \mathbb{P}[i = n : \mathcal{B}_0^{LR^1_{S\mathcal{E}}} = 1] - \mathbb{P}[i = n : \mathcal{B}_0^{LR^0_{S\mathcal{E}}} = 1] \right. \\ &\quad + \mathbb{P}[i = n - 1 : \mathcal{B}_0^{LR^1_{S\mathcal{E}}} = 1] - \mathbb{P}[i = n - 1 : \mathcal{B}_0^{LR^0_{S\mathcal{E}}} = 1] \\ &\quad + \mathbb{P}[i = n - 2 : \mathcal{B}_0^{LR^1_{S\mathcal{E}}} = 1] - \\ &\quad \vdots \\ &\quad \left. - \mathbb{P}[i = 1 : \mathcal{B}_0^{LR^0_{S\mathcal{E}}} = 1] \right. \\ &\quad + \mathbb{P}[i = 1 : \mathcal{B}_1^{LR^0_{S\mathcal{E}}} = 1] - \mathbb{P}[i = 1 : \mathcal{B}_1^{LR^1_{S\mathcal{E}}} = 1] \\ &\quad + \mathbb{P}[i = 2 : \mathcal{B}_1^{LR^0_{S\mathcal{E}}} = 1] - \\ &\quad \vdots \\ &\quad \left. - \mathbb{P}[i = n - 1 : \mathcal{B}_1^{LR^1_{S\mathcal{E}}} = 1] \right. \\ &\quad \left. + \mathbb{P}[i = n : \mathcal{B}_1^{LR^0_{S\mathcal{E}}} = 1] - \mathbb{P}[i = n : \mathcal{B}_1^{LR^1_{S\mathcal{E}}} = 1] \right| \\ &\leq n \text{Adv}_{S\mathcal{E}, \mathcal{B}_1}^{cpa}(\eta) + n \text{Adv}_{S\mathcal{E}, \mathcal{B}_0}^{cpa}(\eta) \end{aligned}$$

As the advantage against IND-CPA is negligible the advantage of \mathcal{A} against soundness of \equiv_{pat} is negligible as well.

Proposition 13.1 now follows directly from Lemmas 13.1 and 13.2.

Exercise 13.2 For our soundness result we need the hypothesis that frames do not contain key cycles. The hypothesis is indeed used in the proof. However, not only the proof would be wrong, but also the result. Show that there exists an IND-CPA secure encryption scheme such that A_η relies on this scheme and

$$\llbracket \{0\}_k \rrbracket_{A_\eta} \not\approx \llbracket \{k\}_k \rrbracket_{A_\eta}.$$

Exercise 13.3 Suppose that we change the definition of patterns (Definition 13.5) by replacing the line

$$\text{pat}(\{t\}_k, K) = \{\square\}_k \quad \text{if } k \notin K$$

with

$$\text{pat}(\{t\}_k, K) = \square \quad \text{if } k \notin K$$

We will denote by pat' the resulting definition.

1. Give an example of two frames φ_1, φ_2 such that $\text{pat}'(\varphi_1) = \text{pat}'(\varphi_2)$ and $\text{pat}(\varphi_1) \neq \text{pat}(\varphi_2)$.
2. Show that Lemma 13.2 (soundness of pattern equivalence) does not hold for the above definition of patterns.
3. How could one strengthen the definition IND-CPA security for such a soundness result to hold?

Chapter 14

Soundness in presence of an Adaptive Adversary

14.1 Adaptive Adversary

In the previous sections we only considered a purely passive adversary, i.e., the adversary observes the computational value of a fixed frame. A stronger adversary model is that of an *adaptive* adversary: we allow the adversary to stepwise construct the frames which are submitted to a left-right oracle. Let (A_η) be a family of computational algebras and \mathcal{A} be an adversary. \mathcal{A} has access to a left-right evaluation oracle \mathcal{O}_{LR} which given a pair of symbolic terms (t_0, t_1) outputs either the implementation of t_0 or of t_1 . This oracle depends on a selection bit b and uses a local store in order to record values generated for the different names (these values are used when processing further queries). With a slight abuse of notation, we omit this store and write:

$$\mathcal{O}_{LR, A_\eta}^b(t_0, t_1) = \llbracket t_b \rrbracket_{A_\eta}$$

Adversary \mathcal{A} plays an indistinguishability game and its objective is to find the value of b . The advantage of \mathcal{A} is defined by:

$$\text{Adv}_{\mathcal{A}, A_\eta}^{\text{adpt}}(\eta) = \left| \mathbb{P}[\mathcal{A}^{\mathcal{O}_{LR, A_\eta}^1}(\eta) = 1] - \mathbb{P}[\mathcal{A}^{\mathcal{O}_{LR, A_\eta}^0}(\eta) = 1] \right|$$

Without further restrictions on the queries made by the adversary, having a non-negligible advantage is easy in most cases. For example the adversary could submit a pair $(0, 1)$ to his oracle. We therefore require the adversary to be *legal*.

Definition 14.1 *An adversary \mathcal{A} is legal if for any sequence of queries $(t_0^i, t_1^i)_{1 \leq i \leq n}$ made by \mathcal{A} to its left-right oracle, queries are statically equivalent:*

$$\{x_1 \mapsto t_0^1, \dots, x_n \mapsto t_0^n\} \sim_{\mathcal{E}} \{x_1 \mapsto t_1^1, \dots, x_n \mapsto t_1^n\}$$

A family of computational algebras (A_η) is $\sim_{\mathcal{E}}$ -ad-sound iff the advantage $\text{Adv}_{\mathcal{A}, A_\eta}^{\text{adpt}}(\eta)$ of any probabilistic polynomial time legal adversary \mathcal{A} is negligible.

We indeed have that adaptive soundness is strictly stronger than the non-adaptive soundness.

Proposition 14.1 *Let (A_η) be a family of computational algebras. If (A_η) is $\sim_\mathcal{E}$ -ad-sound then (A_η) is also $\sim_\mathcal{E}$ -sound. The converse is not true in general.*

Proof : The implication that adaptive soundness implies soundness is easy to prove: let $\varphi = \{x_i \mapsto t_i\}_{1 \leq i \leq n}$ and $\varphi' = \{x_i \mapsto u_i\}_{1 \leq i \leq n}$ be two frames such that $\varphi \sim_\mathcal{E} \varphi'$. Let \mathcal{A} be an adversary against soundness for φ and φ' , (i.e., \mathcal{A} tries to distinguish implementations of φ from implementations of φ') we build an adversary \mathcal{B} against adaptive soundness which uses his oracle to obtain an implementation of φ or φ' .

Adversary $\mathcal{B}^{\mathcal{O}_{LR}}$:

for i **from** 1 **to** n
 $bs_i \leftarrow \mathcal{O}_{LR}(t_i, u_i)$
return $\mathcal{A}(\{x_i \mapsto bs_i\}_{1 \leq i \leq n})$

When the challenge bit of \mathcal{B} is 0, \mathcal{A} is given an implementation of φ whereas when it is 1, \mathcal{A} is given an implementation of φ' . The advantage of \mathcal{B} is:

$$\begin{aligned} \text{Adv}_{\mathcal{B}, A_\eta}^{\text{adpt}}(\eta) &= \left| \mathbb{P}[\mathcal{B}^{\mathcal{O}_{LR, A_\eta}^1} = 1] - \mathbb{P}[\mathcal{B}^{\mathcal{O}_{LR, A_\eta}^0} = 1] \right| \\ &= \left| \mathbb{P}[\hat{\phi} \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \mathcal{A}(\eta, \hat{\phi})] - \mathbb{P}[\hat{\phi} \leftarrow \llbracket \varphi' \rrbracket_{A_\eta} : \mathcal{A}(\eta, \hat{\phi})] \right| \\ &= \text{Adv}_{\mathcal{A}}^{\text{IND}}(\eta, \llbracket \varphi \rrbracket_{A_\eta}, \llbracket \varphi' \rrbracket_{A_\eta}) \end{aligned}$$

As we assume (A_η) to be $\sim_\mathcal{E}$ -ad-sound, the advantage of \mathcal{B} is negligible, so we obtain that $\llbracket \varphi \rrbracket_{A_\eta}$ and $\llbracket \varphi' \rrbracket_{A_\eta}$ are indistinguishable and (A_η) is $\sim_\mathcal{E}$ -sound.

