



**HAL**  
open science

# Formal certification of game-based cryptographic proofs

Santiago Zanella-Béguelin

► **To cite this version:**

Santiago Zanella-Béguelin. Formal certification of game-based cryptographic proofs. Computer Science and Game Theory [cs.GT]. École Nationale Supérieure des Mines de Paris, 2010. English. NNT : 2010ENMP0050 . pastel-00584350

**HAL Id: pastel-00584350**

**<https://pastel.hal.science/pastel-00584350>**

Submitted on 8 Apr 2011

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

École doctorale n°84 :  
Sciences et technologies de l'information et de la communication

**Doctorat ParisTech**

**T H È S E**

pour obtenir le grade de docteur délivré par

**l'École nationale supérieure des mines de Paris**

**Spécialité « Informatique temps-réel, robotique et automatique »**

*présentée et soutenue publiquement par*

**Santiago José ZANELLA BÉGUELIN**

le 9 décembre 2010

**Certification Formelle de Preuves Cryptographiques  
Basées sur les Séquences de Jeux**

**Formal Certification of Game-Based Cryptographic Proofs**

Directeur de thèse : **Gilles BARTHE**

**Jury**

**M. Michael BACKES**, Professeur, Universität des Saarlandes  
**M. Yassine LAKHNECH**, Professeur, Université Joseph Fourier, Verimag  
**Mme. Christine PAULIN-MOHRING**, Professeur, Université Paris-Sud  
**M. David POINTCHEVAL**, Directeur de Recherche, École Normale Supérieure  
**M. Nick BENTON**, Docteur, Microsoft Research  
**M. Cédric FOURNET**, Docteur, Microsoft Research  
**M. Jean-Jacques LÉVY**, Directeur de Recherche, INRIA

Rapporteur  
Rapporteur  
Rapporteur  
Rapporteur  
Examineur  
Examineur  
Examineur

**T  
H  
È  
S  
E**

**MINES ParisTech**

**Centre de Mathématiques Appliquées**

Rue Claude Daunesse B.P. 207, 06904 Sophia Antipolis Cedex, France



## Résumé

Les séquences de jeux sont une méthodologie établie pour structurer les preuves cryptographiques. De telles preuves peuvent être formalisées rigoureusement en regardant les jeux comme des programmes probabilistes et en utilisant des méthodes de vérification de programmes. Cette thèse décrit CertiCrypt, un outil permettant la construction et vérification automatique de preuves basées sur les jeux. CertiCrypt est implémenté dans l'assistant à la preuve Coq, et repose sur de nombreux domaines, en particulier les probabilités, la complexité, l'algèbre, et la sémantique des langages de programmation. CertiCrypt fournit des outils certifiés pour raisonner sur l'équivalence de programmes probabilistes, en particulier une logique de Hoare relationnelle, une théorie équationnelle pour l'équivalence observationnelle, une bibliothèque de transformations de programme, et des techniques propres aux preuves cryptographiques, permettant de raisonner sur les événements. Nous validons l'outil en formalisant les preuves de sécurité de plusieurs exemples emblématiques, notamment le schéma de chiffrement OAEP et le schéma de signature FDH.



## Abstract

The game-based approach is a popular methodology for structuring cryptographic proofs as sequences of games. Game-based proofs can be rigorously formalized by taking a code-centric view of games as probabilistic programs and relying on programming language techniques to justify proof steps. In this dissertation we present CertiCrypt, a framework that enables the machine-checked construction and verification of game-based cryptographic proofs. CertiCrypt is built upon the general-purpose proof assistant Coq, from which it inherits the ability to provide independently verifiable evidence that proofs are correct, and draws on many areas, including probability and complexity theory, algebra, and semantics of programming languages. The framework provides certified tools to reason about the equivalence of probabilistic programs, including a relational Hoare logic, a theory of observational equivalence, verified program transformations, and ad-hoc programming language techniques of particular interest in cryptographic proofs, such as reasoning about failure events. We validate our framework through the formalization of several significant case studies, including proofs of security of the Optimal Asymmetric Encryption Padding scheme against adaptive chosen-ciphertext attacks, and of existential unforgeability of Full-Domain Hash signatures.