To prove that the converse statement does not hold, we give a counter example based on the following theory. We consider five sorts: Nonces, Bit, Bool, BS and Data. We suppose that Bit is a subsort of BS (hence we require that $\llbracket \text{Bit} \rrbracket_A \subseteq \llbracket \text{BS} \rrbracket_A$) and Data is a supersort of all others ($\llbracket \text{Data} \rrbracket_A$ is the set of all bitstrings). Consider the following symbols:

$$\begin{aligned} 0, 1 &: \text{Bit} \\ \text{cons} &: \text{Bit} \times \text{BS} \rightarrow \text{BS} \\ \text{eqnonce} &: \text{BS} \times \text{Nonces} \times \text{Nonces} \rightarrow \text{Bool} \\ \text{tbool} &: \text{Data} \rightarrow \text{Bit} \end{aligned}$$

For soundness, we consider frames to be well-formed if they contain terms of sort Nonces or of sort Bool and no occurrences of **tbool**. In particular we assume that the frame does not contain any names that are not of sort Nonces. The equational theory is generated by the equation $\text{tbool}(\text{eqnonce}(x, y)) = 1$.

As an implementation, we define the computational algebras A_η : the concrete domain $\llbracket \text{Nonces} \rrbracket_{A_\eta}$ is $\{0, 1\}^\eta$ equipped with the uniform distribution; the constants 0 and 1 are respectively interpreted with 0 and 1. The implementation of sort Bool contains the same values 0 and 1. **cons** is interpreted by bit concatenation, $\llbracket \text{eqnonce} \rrbracket(bs, N, N')$ outputs 1 if the computational interpretations of bs and N are the same, 0 otherwise (it ignores the nonce N'). We easily build an adversary \mathcal{A} against adaptive soundness that has an advantage of 1: the first query of \mathcal{A} is (N, N) . \mathcal{A} obtains the bit-string value bs of N and produces term t of type BS which has the same implementation. \mathcal{A} also produces a term t' equal to t except that its first bit has been flipped. Then \mathcal{A} queries his oracle with $(\text{eqnonce}(t', N, N'), \text{eqnonce}(t, N, N'))$. (Note that the second nonce argument of **eqnonce** is needed to ensure that the queries yield

statically equivalent frames). If the oracle answer is 1, \mathcal{A} outputs 1, else \mathcal{A} outputs 0, in both cases \mathcal{A} has correctly guessed his challenge bit. Hence his advantage is 1.

However we prove that (A_η) is $\sim_{\mathcal{E}}$ -sound, *i.e.*, the advantage of any adversary \mathcal{A} in distinguishing two statically equivalent frame is negligible. We consider two statically equivalent frames φ and φ' . Without loss of generality (renaming of nonces and variables), φ has the form:

$$\{x_1 \mapsto N_1, \dots, x_k \mapsto N_k, y_1 \mapsto \text{eqnonce}(bs_1, M_{1,1}, M_{1,2}), \dots, y_\ell \mapsto \text{eqnonce}(bs_\ell, M_{\ell,1}, M_{\ell,2})\}$$

and φ' has the form:

$$\{x_1 \mapsto N_1, \dots, x_k \mapsto N_k, y_1 \mapsto \text{eqnonce}(bs'_1, M'_{1,1}, M'_{1,2}), \dots, y_\ell \mapsto \text{eqnonce}(bs'_\ell, M'_{\ell,1}, M'_{\ell,2})\}$$

For any i , the probability to have y_i different from 0 in $\llbracket \varphi \rrbracket_{A_\eta}$ is lower than $1/2^\eta$. Hence the probability for one of the y_i to be different from 0 is lower than $\ell/2^\eta$. Let us denote by $P(\psi)$ the event that for all i ($1 \leq i \leq \ell$) $y_i\psi = 0$. We have that

$$\mathbb{P}[\psi \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \neg P(\psi)] \leq \frac{\ell}{2^\eta}$$

Thus with overwhelming probability all the y_i are equal to 0. The same holds for φ' . This allows us to conclude that the advantage of \mathcal{A} is negligible.

$$\begin{aligned} \text{Adv}_{\mathcal{A}}^{\text{IND}} &= |\mathbb{P}[\psi \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1] - \mathbb{P}[\psi' \leftarrow \llbracket \varphi' \rrbracket_{A_\eta} : \mathcal{A}(\psi') = 1]| \\ &= |\mathbb{P}[\psi \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1 \wedge P(\psi)] + \mathbb{P}[\psi \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1 \wedge \neg P(\psi)] \\ &\quad - \mathbb{P}[\psi' \leftarrow \llbracket \varphi' \rrbracket_{A_\eta} : \mathcal{A}(\psi') = 1 \wedge P(\psi')] - \mathbb{P}[\psi' \leftarrow \llbracket \varphi' \rrbracket_{A_\eta} : \mathcal{A}(\psi') = 1 \wedge \neg P(\psi')]| \\ &\leq |\mathbb{P}[\psi \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1 \wedge P(\psi)] - \mathbb{P}[\psi' \leftarrow \llbracket \varphi' \rrbracket_{A_\eta} : \mathcal{A}(\psi') = 1 \wedge P(\psi')]| + \\ &\quad |\mathbb{P}[\psi \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1 \wedge \neg P(\psi)] - \mathbb{P}[\psi' \leftarrow \llbracket \varphi' \rrbracket_{A_\eta} : \mathcal{A}(\psi') = 1 \wedge \neg P(\psi')]| \\ &\leq |\mathbb{P}[\psi \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1 \wedge P(\psi)] - \mathbb{P}[\psi' \leftarrow \llbracket \varphi' \rrbracket_{A_\eta} : \mathcal{A}(\psi') = 1 \wedge P(\psi)]| + \\ &\quad |\mathbb{P}[\psi \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1 \wedge \neg P(\psi)] + \mathbb{P}[\psi' \leftarrow \llbracket \varphi' \rrbracket_{A_\eta} : \mathcal{A}(\psi') = 1 \wedge \neg P(\psi')]| \\ &\leq |\mathbb{P}[\psi \leftarrow \llbracket \varphi \rrbracket_{A_\eta} : \mathcal{A}(\psi) = 1 \wedge P(\psi)] - \mathbb{P}[\psi' \leftarrow \llbracket \varphi' \rrbracket_{A_\eta} : \mathcal{A}(\psi') = 1 \wedge P(\psi')]| + \frac{2\ell}{2^\eta} \end{aligned}$$

As the distributions of ψ and ψ' when $P(\psi)$ and $P(\psi')$ are true are identical, the advantage of \mathcal{A} is bounded by $2\ell/2^\eta$, which is negligible. Hence, (A_η) is $\sim_{\mathcal{E}}$ -sound.

Let us now get back to the equational theory sym and its corresponding computational algebra which we showed to be \sim_{sym} -sound in the previous section. We will also show adaptive soundness for this computational algebra.

Well formed frames and adversaries. When showing non-adaptive soundness we have already seen that we need to make some hypotheses on frames. In particular, we needed to forbid key cycles. As the adversary can make queries dynamically we need to fix a total order on keys \prec a priori. Then a frame φ is acyclic for \prec if whenever $\{t\}_{k_2} \in st(\varphi)$ and $k_1 \in st(\varphi)$ we have that $k_1 \prec k_2$.

Definition 14.2 A frame $\varphi = \{x_1 \mapsto t_1, \dots, x_n \mapsto t_n\}$ is well-formed for \prec if

- $K(\varphi)$ only contains names;
- φ is acyclic for \prec ;
- the terms t_i do not contain destructors, i.e. they only use symbols `senc`, `pair`, `0` and `1`,
- if k is used as a plaintext in t_i then k cannot be used at a key position in t_j for $j < i$.

An adversary is well-formed for \prec if his queries $(t_1^0, t_1^1), \dots, (t_n^0, t_n^1)$ are such that $\varphi_i = \{x_1 \mapsto t_1^i, \dots, x_n \mapsto t_n^i\}$ for $i \in \{0, 1\}$ are well-formed frames for \prec .

We will see in the proof that the additional assumption on frames “if k is used as a plaintext in t_i then k cannot be used at a key position in t_j for $j < i$ ” is needed to avoid a *selective decommitment* problem.

Proposition 14.2 *Let A_η be a computational algebra based on an IND-CPA secure encryption scheme and let \prec be a total order on names. Then we have that (A_η) is $\sim_{\text{sym-ad-sound}}$ when we consider well-formed adversaries, i.e., for any well-formed probabilistic polynomial time legal adversary \mathcal{A} the advantage $\text{Adv}_{\mathcal{A}, A_\eta}^{\text{adpt}}(\eta)$ of is negligible in η .*

Proof : The proof is an easy adaptation of the proof of Proposition 13.1. We have that the number of keys n appearing in the queries is polynomial in η (as the adversary \mathcal{A} is polynomially bounded in η). We then denote the keys as k_1, \dots, k_n such that $i < j$ implies that $i \prec j$. The rest of the proof is similar.

In particular we note that for an adaptive adversary we need the additional assumption “if k is used as a plaintext in t_i then k cannot be used at a key position in t_j for $j < i$ ”. Otherwise an adversary could make requests $\{k_2\}_{k_3}$ followed by $\{k_3\}_{k_1}$. While the corresponding frame does not contain a key cycle an adversary \mathcal{B} which has drawn the challenge key to be k_3 would not be able to respond in a coherent way: the first encryption would be answered by a call to the encryption oracle, while the adversary would need to know the value of k_3 to answer the second query. Also note that it may be possible that \mathcal{B} uses the value it generated himself for the challenge key k_i . However, in that case, the condition ensures that \mathcal{B} will never query the oracle.

Exercise 14.1 *Let IND-CPA-1 be a 1 query variant of IND-CPA. The definition of IND-CPA-1, uses a left-right oracle $LR1_{\mathcal{SE}}^b$. This oracle first generates a key k using \mathcal{KG} . Then it answers a single query of the form $(bs_0^1, bs_1^1, \dots, bs_0^n, bs_1^n)$, where bs_0^i and bs_1^i are bit-strings. The oracle returns the sequence of ciphertexts $\mathcal{E}(bs_b^1, k), \dots, \mathcal{E}(bs_b^n, k)$. The goal of the adversary \mathcal{A} is to guess the value of bit b . His advantage is defined as:*

$$\text{Adv}_{\mathcal{SE}, \mathcal{A}}^{\text{cpa1}}(\eta) = \left| \mathbb{P}[\mathcal{A}^{LR1_{\mathcal{SE}}^1}(\eta) = 1] - \mathbb{P}[\mathcal{A}^{LR1_{\mathcal{SE}}^0}(\eta) = 1] \right|$$

Encryption scheme \mathcal{SE} is IND-CPA-1 secure if the advantage of any adversary \mathcal{A} is negligible in η .

1. Let φ_1, φ_2 be two well-formed frames (Definition 13.4) and A_η a computational algebra based on an IND-CPA-1 secure encryption scheme. Show that $\varphi_1 \sim_{\text{sym}} \varphi_2$ implies that $\llbracket \varphi_1 \rrbracket_{A_\eta} \approx \llbracket \varphi_2 \rrbracket_{A_\eta}$.

2. *Show that if \mathcal{SE} is an IND-CPA secure encryption scheme then \mathcal{SE} is also IND-CPA-1 secure.*
3. *Give an example of an IND-CPA-1 secure encryption scheme which is not IND-CPA secure.*

14.2 Further readings

The material covered in this chapter is based on [?, ?, ?]. [?] is actually the first soundness of this kind. The paper includes a detailed proof and many examples. The soundness result is directly on a dedicated pattern equivalence and not on static equivalence. A more detailed course focusing exclusively on soundness results in the presence of a passive adversary is available in [?]. For an intensive survey with many more references we refer the interested reader to the survey paper [?]. This survey discusses extensions of such soundness results to many more primitives. It also gives many references to soundness results in the case of an active adversary. In particular, in [?] a soundness result for symmetric encryption and an active adversary is presented for a process calculus close to the one presented in these lecture notes.

Chapter 15

Soundness in presence of an Active Adversary

Part V

Implementation

Bibliography

- [1] Spore, the security protocol open repository. <http://www.lsv.ens-cachan.fr/spore>.
- [2] M. Abadi and B. Blanchet. Secrecy types for asymmetric communication. In F. Honsell and M. Miculan, editors, *Foundations of Software Science and Computation Structures (FoSSaCS 2001)*, volume 2030 of *Lecture Notes on Computer Science*, pages 25–41, Genova, Italy, Apr. 2001. Springer.
- [3] M. Abadi and B. Blanchet. Analyzing security protocols with secrecy types and logic programs. *Journal of the ACM*, 52(1):102–146, Jan. 2005.
- [4] M. Abadi and B. Blanchet. Computer-assisted verification of a protocol for certified email. *Science of Computer Programming*, 58(1-2):3–27, Oct. 2005. Special issue SAS’03.
- [5] M. Abadi, B. Blanchet, and C. Fournet. Just fast keying in the pi calculus. *ACM Transactions on Information and System Security (TISSEC)*, 10(3):1–59, 2007.
- [6] M. Abadi and V. Cortier. Deciding knowledge in security protocols under equational theories. *Theoretical Computer Science*, 387(1-2):2–32, 2006.
- [7] M. Abadi and C. Fournet. Mobile values, new names, and secure communication. In *Proc. 28th Symposium on Principles of Programming Languages (POPL’01)*, pages 104–115. ACM Press, 2001.
- [8] M. Abadi and C. Fournet. Private authentication. *Theoretical Computer Science*, 322(3):427–476, 2004.
- [9] M. Abadi, N. Glew, B. Horne, and B. Pinkas. Certified email with a light on-line trusted third party: Design and implementation. In *11th International World Wide Web Conference*, pages 387–395, Honolulu, Hawaii, May 2002. ACM Press.
- [10] M. Abadi and A. Gordon. A calculus for cryptographic protocols: the spi calculus. *Information and Computation*, 148(1), 1999.
- [11] M. Abadi and R. M. Needham. Prudent engineering practice for cryptographic protocols. *IEEE Trans. Software Eng.*, 22(1):6–15, 1996.
- [12] M. Abadi and P. Rogaway. Reconciling two views of cryptography (the computational soundness of formal encryption). *Journal of Cryptology*, 15(2):103–127, 2002.

- [13] W. Aiello, S. M. Bellovin, M. Blaze, R. Canetti, J. Ioannidis, K. Keromytis, and O. Rein-gold. Just Fast Keying: Key agreement in a hostile Internet. *ACM Transactions on Information and System Security*, 7(2):242–273, May 2004.
- [14] X. Allamigeon and B. Blanchet. Reconstruction of attacks against cryptographic protocols. In *18th IEEE Computer Security Foundations Workshop (CSFW-18)*, pages 140–154, Aix-en-Provence, France, June 2005. IEEE.
- [15] American National Standards Institute. Public Key Cryptography for the Financial Services Industry: The Elliptic Curve Digital Signature Algorithm. ANSI X9.62-1998. January 1999.
- [16] R. Anderson and R. Needham. Programming Satan’s computer. In J. van Leeuwen, editor, *Computer Science Today: Recent Trends and Developments*, volume 1000 of *Lecture Notes on Computer Science*, pages 426–440. Springer, 1995.
- [17] S. Andova, C. Cremers, K. G. Steen, S. Mauw, S. M. Isnes, and S. Radomirović. Sufficient conditions for composing security protocols. *Information and Computation*, 206(2-4):425–459, 2008.
- [18] M. Arapinis, S. Delaune, and S. Kremer. From one session to many: Dynamic tags for security protocols. In I. Cervesato, H. Veith, and A. Voronkov, editors, *Proceedings of the 15th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR’08)*, volume 5330 of *Lecture Notes in Artificial Intelligence*, pages 128–142, Doha, Qatar, Nov. 2008. Springer.
- [19] M. Arapinis and M. Duflot. Bounding messages for free in security protocols. In *Proc. 27th Conference on Foundations of Software Technology and Theoretical Computer Science (FST&TCS’07)*, volume 4855 of *LNCS*, pages 376–387. Springer, 2007.
- [20] A. Armando, D. A. Basin, Y. Boichut, Y. Chevalier, L. Compagna, J. Cuéllar, P. H. Drielsma, P.-C. Héam, O. Kouchnarenko, J. Mantovani, S. Mödersheim, D. von Oheimb, M. Rusinowitch, J. Santiago, M. Turuani, L. Viganò, and L. Vigneron. The avispa tool for the automated validation of internet security protocols and applications. In *Computer Aided Verification*, volume 3576 of *Lecture Notes in Computer Science*. Springer, 2005.
- [21] M. Arnaud, V. Cortier, and S. Delaune. Combining algorithms for deciding knowledge in security protocols. In F. Wolter, editor, *Proceedings of the 6th International Symposium on Frontiers of Combining Systems (FroCoS’07)*, volume 4720 of *Lecture Notes in Artificial Intelligence*, pages 103–117, Liverpool, UK, Sept. 2007. Springer.
- [22] L. Bachmair and H. Ganzinger. Resolution theorem proving. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume 1, chapter 2, pages 19–100. North Holland, 2001.
- [23] M. Backes, C. Hritcu, and M. Maffei. Automated verification of remote electronic voting protocols in the applied pi-calculus. In *Proc. 21st IEEE Computer Security Foundations Symposium, (CSF’08)*, pages 195–209. IEEE Comp. Soc. Press, 2008.