## Disclaimer

This dissertation builds on several published works that I co-authored. In chronological order,

- *Formal certification of ElGamal encryption. A gentle introduction to CertiCrypt.*  
With Gilles Barthe, Benjamin Grégoire and Sylvain Héraud.  
In proceedings of the *5th International workshop on Formal Aspects in Security and Trust, FAST 2008*, volume 5491 of *Lecture Notes in Computer Science*, pages 1–19, Berlin, 2009. Springer-Verlag.
- *Formal certification of code-based cryptographic proofs.*  
With Gilles Barthe and Benjamin Grégoire.  
In proceedings of the *36th ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages, POPL 2009*, pages 90–101, New York, 2009. ACM Press.
- *Formally certifying the security of digital signature schemes.*  
With Gilles Barthe, Benjamin Grégoire and Federico Olmedo.  
In proceedings of the *30th IEEE symposium on Security and Privacy, S&P 2009*, pages 237–250, Los Alamitos, California, 2009. IEEE Computer Society.
- *Programming language techniques for cryptographic proofs.*  
With Gilles Barthe and Benjamin Grégoire.  
In proceedings of the *1st International conference on Interactive Theorem Proving, ITP 2010*, volume 6172 of *Lecture Notes in Computer Science*, pages 115–130, Berlin, 2010. Springer-Verlag.
- *A machine-checked formalization of Sigma-protocols.*  
With Gilles Barthe, Benjamin Grégoire, Sylvain Héraud and Daniel Hedin.  
In proceedings of the *23rd IEEE Computer Security Foundations symposium, CSF 2010*, pages 246–260, Los Alamitos, California, 2010. IEEE Computer Society.
- *Beyond Provable Security. Verifiable IND-CCA Security of OAEP.*  
With Gilles Barthe, Benjamin Grégoire and Yassine Lakhnech.  
To appear in *Topics in Cryptology – CT-RSA 2011, The Cryptographers’ Track at the RSA Conference 2011*.

I contributed in the elaboration of all of them, as well as in the development of the supporting machine-checked proofs.

A Coq development accompanying this dissertation is available upon request.





# Contents

<b>1</b>	<b>Introduction</b> . . . . .	1
1.1	The CertiCrypt Framework . . . . .	2
1.2	Organization of the Dissertation . . . . .	3
1.3	A Primer on Formal Proofs . . . . .	4
1.4	Introductory Examples . . . . .	5
1.4.1	The ElGamal Encryption Scheme . . . . .	7
1.4.2	The Hashed ElGamal Encryption Scheme . . . . .	9
<b>2</b>	<b>A Language for Describing Games</b> . . . . .	15
2.1	Mathematical Preliminaries . . . . .	15
2.1.1	The Unit Interval . . . . .	15
2.1.2	The Measure Monad . . . . .	16
2.1.3	Lifting Predicates and Relations to Distributions . . . . .	17
2.2	The pWHILE Language . . . . .	20
2.2.1	Syntax . . . . .	21
2.2.2	Semantics . . . . .	25
2.3	Probabilistic Polynomial-Time Programs . . . . .	28
2.4	Adversaries . . . . .	31
2.5	Making Security Properties Precise . . . . .	32
2.5.1	EF-CMA Security . . . . .	32
2.5.2	IND-CCA2 Security . . . . .	35
<b>3</b>	<b>Reasoning about Games</b> . . . . .	41
3.1	Probabilistic Relational Hoare Logic (pRHL) . . . . .	41
3.1.1	Observational Equivalence . . . . .	46
3.2	Bridging Steps . . . . .	47
3.2.1	Certified Program Transformations . . . . .	47
3.2.2	Algebraic Equivalences . . . . .	49
3.2.3	Inter-procedural Code Motion . . . . .	50
3.3	Reasoning about Failure Events . . . . .	53
3.3.1	A Logic for Bounding the Probability of Events . . . . .	53
3.3.2	Automation . . . . .	55

<b>4</b>	<b>The PRP/PRF Switching Lemma</b> .....	57
4.1	Pseudorandom Functions .....	58
4.2	Pseudorandom Permutations .....	58
4.3	The PRP/PRF Switching Lemma .....	59
4.3.1	A Proof Based on Eager Sampling .....	60
4.3.2	A Proof Based on the Failure Event Lemma .....	61
4.3.3	Comparison of Both Proofs .....	62
4.4	Pseudorandom Permutations as Pseudorandom Functions .....	62
4.5	Discussion .....	63
4.6	Related Work .....	65
<b>5</b>	<b>Unforgeability of Full-Domain Hash Signatures</b> .....	67
5.1	The Original Proof .....	69
5.2	Improved Bound .....	75
5.3	Practical Interpretation .....	81
5.4	Discussion and Related Work .....	83
<b>6</b>	<b>Ciphertext Indistinguishability of OAEP</b> .....	87
6.1	Indistinguishability under Chosen-Plaintext Attacks .....	90
6.2	Indistinguishability under Chosen-Ciphertext Attacks .....	101
6.2.1	Proof Outline .....	101
6.2.2	Notes about the Proved Security Bound .....	108
6.3	About the Security of OAEP in the Standard Model .....	109
<b>7</b>	<b>Machine-Checked Formalization of Zero-Knowledge Protocols</b> .	111
7.1	Sigma-Protocols .....	112
7.1.1	Relation between sHVZK and HVZK .....	114
7.2	Sigma Protocols Based on Special Homomorphisms .....	116
7.2.1	Concrete Instances of Sigma-Phi Protocols .....	120
7.2.2	Composition of Sigma-Phi Protocols .....	121
7.3	Sigma Protocols Based on Claw-Free Permutations .....	122
7.4	Combination of Sigma-Protocols .....	124
7.4.1	AND-Combination .....	125
7.4.2	OR-Combination .....	126
7.5	Related Work .....	130
7.6	Perspectives .....	131
<b>8</b>	<b>Related Work and Conclusion</b> .....	133
8.1	Related Work .....	133
8.2	Conclusion .....	135
	<b>References</b> .....	137
	<b>Index</b> .....	145

# 1

## Introduction

DESIGNING secure cryptographic systems is a notoriously difficult task. Indeed, the history of modern cryptography is fraught with examples of cryptographic systems that had been thought secure for a long time before being broken and with flawed security proofs that stood unchallenged for years. Provable security [Goldwasser and Micali 1984; Stern 2003] is an approach that aims to establish the security of cryptographic systems through a rigorous analysis in the form of a mathematical proof, borrowing techniques from complexity theory. In a typical provable security argument, security is proved by reduction, showing that any attack against the security of the system would lead to an efficient way to solve some computationally hard problem.

Provable security holds the promise of delivering strong guarantees that cryptographic schemes meet their goals and is becoming unavoidable in the design and evaluation of new schemes. Yet provable security *per se* does not provide specific tools for managing the complexity of proofs and as a result several purported security arguments that followed the approach have been shown to be flawed. Consequently, the cryptographic community is increasingly aware of the necessity of developing methodologies that systematize the type of reasoning that pervade cryptographic proofs, and that guarantee that such reasoning is applied correctly.

One prominent method for achieving a high degree of confidence in cryptographic proofs is to cast security as a program verification problem: this is achieved by formulating goals and hypotheses in terms of probabilistic programs, and defining the adversarial model in terms of complexity classes, e.g. probabilistic polynomial-time programs. This code-centric view leads to statements that are unambiguous and amenable to formalization. However, standard methods to verify programs (e.g. in terms of program logics) are ineffective to directly address the kind of verification goals that arise from cryptographic statements. The game-based approach [Bellare and Rogaway 2006; Halevi 2005; Shoup 2004] is an alternative to standard program verification methods that establishes the verification goal through successive program transformations. In a nutshell, a game-based proof is structured