- [24] M. Backes, M. Maffei, and D. Unruh. Zero-knowledge in the applied pi-calculus and automated verification of the direct anonymous attestation protocol. In *29th IEEE Symposium on Security and Privacy*, pages 202–215, Oakland, CA, May 2008. IEEE. Technical report version available at <http://eprint.iacr.org/2007/289>.
- [25] B. Barak, R. Canetti, J. Nielsen, and R. Pass. Universally composable protocols with relaxed set-up assumptions. In *Proc. 45th Symposium on Foundations of Computer Science (FOCS'04)*, pages 186–195. IEEE Comp. Soc. Press, 2004.
- [26] M. Baudet. Deciding security of protocols against off-line guessing attacks. In *Proc. 12th Conference on Computer and Communications Security*, pages 16–25. ACM Press, 2005.
- [27] M. Baudet. *Sécurité des protocoles cryptographiques : aspects logiques et calculatoires*. Thèse de doctorat, LSV, ENS Cachan, France, Jan. 2007.
- [28] M. Baudet, V. Cortier, and S. Delaune. YAPA: A generic tool for computing intruder knowledge. In R. Treinen, editor, *Proceedings of the 20th International Conference on Rewriting Techniques and Applications (RTA'09)*, volume 5595 of *Lecture Notes in Computer Science*, pages 148–163, Brasília, Brazil, June-July 2009. Springer.
- [29] M. Baudet, V. Cortier, and S. Kremer. Computationally sound implementations of equational theories against passive adversaries. *Information and Computation*, 207(4):496–520, Apr. 2009.
- [30] O. Baudron, D. Pointcheval, and J. Stern. Extended Notions of Security for Multicast Public Key Cryptosystems. In *Proc. of the 27th ICALP*, LNCS 1853, pages 499–511. Springer-Verlag, Berlin, 2000.
- [31] M. Bellare. Practice-Oriented Provable Security. In *ISW '97*, LNCS 1396. Springer-Verlag, Berlin, 1997.
- [32] M. Bellare, A. Boldyreva, and S. Micali. Public-key Encryption in a Multi-User Setting: Security Proofs and Improvements. In *Eurocrypt '00*, LNCS 1807, pages 259–274. Springer-Verlag, Berlin, 2000.
- [33] M. Bellare, A. Boldyreva, and A. Palacio. A Separation between the Random-Oracle Model and the Standard Model for a Hybrid Encryption Problem, 2003. Cryptology ePrint Archive 2003/077.
- [34] M. Bellare, A. Desai, D. Pointcheval, and P. Rogaway. Relations among Notions of Security for Public-Key Encryption Schemes. In *Crypto '98*, LNCS 1462, pages 26–45. Springer-Verlag, Berlin, 1998.
- [35] M. Bellare and A. Palacio. GQ and Schnorr Identification Schemes: Proofs of Security against Impersonation under Active and Concurrent Attacks. In *Crypto '02*, LNCS 2442, pages 162–177. Springer-Verlag, Berlin, 2002.
- [36] M. Bellare, D. Pointcheval, and P. Rogaway. Authenticated Key Exchange Secure Against Dictionary Attacks. In *Eurocrypt '00*, LNCS 1807, pages 139–155. Springer-Verlag, Berlin, 2000.

- [37] M. Bellare and P. Rogaway. Random Oracles Are Practical: a Paradigm for Designing Efficient Protocols. In *Proc. of the 1st CCS*, pages 62–73. ACM Press, New York, 1993.
- [38] M. Bellare and P. Rogaway. Optimal Asymmetric Encryption – How to Encrypt with RSA. In *Eurocrypt '94*, LNCS 950, pages 92–111. Springer-Verlag, Berlin, 1995.
- [39] M. Bellare and P. Rogaway. The Exact Security of Digital Signatures – How to Sign with RSA and Rabin. In *Eurocrypt '96*, LNCS 1070, pages 399–416. Springer-Verlag, Berlin, 1996.
- [40] M. Bellare and P. Rogaway. The Game-Playing Technique and its Application to Triple Encryption, 2004. Cryptology ePrint Archive 2004/331.
- [41] S. M. Bellovin and M. Merritt. Encrypted key exchange: Password-based protocols secure against dictionary attacks. In *Proc. Symposium on Security and Privacy (SP'92)*, pages 72–84. IEEE Comp. Soc., 1992.
- [42] K. Bhargavan, R. Corin, C. Fournet, and A. Gordon. Secure sessions for web services. In *ACM Workshop on Secure Web Services (SWS'04)*, Washington DC, Oct. 2004.
- [43] K. Bhargavan, R. Corin, C. Fournet, and E. Zălinescu. Cryptographically verified implementations for TLS. In *Proceedings of the 15th ACM Conference on Computer and Communications Security (CCS'08)*, pages 459–468. ACM, Oct. 2008.
- [44] K. Bhargavan, C. Fournet, and A. Gordon. Verifying policy-based security for web services. In *ACM Conference on Computer and Communications Security (CCS'04)*, pages 268–277, Washington DC, Oct. 2004. ACM.
- [45] K. Bhargavan, C. Fournet, and A. Gordon. Verified reference implementations of WS-Security protocols. In M. Bravetti, M. Núñez, and G. Zavattaro, editors, *3rd International Workshop on Web Services and Formal Methods (WS-FM 2006)*, volume 4184 of *Lecture Notes on Computer Science*, pages 88–106, Vienna, Austria, Sept. 2006. Springer.
- [46] K. Bhargavan, C. Fournet, A. Gordon, and N. Swamy. Verified implementations of the information card federated identity-management protocol. In *ACM Symposium on Information, Computer and Communications Security (ASIACCS'08)*, pages 123–135, Tokyo, Japan, Mar. 2008. ACM.
- [47] K. Bhargavan, C. Fournet, A. Gordon, and S. Tse. Verified interoperable implementations of security protocols. In *19th IEEE Computer Security Foundations Workshop (CSFW'06)*, pages 139–152, Venice, Italy, July 2006. IEEE Computer Society.
- [48] K. Bhargavan, C. Fournet, A. D. Gordon, and R. Pucella. TulaFale: A security tool for web services. In F. S. de Boer, M. M. Bonsangue, S. Graf, and W.-P. de Roever, editors, *Formal Methods for Components and Objects (FMCO 2003)*, volume 3188 of *Lecture Notes on Computer Science*, pages 197–222, Leiden, The Netherlands, Nov. 2003. Springer. Paper and tool available at <http://securing.ws/>.
- [49] E. Biham and A. Shamir. Differential Fault Analysis of Secret Key Cryptosystems. In *Crypto '97*, LNCS 1294, pages 513–525. Springer-Verlag, Berlin, 1997.

- [50] B. Blanchet. *ProVerif Automatic Cryptographic Protocol Verifier User Manual*. <http://www.proverif.ens.fr/proverif-manual.pdf>.
- [51] B. Blanchet. An efficient cryptographic protocol verifier based on Prolog rules. In *Computer Security Foundations Workshop (CSFW'01)*, 2001.
- [52] B. Blanchet. Security protocols: From linear to classical logic by abstract interpretation. *Information Processing Letters*, 95(5):473–479, Sept. 2005.
- [53] B. Blanchet. Automatic verification of correspondences for security protocols. Report arXiv:0802.3444v1, 2008. Available at <http://arxiv.org/abs/0802.3444v1>.
- [54] B. Blanchet. Automatic verification of correspondences for security protocols. *Journal of Computer Security*, 17(4):363–434, July 2009.
- [55] B. Blanchet, M. Abadi, and C. Fournet. Automated verification of selected equivalences for security protocols. *Journal of Logic and Algebraic Programming*, 75(1):3–51, Feb.–Mar. 2008.
- [56] B. Blanchet and A. Chaudhuri. Automated formal analysis of a protocol for secure file sharing on untrusted storage. In *IEEE Symposium on Security and Privacy*, pages 417–431, Oakland, CA, May 2008. IEEE.
- [57] B. Blanchet and A. Podelski. Verification of cryptographic protocols: Tagging enforces termination. *Theoretical Computer Science*, 333(1-2):67–90, Mar. 2005. Special issue FoSSaCS'03.
- [58] D. Bleichenbacher. Generating El Gamal Signatures without Knowing the Secret Key. In *Eurocrypt '96*, LNCS 1070, pages 10–18. Springer-Verlag, Berlin, 1996.
- [59] D. Bleichenbacher. A Chosen Ciphertext Attack against Protocols based on the RSA Encryption Standard PKCS #1. In *Crypto '98*, LNCS 1462, pages 1–12. Springer-Verlag, Berlin, 1998.
- [60] C. Bodei. *Security Issues in Process Calculi*. PhD thesis, Università di Pisa, Jan. 2000.
- [61] D. Boneh, R. DeMillo, and R. Lipton. On the Importance of Checking Cryptographic Protocols for Faults. In *Eurocrypt '97*, LNCS 1233, pages 37–51. Springer-Verlag, Berlin, 1997.
- [62] J. Borgström. Static equivalence is harder than knowledge. *Electr. Notes Theor. Comput. Sci.*, 154(3):45–57, 2006.
- [63] X. Boyen. The uber-assumption family. In *Pairing '08*, LNCS 5209, pages 39–56. Springer-Verlag, Berlin, 2008.
- [64] E. Brickell, D. Pointcheval, S. Vaudenay, and M. Yung. Design Validations for Discrete Logarithm Based Signature Schemes. In *PKC '00*, LNCS 1751, pages 276–292. Springer-Verlag, Berlin, 2000.
- [65] D. R. L. Brown and D. B. Johnson. Formal Security Proofs for a Signature Scheme with Partial Message Recovery. In *CT – RSA '01*, LNCS 2020, pages 126–142. Springer-Verlag, Berlin, 2001.

- [66] M. Burrows, M. Abadi, and R. Needham. A logic of authentication. *Proceedings of the Royal Society of London A*, 426(1871):233–271, dec 1989. A preliminary version appeared as Digital Equipment Corporation Systems Research Center report No. 39, February 1989.
- [67] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *Proc. 42nd Annual Symposium on Foundations of Computer Science (FOCS'01)*, pages 136–145. IEEE Comp. Soc., 2001.
- [68] R. Canetti, Y. Dodis, R. Pass, and S. Walfish. Universally composable security with global setup. In *Proc. 4th Theory of Cryptography Conference, (TCC'07)*, LNCS, pages 61–85. Springer, 2007.
- [69] R. Canetti, O. Goldreich, and S. Halevi. The Random Oracle Methodology, Revisited. In *Proc. of the 30th STOC*, pages 209–218. ACM Press, New York, 1998.
- [70] R. Canetti and J. Herzog. Universally composable symbolic analysis of mutual authentication and key exchange protocols. In S. Halevi and T. Rabin, editors, *Proceedings, Theory of Cryptography Conference (TCC'06)*, volume 3876 of *Lecture Notes on Computer Science*, pages 380–403, New York, NY, Mar. 2006. Springer. Extended version available at <http://eprint.iacr.org/2004/334>.
- [71] R. Canetti, C. Meadows, and P. F. Syverson. Environmental requirements for authentication protocols. In *Proc. Symposium on Software Security – Theories and Systems*, volume 2609 of *LNCS*, pages 339–355. Springer, 2002.
- [72] R. Canetti and T. Rabin. Universal composition with joint state. In *Proc. 23rd International Cryptology Conference (CRYPTO'03)*, LNCS, pages 265–281. Springer, 2003.
- [73] L. Cardelli, G. Ghelli, and A. D. Gordon. Secrecy and group creation. In C. Palamidessi, editor, *CONCUR 2000: Concurrency Theory*, volume 1877 of *Lecture Notes on Computer Science*, pages 365–379. Springer, Aug. 2000.
- [74] S. Cavallar, B. Dodson, A. K. Lenstra, W. Lioen, P. L. Montgomery, B. Murphy, H. te Riele, K. Aardal, J. Gilchrist, G. Guillerm, P. Leyland, J. Marchand, F. Morain, A. Muffett, C. Putnam, C. Putnam, and P. Zimmermann. Factorization of a 512-bit RSA Modulus. In *Eurocrypt '00*, LNCS 1807, pages 1–18. Springer-Verlag, Berlin, 2000.
- [75] Y. Chevalier, R. Küsters, M. Rusinowitch, and M. Turuani. Deciding the security of protocols with Diffie-Hellman exponentiation and products in exponents. In P. K. Pandya and J. Radhakrishnan, editors, *FST TCS 2003: Foundations of Software Technology and Theoretical Computer Science, 23rd Conference*, volume 2914 of *Lecture Notes on Computer Science*, pages 124–135, Mumbai, India, Dec. 2003. Springer.
- [76] Y. Chevalier, R. Küsters, M. Rusinowitch, and M. Turuani. An NP decision procedure for protocol insecurity with XOR. *Theoretical Computer Science*, 338(1–3):247–274, June 2005.
- [77] B. Chor and R. L. Rivest. A Knapsack Type Public Key Cryptosystem based on Arithmetic in Finite Fields. In *Crypto '84*, LNCS 196, pages 54–65. Springer-Verlag, Berlin, 1985.

- [78] Ș. Ciobâcă and V. Cortier. Protocol composition for arbitrary primitives. In *Proceedings of the 23rd IEEE Computer Security Foundations Symposium (CSF'10)*, Edinburgh, Scotland, UK, July 2010. IEEE Computer Society Press. To appear.
- [79] Ș. Ciobâcă, S. Delaune, and S. Kremer. Computing knowledge in security protocols under convergent equational theories. In R. Schmidt, editor, *Proceedings of the 22nd International Conference on Automated Deduction (CADE'09)*, Lecture Notes in Artificial Intelligence, pages 355–370, Montreal, Canada, Aug. 2009. Springer.
- [80] J. Clark and J. Jacob. A survey of authentication protocol literature. <http://www.cs.york.ac.uk/~jac/papers/drareviewps.ps>, 1997.
- [81] J. Clulow. The design and analysis of cryptographic APIs for security devices. Master's thesis, University of Natal, Durban, South Africa, 2003. Chapter 3.
- [82] E. Cohen. Proving protocols safe from guessing. In *Foundations of Computer Security*, Copenhagen, Denmark, July 2002.
- [83] H. Comon-Lundh. Soundness of abstract cryptography. Lecture Notes, 2007. <http://www.lsv.ens-cachan.fr/~comon/Soundness/>.
- [84] H. Comon-Lundh and V. Cortier. New decidability results for fragments of first-order logic and application to cryptographic protocols. In R. Nieuwenhuis, editor, *14th Int. Conf. Rewriting Techniques and Applications (RTA'2003)*, volume 2706 of *Lecture Notes on Computer Science*, pages 148–164, Valencia, Spain, June 2003. Springer.
- [85] H. Comon-Lundh and V. Cortier. Security properties: two agents are sufficient. *Science of Computer Programming*, 50(1-3):51–71, 2004.
- [86] H. Comon-Lundh and V. Cortier. Computational soundness of observational equivalence. In *Proceedings of the 15th ACM Conference on Computer and Communications Security (CCS'08)*, pages 109–118, Alexandria, Virginia, USA, Oct. 2008. ACM Press.
- [87] H. Comon-Lundh, V. Cortier, and E. Zlinescu. Deciding security properties of cryptographic protocols. application to key cycles. *Transaction on Computational Logic*, 11(2), 2010.
- [88] H. Comon-Lundh and V. Shmatikov. Intruder deductions, constraint solving and insecurity decision in presence of exclusive or. In P. Kolaitis, editor, *Eighteenth Annual IEEE Symposium on Logic in Computer Science*, Ottawa, Canada, June 2003. IEEE Computer Society.
- [89] R. Corin, J. Doumen, and S. Etalle. Analysing password protocol security against off-line dictionary attacks. *ENTCS*, 121:47–63, 2005.
- [90] R. Corin, S. Malladi, J. Alves-Foss, and S. Etalle. Guess what? here is a new tool that finds some new guessing attacks. In R. Gorrieri, editor, *Workshop on Issues in the Theory of Security (WITS'03)*, Warsaw, Poland, Apr. 2003.
- [91] J.-S. Coron. On the Exact Security of Full-Domain-Hash. In *Crypto '00*, LNCS 1880, pages 229–235. Springer-Verlag, Berlin, 2000.