as a sequence of transformations of the form  $G, A \rightarrow^h G', A'$ , where  $G$  and  $G'$  are probabilistic programs,  $A$  and  $A'$  are events, and  $h$  is a monotonic function such that  $\Pr[G : A] \leq h(\Pr[G' : A'])$ . When the security of a scheme is expressed as an inequality of the form  $\Pr[G_0 : A_0] \leq p$ , it can be proved by exhibiting a sequence of transformations

$$G_0, A_0 \rightarrow^{h_1} G_1, A_1 \rightarrow \cdots \rightarrow^{h_n} G_n, A_n$$

and proving that  $h_1 \circ \cdots \circ h_n(\Pr[G_n : A_n]) \leq p$ . Reductionist arguments can be naturally formulated in this manner by exhibiting a sequence of games where  $\Pr[G_n : A_n]$  encodes the probability of success of some efficient algorithm in solving a problem assumed to be hard. Under this code-centric view of games, game transformations become program transformations and can be justified rigorously by semantic means; in particular, many transformations can be viewed as common program optimizations.

## 1.1 The CertiCrypt Framework

Whereas Bellare and Rogaway [2006] already observed that code-based proofs could be more easily amenable to machine-checking, Halevi [2005] argued that formal verification techniques should be used to improve trust in cryptographic proofs, and set up a program for building a tool that could be used by the cryptographic community to mechanize their proofs. We take a first step towards Halevi's ambitious program by presenting CertiCrypt [Barthe et al. 2009c], a fully machine-checked framework for constructing and verifying game-based cryptographic proofs. CertiCrypt builds on top of the Coq proof assistant [The Coq development team 2009] a broad set of reasoning principles used by cryptographers, drawing on program verification, algebraic reasoning, and probability and complexity theory. The most notable features of the framework are:

*Faithful and rigorous encoding of games.* In order to be readily accessible to cryptographers, we adopt a formalism that is commonly used to describe games. Concretely, the lowest layer of CertiCrypt is an imperative programming language with probabilistic assignments, structured datatypes, and procedure calls. We formalize the syntax and semantics of programs; the latter uses the measure monad of Audebaud and Paulin-Mohring [2009]. (For the connoisseur, we provide a deep and dependently-typed embedding of the syntax; thanks to dependent types, the typeability of programs is obtained for free.) The semantics is instrumented to calculate the cost of running programs; this offers the means to define complexity classes, and in particular to define formally the notion of efficient (probabilistic polynomial-time) adversary. We provide in addition a precise formalization of the adversarial model that captures many assumptions left informal in proofs, notably including policies on memory access.

*Exact security.* Many security proofs only show that the advantage of any efficient adversary against the security of a cryptographic system is asymptotically

negligible with respect to a security parameter (which typically determines the length of keys or messages). However, the cryptographic community is increasingly focusing on exact security, a more useful goal since it gives hints as to how to choose system parameters in practice to satisfy a security requirement. The goal of exact security is to provide a concrete upper bound for the advantage of an adversary executing in a given amount of time. This is in general done by reduction, constructing an algorithm that solves a problem believed to be hard and giving a lower bound for its success probability and an upper bound for its execution time in terms of the advantage and execution time of the original adversary. We focus on bounding the success probability (and only provide automation to bound the execution time asymptotically) since it is arguably where lies most of the difficulty of a cryptographic proof.

*Full and independently verifiable proofs.* We adopt a formal semanticist perspective and go beyond Halevi’s vision in two respects. First, we provide a unified framework to carry out full proofs; all intermediate steps of reasoning can be justified formally, including complex side conditions that justify the correctness of transformations (about probabilities, algebra, complexity, etc.). Second, one notable feature of Coq, and thus CertiCrypt, is to deliver independently verifiable proofs, an important motivation behind the game-based approach. More concretely, every proof yields a proof object that can be checked automatically by a (small and trustworthy) proof checking engine. In order to trust a cryptographic proof, one only needs to check its statement and not its details.

*Powerful and automated reasoning methods.* We formalize a relational Hoare logic and a theory of observational equivalence, and use them as stepping stones to support the main tools of code-based reasoning. In particular, we prove that many transformations used in code-based proofs, including common optimizations, are semantics-preserving. In addition, we mechanize reasoning patterns used ubiquitously in cryptographic proofs, such as reasoning about failure events (the so-called fundamental lemma of game-playing), and a logic for inter-procedural code-motion (used to justify the eager/lazy sampling of random values).

## 1.2 Organization of the Dissertation

The purpose of this dissertation is to provide a high-level description of the CertiCrypt framework, overview the case studies that have been formalized, and stir further interest in machine-checked cryptographic proofs. The rest of the dissertation is organized as follows:

- In the rest of this chapter we briefly discuss the motivation behind formal proofs and the features of modern proof assistants. We then present two introductory examples of game-based proofs, namely the semantic security of the ElGamal and Hashed ElGamal encryption schemes. These examples, although simple,

nically illustrate the kind of verification problem that we study and the techniques that we use to mechanize the construction of proofs.

- In Chapter 2 we overview the mathematical background behind our formalization and present the probabilistic language that we use to describe games and its semantics. We also discuss our model of adversaries and the notions of complexity and termination of programs in a probabilistic setting;
- In Chapter 3 we present the probabilistic relational Hoare logic that constitutes the core of the framework, and describe the formulation and mechanization of game transformations. We describe as well two ubiquitous reasoning patterns of cryptographic proofs and specializations of the relational Hoare logic that can be used to automate them;
- In Chapter 4 we give a detailed description of two different formalizations of the PRP/PRF switching lemma, an important result in cryptography that admits an elegant proof using games;
- In Chapter 5 we describe two different machine-checked proofs of the existential unforgeability against chosen-message attacks of the Full-Domain Hash digital signature scheme. We compare the resulting security bounds and discuss the practical importance of tight reductions and the role that exact security plays in choosing adequate parameters when instantiating schemes;
- In Chapter 6 we describe in some detail a machine-checked proof of the semantic security of the Optimal Asymmetric Encryption Padding scheme against chosen-plaintext attacks, and we report on a significantly more challenging proof of the security of the same scheme under adaptive chosen-ciphertext attacks;
- In Chapter 7 we overview a machine-checked theory of a large class of zero-knowledge protocols. We illustrate how to use this theory to obtain short proofs of several well-known zero-knowledge protocols from the literature;
- We conclude in Chapter 8 with a survey of related work in the area, a discussion of the lessons we learned while building CertiCrypt and perspectives to improve automation and further this line of research.

### 1.3 A Primer on Formal Proofs

Proof assistants are programs designed to support interactive construction and automatic verification of mathematical statements (understood in a broad sense). Initially developed by logicians to experiment with the expressive power of their foundational formalisms, proof assistants are now emerging as a mature technology that can be used effectively for verifying intricate mathematical proofs, such as the Four Color theorem [Gonthier 2008] or the Kepler conjecture [Hales 2008; Hales et al. 2010], or complex software systems, such as operating systems [Klein et al. 2009], virtual machines [Klein and Nipkow 2006] and optimizing compilers [Leroy 2006]. In the realm of cryptography, proof assistants have been used to formally verify secrecy and authenticity properties of protocols [Paulson 1998].

Proof assistants rely on expressive specification languages that allow formalizing arbitrary mathematical notions, and that provide a formal representation of proofs

as proof objects. Their architecture is organized into two layers: a kernel, and a proof engine.

- The kernel is the cornerstone for correctness. Its central component is a checker for verifying the consistency of formal theories, including definitions and proofs. In particular, the checker guarantees that definitions and proofs are well-typed, that there are no missing cases or undefined notions in definitions, and that all proofs are built from valid elementary logical steps and make a correct use of assumptions.
- In contrast, the proof engine assists proof construction. The proof engine embraces a variety of tools. The primary tools are a set of pre-defined tactics; a language for writing user-defined tactics is usually provided. Tactics allow to reduce a proof goal to simpler ones. When invoked on a proof goal  $Q$ , a tactic will compute a new set of goals  $P_1 \dots P_n$ , and a proof that  $P_1 \wedge \dots \wedge P_n \implies Q$ . At the end of each demonstration, the proof engine outputs a proof object.