- [92] V. Cortier, J. Delaitre, and S. Delaune. Safely composing security protocols. In *Proc. 27th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'07)*, volume 4855 of *LNCS*, pages 352–363. Springer, 2007.
- [93] V. Cortier and S. Delaune. Deciding knowledge in security protocols for monoidal equational theories. In N. Dershowitz and A. Voronkov, editors, *Proceedings of the 14th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR'07)*, volume 4790 of *Lecture Notes in Artificial Intelligence*, pages 196–210, Yerevan, Armenia, Oct. 2007. Springer.
- [94] V. Cortier and S. Delaune. A method for proving observational equivalence. In *Proceedings of the 22nd IEEE Computer Security Foundations Symposium (CSF'09)*, pages 266–276, Port Jefferson, NY, USA, July 2009. IEEE Computer Society Press.
- [95] V. Cortier and S. Delaune. Safely composing security protocols. *Formal Methods in System Design*, 34(1):1–36, Feb. 2009.
- [96] V. Cortier, S. Delaune, and P. Lafourcade. A survey of algebraic properties used in cryptographic protocols. *Journal of Computer Security*, 14(1):1–43, 2006.
- [97] V. Cortier, S. Kremer, and B. Warinschi. A survey of symbolic methods in computational analysis of cryptographic systems. Research Report RR-6912, INRIA, Apr. 2009.
- [98] R. Cramer and V. Shoup. A Practical Public Key Cryptosystem Provably Secure against Adaptive Chosen Ciphertext Attack. In *Crypto '98*, LNCS 1462, pages 13–25. Springer-Verlag, Berlin, 1998.
- [99] A. Datta, A. Derek, J. C. Mitchell, and A. Roy. Protocol composition logic (PCL). *Electr. Notes Theoretical Computer Science*, 172:311–358, 2007.
- [100] H. de Nivelle. *Ordering Refinements of Resolution*. PhD thesis, Technische Universiteit Delft, Oct. 1995.
- [101] S. Delaune and F. Jacquemard. A theory of dictionary attacks and its complexity. In *17th IEEE Computer Security Foundations Workshop*, pages 2–15, Pacific Grove, CA, June 2004. IEEE.
- [102] S. Delaune and F. Jacquemard. Decision procedures for the security of protocols with probabilistic encryption against offline dictionary attacks. *Journal of Automated Reasoning*, 36(1-2):85–124, Jan. 2006.
- [103] S. Delaune, S. Kremer, and M. D. Ryan. Composition of password-based protocols. In *Proceedings of the 21st IEEE Computer Security Foundations Symposium (CSF'08)*, pages 239–251, Pittsburgh, PA, USA, June 2008. IEEE Computer Society Press.
- [104] S. Delaune, S. Kremer, and M. D. Ryan. Symbolic bisimulation for the applied pi calculus. *Journal of Computer Security*, 2009. To appear.
- [105] S. Delaune, S. Kremer, and M. D. Ryan. Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security*, 17(4):435–487, July 2009.

- [106] G. Denker, J. Meseguer, and C. Talcott. Protocol specification and analysis in Maude. In N. Heintze and J. Wing, editors, *Workshop on Formal Methods and Security Protocols*, Indianapolis, Indiana, 25 June 1998.
- [107] D. E. Denning and G. M. Sacco. Timestamps in key distribution protocols. *Commun. ACM*, 24(8):533–536, Aug. 1981.
- [108] W. Diffie and M. E. Hellman. New Directions in Cryptography. *IEEE Transactions on Information Theory*, IT-22(6):644–654, November 1976.
- [109] D. Dolev, C. Dwork, and M. Naor. Non-Malleable Cryptography. *SIAM Journal on Computing*, 30(2):391–437, 2000.
- [110] D. Dolev and A. C. Yao. On the security of public key protocols. *IEEE Transactions on Information Theory*, IT-29(12):198–208, Mar. 1983.
- [111] P. H. Drielsma, S. Mödersheim, and L. Viganò. A formalization of off-line guessing for security protocol analysis. In F. Baader and A. Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning: 11th International Conference, LPAR 2004*, volume 3452 of *Lecture Notes on Computer Science*, pages 363–379, Montevideo, Uruguay, Mar. 2005. Springer.
- [112] N. A. Durgin, P. D. Lincoln, J. C. Mitchell, and A. Scedrov. Undecidability of bounded security protocols. In *Workshop on Formal Methods and Security Protocols (FMSP'99)*, Trento, Italy, 5 July 1999.
- [113] T. El Gamal. A Public Key Cryptosystem and a Signature Scheme Based on Discrete Logarithms. *IEEE Transactions on Information Theory*, IT-31(4):469–472, July 1985.
- [114] F. J. T. Fábrega, J. C. Herzog, and J. D. Guttman. Strand spaces: Proving security protocols correct. *Journal of Computer Security*, 7(2/3):191–230, 1999.
- [115] A. Fiat and A. Shamir. How to Prove Yourself: Practical Solutions of Identification and Signature Problems. In *Crypto '86*, LNCS 263, pages 186–194. Springer-Verlag, Berlin, 1987.
- [116] C. Fournet and M. Abadi. Hiding names: Private authentication in the applied pi calculus. In *Proc. International Symposium on Software Security (ISSS'02)*, volume 2609 of *Lecture Notes in Computer Science*, pages 317–338. Springer, 2003.
- [117] E. Fujisaki, T. Okamoto, D. Pointcheval, and J. Stern. RSA–OAEP is Secure under the RSA Assumption. In *Crypto '01*, LNCS 2139, pages 260–274. Springer-Verlag, Berlin, 2001. Also appeared as *RSA–OAEP is Still Alive* in the Cryptology ePrint Archive 2000/061. November 2000.
Available from <http://eprint.iacr.org/>.
- [118] E. Fujisaki, T. Okamoto, D. Pointcheval, and J. Stern. RSA–OAEP is Secure under the RSA Assumption. *Journal of Cryptology*, 17(2):81–104, 2004.
- [119] J. C. Godskesen. Formal verification of the ARAN protocol using the applied pi-calculus. In *Proceedings of the Sixth International IFIP WG 1.7 Workshop on Issues in the Theory of Security (WITS'06)*, pages 99–113, Vienna, Austria, Mar. 2006.

- [120] O. Goldreich, S. Goldwasser, and S. Micali. How to Construct Random Functions. *Journal of the ACM*, 33(4):792–807, 1986.
- [121] S. Goldwasser and S. Micali. Probabilistic Encryption. *Journal of Computer and System Sciences*, 28:270–299, 1984.
- [122] S. Goldwasser, S. Micali, and C. Rackoff. The Knowledge Complexity of Interactive Proof Systems. In *Proc. of the 17th STOC*, pages 291–304. ACM Press, New York, 1985.
- [123] S. Goldwasser, S. Micali, and R. Rivest. A “Paradoxical” Solution to the Signature Problem. In *Proc. of the 25th FOCS*, pages 441–448. IEEE, New York, 1984.
- [124] S. Goldwasser, S. Micali, and R. Rivest. A Digital Signature Scheme Secure Against Adaptative Chosen-Message Attacks. *SIAM Journal of Computing*, 17(2):281–308, April 1988.
- [125] L. Gong. Verifiable-text attacks in cryptographic protocols. In *INFOCOM ’90, The Conference on Computer Communications*, pages 686–693, San Francisco, CA, June 1990. IEEE.
- [126] L. Gong, T. M. A. Lomas, R. M. Needham, and J. H. Saltzer. Protecting poorly chosen secrets from guessing attacks. *IEEE Journal on Selected Areas in Communications*, 11(5):648–656, June 1993.
- [127] L. Gong and P. Syverson. Fail-stop protocols: An approach to designing secure protocols. In *Proc. 5th International Working Conference on Dependable Computing for Critical Applications*, pages 44–55, 1995.
- [128] A. Gordon and A. Jeffrey. Authenticity by typing for security protocols. *Journal of Computer Security*, 11(4):451–521, 2003.
- [129] J. Goubault-Larrecq. Deciding \mathcal{H}_1 by resolution. *Information Processing Letters*, 95(3):401–408, Aug. 2005.
- [130] J. Goubault-Larrecq and F. Parrennes. Cryptographic protocol analysis on real C code. In R. Cousot, editor, *Proceedings of the 6th International Conference on Verification, Model Checking and Abstract Interpretation (VMCAI’05)*, volume 3385 of *Lecture Notes on Computer Science*, pages 363–379, Paris, France, Jan. 2005. Springer.
- [131] J. D. Guttman and F. J. Thayer. Protocol independence through disjoint encryption. In *Proc. 13th Computer Security Foundations Workshop (CSFW’00)*, pages 24–34. IEEE Comp. Soc. Press, 2000.
- [132] C. Hall, I. Goldberg, and B. Schneier. Reaction Attacks Against Several Public-Key Cryptosystems. In *Proc. of ICICS ’99*, LNCS, pages 2–12. Springer-Verlag, 1999.
- [133] J. Håstad. Solving Simultaneous Modular Equations of Low Degree. *SIAM Journal of Computing*, 17:336–341, 1988.
- [134] J. Heather, G. Lowe, and S. Schneider. How to prevent type flaw attacks on security protocols. In *Proc. 13th Computer Security Foundations Workshop (CSFW’01)*, pages 255–268. IEEE Comp. Soc. Press, 2000.

- [135] A. Joux and R. Lercier. Improvements to the general Number Field Sieve for discrete logarithms in prime fields. *Mathematics of Computation*, 2000. to appear.
- [136] M. Joye, J. J. Quisquater, and M. Yung. On the Power of Misbehaving Adversaries and Security Analysis of the Original EPOC. In *CT – RSA ’01*, LNCS 2020, pages 208–222. Springer-Verlag, Berlin, 2001.
- [137] M. Kallahalla, E. Riedel, R. Swaminathan, Q. Wang, and K. Fu. Plutus: Scalable secure file sharing on untrusted storage. In *2nd Conference on File and Storage Technologies (FAST’03)*, pages 29–42, San Francisco, CA, Apr. 2003. Usenix.
- [138] KCDSA Task Force Team. The Korean Certificate-based Digital Signature Algorithm. Submission to IEEE P1363a. August 1998.
Available from <http://grouper.ieee.org/groups/1363/>.
- [139] J. Kelsey, B. Schneier, and D. Wagner. Protocol interactions and the chosen protocol attack. In *Proc. 5th International Workshop on Security Protocols*, volume 1361 of *LNCS*, pages 91–104. Springer, 1997.
- [140] H. Khurana and H.-S. Hahm. Certified mailing lists. In *Proceedings of the ACM Symposium on Communication, Information, Computer and Communication Security (ASI-ACCS’06)*, pages 46–58, Taipei, Taiwan, Mar. 2006. ACM.
- [141] T. Kleinjung, K. Aoki, J. Franke, A. Lenstra, E. Thomé, J. Bos, P. Gaudry, A. Kruppa, P. Montgomery, D. A. Osvik, H. te Riele, A. Timofeev, and P. Zimmermann. Factorization of a 768-bit RSA modulus, 2010. Cryptology ePrint Archive 2010/006.
- [142] P. C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Crypto ’96*, LNCS 1109, pages 104–113. Springer-Verlag, Berlin, 1996.
- [143] P. C. Kocher, J. Jaffe, and B. Jun. Differential Power Analysis. In *Crypto ’99*, LNCS 1666, pages 388–397. Springer-Verlag, Berlin, 1999.
- [144] H. Krawczyk. SKEME: A versatile secure key exchange mechanism for internet. In *Internet Society Symposium on Network and Distributed Systems Security*, Feb. 1996. Available at <http://bilbo.isu.edu/sndss/sndss96.html>.
- [145] S. Kremer and L. Mazaré. Adaptive soundness of static equivalence. In *European Symposium on Research in Computer Security (ESORICS’07)*, volume 4734 of *LNCS*, pages 610–625, 2007.
- [146] S. Kremer and M. D. Ryan. Analysis of an electronic voting protocol in the applied pi-calculus. In M. Sagiv, editor, *Programming Languages and Systems — Proceedings of the 14th European Symposium on Programming (ESOP’05)*, volume 3444 of *Lecture Notes in Computer Science*, pages 186–200, Edinburgh, U.K., Apr. 2005. Springer.
- [147] R. Küsters and T. Truderung. Reducing protocol analysis with XOR to the XOR-free case in the Horn theory based approach. In *Proceedings of the 15th ACM conference on Computer and communications security (CCS’08)*, pages 129–138, Alexandria, Virginia, USA, Oct. 2008. ACM.

- [148] R. Küsters and T. Truderung. Using ProVerif to analyze protocols with Diffie-Hellman exponentiation. In *22nd IEEE Computer Security Foundations Symposium (CSF'09)*, pages 157–171, Port Jefferson, New York, USA, July 2009. IEEE.
- [149] A. Lenstra and H. Lenstra. *The Development of the Number Field Sieve*, volume 1554 of *Lecture Notes in Mathematics*. Springer-Verlag, 1993.
- [150] A. Lenstra and E. Verheul. Selecting Cryptographic Key Sizes. In *PKC '00*, LNCS 1751, pages 446–465. Springer-Verlag, Berlin, 2000.
- [151] H. Lenstra. On the Chor-Rivest Knapsack Cryptosystem. *Journal of Cryptology*, 3:149–155, 1991.
- [152] G. Lowe. An attack on the Needham-Schroeder public-key authentication protocol. *Information Processing Letters*, 56(3):131–133, 1996.
- [153] G. Lowe. Breaking and fixing the Needham-Schroeder public-key protocol using FDR. In Margaria and Steffen, editors, *Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, volume 1055 of *Lecture Notes in Computer Science*, pages 147–166, 1996.
- [154] G. Lowe. A hierarchy of authentication specification. In *10th Computer Security Foundations Workshop (CSFW '97), June 10-12, 1997, Rockport, Massachusetts, USA*, pages 31–44. IEEE Computer Society, 1997.
- [155] G. Lowe. Towards a completeness result for model checking of security protocols. *Journal of Computer Security*, 7(1), 1999.
- [156] G. Lowe. Analysing protocols subject to guessing attacks. *Journal of Computer Security*, 12(1):83–98, 2004.
- [157] K. D. Lux, M. J. May, N. L. Bhattad, and C. A. Gunter. WSEmail: Secure internet messaging based on web services. In *International Conference on Web Services (ICWS'05)*, pages 75–82, Orlando, Florida, July 2005. IEEE Computer Society.
- [158] C. Lynch. Oriented equational logic programming is complete. *Journal of Symbolic Computation*, 21(1):23–45, 1997.
- [159] J. Manger. A Chosen Ciphertext Attack on RSA Optimal Asymmetric Encryption Padding (OAEP) as Standardized in PKCS #1. In *Crypto '01*, LNCS 2139, pages 230–238. Springer-Verlag, Berlin, 2001.
- [160] D. McAllester. Automatic recognition of tractability in inference relations. *J. ACM*, 40(2), 1993.
- [161] C. Meadows and P. Narendran. A unification algorithm for the group Diffie-Hellman protocol. In *Workshop on Issues in the Theory of Security (WITS'02)*, Portland, Oregon, Jan. 2002.
- [162] J. Millen and V. Shmatikov. Constraint solving for bounded-process cryptographic protocol analysis. In *Proc. 8th ACM Conference on Computer and Communications Security*, 2001.

- [163] G. Miller. Riemann's Hypothesis and Tests for Primality. *Journal of Computer and System Sciences*, 13:300–317, 1976.
- [164] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, i & ii. *Inf. Comput.*, 100(1):1–77, 1992.
- [165] D. M'Raihi, D. Naccache, D. Pointcheval, and S. Vaudenay. Computational Alternatives to Random Number Generators. In *Fifth Annual Workshop on Selected Areas in Cryptography (SAC '98)*, LNCS 1556, pages 72–80. Springer-Verlag, Berlin, 1998.
- [166] M. Naor and M. Yung. Public-Key Cryptosystems Provably Secure against Chosen Ciphertext Attacks. In *Proc. of the 22nd STOC*, pages 427–437. ACM Press, New York, 1990.
- [167] V. I. Nechaev. Complexity of a Determinate Algorithm for the Discrete Logarithm. *Mathematical Notes*, 55(2):165–172, 1994.
- [168] R. Needham and M. Schroeder. Using encryption for authentication in large networks of computers. *Communications of the ACM*, 21(12):993–999, 1978.
- [169] R. M. Needham and M. D. Schroeder. Authentication revisited. *Operating Systems Review*, 21(1):7, 1987.
- [170] J. B. Nielsen. Separating Random Oracle Proofs from Complexity Theoretic Proofs: The Non-committing Encryption Case. In *Crypto '02*, LNCS 2442, pages 111–126. Springer-Verlag, Berlin, 2002.
- [171] NIST. Digital Signature Standard (DSS). Federal Information Processing Standards Publication 186, November 1994.
- [172] NIST. Secure Hash Standard (SHS). Federal Information Processing Standards Publication 180–1, April 1995.
- [173] K. Ohta and T. Okamoto. On Concrete Security Treatment of Signatures Derived from Identification. In *Crypto '98*, LNCS 1462, pages 354–369. Springer-Verlag, Berlin, 1998.
- [174] T. Okamoto and D. Pointcheval. REACT: Rapid Enhanced-security Asymmetric Cryptosystem Transform. In *CT – RSA '01*, LNCS 2020, pages 159–175. Springer-Verlag, Berlin, 2001.
- [175] T. Okamoto and D. Pointcheval. The Gap-Problems: a New Class of Problems for the Security of Cryptographic Schemes. In *PKC '01*, LNCS 1992. Springer-Verlag, Berlin, 2001.
- [176] D. Otway and O. Rees. Efficient and timely mutual authentication. *Operating Systems Review*, 21(1):8–10, 1987.
- [177] L. C. Paulson. The inductive approach to verifying cryptographic protocols. *Journal of Computer Security*, 6(1–2):85–128, 1998.
- [178] D. Pointcheval and J. Stern. Security Proofs for Signature Schemes. In *Eurocrypt '96*, LNCS 1070, pages 387–398. Springer-Verlag, Berlin, 1996.

- [179] D. Pointcheval and J. Stern. Security Arguments for Digital Signatures and Blind Signatures. *Journal of Cryptology*, 13(3):361–396, 2000.
- [180] D. Pointcheval and S. Vaudenay. On Provable Security for Digital Signature Algorithms. Technical Report LIENS-96-17, LIENS, October 1996.
- [181] J. M. Pollard. Monte Carlo Methods for Index Computation (mod p). *Mathematics of Computation*, 32(143):918–924, July 1978.
- [182] C. Rackoff and D. R. Simon. Non-Interactive Zero-Knowledge Proof of Knowledge and Chosen Ciphertext Attack. In *Crypto '91*, LNCS 576, pages 433–444. Springer-Verlag, Berlin, 1992.
- [183] R. Ramanujam and S. P. Suresh. Tagging makes secrecy decidable for unbounded nonces as well. In *Proc. 23rd Conference on Foundations of Software Technology and Theoretical Computer Science (FST&TCS'03)*, volume 2914 of *LNCS*, pages 363–374. Springer, 2003.
- [184] R. Ramanujam and S. P. Suresh. Decidability of context-explicit security protocols. *Journal of Computer Security*, 13(1):135–165, 2005.
- [185] R. Rivest. The MD5 Message-Digest Algorithm. RFC 1321, The Internet Engineering Task Force, April 1992.
- [186] R. Rivest, A. Shamir, and L. Adleman. A Method for Obtaining Digital Signatures and Public Key Cryptosystems. *Communications of the ACM*, 21(2):120–126, February 1978.
- [187] RSA Data Security, Inc. Public Key Cryptography Standards – PKCS. Available from <http://www.rsa.com/rsalabs/pubs/PKCS/>.
- [188] M. Rusinowitch and M. Turuani. Protocol insecurity with finite number of sessions is np-complete. In *Proc. 14th IEEE Computer Security Foundations Workshop*, Cape Breton, Nova Scotia, June 2001.
- [189] M. Ryan and B. Smyth. Applied pi calculus. In V. Cortier and S. Kremer, editors, *Formal Models and Techniques for Analyzing Security Protocols*. IOS Press, To appear.
- [190] C. P. Schnorr. Efficient Identification and Signatures for Smart Cards. In *Crypto '89*, LNCS 435, pages 235–251. Springer-Verlag, Berlin, 1990.
- [191] C. P. Schnorr. Efficient Signature Generation by Smart Cards. *Journal of Cryptology*, 4(3):161–174, 1991.
- [192] C. P. Schnorr and M. Jakobsson. Security of Signed ElGamal Encryption. In *Asiacrypt '00*, LNCS 1976, pages 458–469. Springer-Verlag, Berlin, 2000.
- [193] D. Shanks. Class Number, a Theory of Factorization, and Genera. In *Proceedings of the Symposium on Pure Mathematics*, volume 20, pages 415–440. AMS, 1971.
- [194] H. Shimizu. On the Improvement of the Håstad Bound. In *1996 IEICE Fall Conference*, Volume A-162, 1996. In Japanese.

- [195] V. Shoup. Lower Bounds for Discrete Logarithms and Related Problems. In *Eurocrypt '97*, LNCS 1233, pages 256–266. Springer-Verlag, Berlin, 1997.
- [196] V. Shoup. OAEP Reconsidered. In *Crypto '01*, LNCS 2139, pages 239–259. Springer-Verlag, Berlin, 2001. Also appeared in the Cryptology ePrint Archive 2000/060. November 2000.
Available from <http://eprint.iacr.org/>.
- [197] V. Shoup. OAEP Reconsidered. *Journal of Cryptology*, 15(4):223–249, September 2002.
- [198] V. Shoup. Sequences of games: a tool for taming complexity in security proofs, 2004. Cryptology ePrint Archive 2004/332.
- [199] J. Stern, D. Pointcheval, J. Malone-Lee, and N. Smart. Flaws in Applying Proof Methodologies to Signature Schemes. In *Crypto '02*, LNCS 2442, pages 93–110. Springer-Verlag, Berlin, 2002.
- [200] Y. Tsiounis and M. Yung. On the Security of El Gamal based Encryption. In *PKC '98*, LNCS. Springer-Verlag, Berlin, 1998.
- [201] S. Vaudenay. Cryptanalysis of the Chor-Rivest Scheme. In *Crypto '98*, LNCS 1462, pages 243–256. Springer-Verlag, Berlin, 1998.
- [202] B. Warinschi. A computational analysis of the Needham-Schröder-(Lowe) protocol. In *Proceedings of the 16th IEEE Computer Security Foundations Workshop (CSFW'03)*. IEEE Computer Society, 2003.
- [203] C. Weidenbach. Towards an automatic analysis of security protocols in first-order logic. In H. Ganzinger, editor, *16th International Conference on Automated Deduction (CADE-16)*, volume 1632 of *Lecture Notes in Artificial Intelligence*, pages 314–328, Trento, Italy, July 1999. Springer.
- [204] T. Y. C. Woo and S. S. Lam. Authentication for distributed systems. *Computer*, 25(1):39–52, Jan. 1992.
- [205] T. Y. C. Woo and S. S. Lam. A semantic model for authentication protocols. In *Proceedings IEEE Symposium on Research in Security and Privacy*, pages 178–194, Oakland, California, May 1993.
- [206] T. Y. C. Woo and S. S. Lam. Authentication for distributed systems. In D. Denning and P. Denning, editors, *Internet Besieged: Countering Cyberspace Scofflaws*, pages 319–355. ACM Press and Addison-Wesley, Oct. 1997.