Proof objects are independently checked by the kernel. Therefore, the proof engine needs not be trusted, and the validity of a formal proof—beyond the accuracy of the statement itself—only depends on the correctness of the kernel. Pleasingly, kernels are extremely reliable programs with restricted functionalities and solid logical foundations.

As with any other mathematical activity, formal proofs strive for elegance and conciseness. In our experience, they also provide a natural setting for improving proofs—in the case of cryptography, improvement can be measured by comparing exact security bounds. Yet, what matters most about a formal proof is that it provides a nearly absolute degree of guarantee, without requiring expensive human verification.

## 1.4 Introductory Examples

This section illustrates the principles of the CertiCrypt framework on two elementary examples of game-based proofs: the semantic security of ElGamal encryption under the Decision Diffie-Hellman assumption, and the semantic security of Hashed ElGamal encryption in the Random Oracle Model under the Computational Diffie-Hellman assumption. The language used to represent games will be formally introduced in the next chapter; an intuitive understanding should suffice to grasp the meaning of the games appearing here. We begin with some background on encryption schemes and their security.

**Definition 1.1 (Asymmetric encryption scheme).** *An asymmetric encryption scheme is composed of a triple of algorithms:*

Key generation: *Given a security parameter  $\eta$ , the key generation algorithm  $\mathcal{KG}(\eta)$  returns a public/secret key pair  $(pk, sk)$ ;*



- Encryption: Given a public key  $pk$  and a plaintext  $m$ , the encryption algorithm  $\mathcal{E}(pk, m)$  computes a ciphertext corresponding to the encryption of  $m$  under  $pk$ ;
- Decryption: Given a secret key  $sk$  and a ciphertext  $c$ , the decryption algorithm  $\mathcal{D}(sk, c)$  returns either the plaintext corresponding to the decryption of  $c$ , if it is a valid ciphertext, or a distinguished value  $\perp$  otherwise.

Key generation and encryption may be probabilistic, while decryption is deterministic. We require that decryption undo encryption: for every pair of keys  $(pk, sk)$  that can be output by the key generation algorithm, and every plaintext  $m$ , it must be the case that  $\mathcal{D}(sk, \mathcal{E}(pk, m)) = m$ .

An asymmetric encryption scheme is said to be semantically secure if it is unfeasible to gain significant information about a plaintext given only a corresponding ciphertext and the public key. Goldwasser and Micali [1984] showed that semantic security is equivalent to the property of ciphertext indistinguishability under chosen-plaintext attacks (IND-CPA, for short). This property can be formally defined in terms of a game played between a challenger and an adversary  $\mathcal{A}$ , represented as a pair of procedures  $(\mathcal{A}_1, \mathcal{A}_2)$  that may share state:

**Game IND-CPA :**  
 $(pk, sk) \leftarrow \mathcal{KG}(\eta)$ ;  
 $(m_0, m_1) \leftarrow \mathcal{A}_1(pk)$ ;  
 $b \xleftarrow{\$} \{0, 1\}$ ;  
 $c \leftarrow \mathcal{E}(pk, m_b)$ ;  
 $\tilde{b} \leftarrow \mathcal{A}_2(c)$

In this game, the challenger first generates a new key pair and gives the public key  $pk$  to the adversary, who returns two plaintexts  $m_0, m_1$  of its choice. The challenger then tosses a fair coin  $b$  and gives the encryption of  $m_b$  back to the adversary, whose goal is to guess which message has been encrypted.

**Definition 1.2 (IND-CPA security).** *The advantage of an adversary  $\mathcal{A}$  in the above experiment is defined as*

$$\mathbf{Adv}_{\mathcal{A}}^{\text{IND-CPA}} = \left| \Pr \left[ \text{IND-CPA} : b = \tilde{b} \right] - \frac{1}{2} \right|$$

An encryption scheme is said to be IND-CPA secure if the advantage of any efficient adversary is a negligible function of the security parameter  $\eta$ , i.e., the adversary cannot do much better than a blind guess.

**Definition 1.3 (Negligible function).** *A function  $\nu : \mathbb{N} \rightarrow \mathbb{R}$  is said to be negligible if it decreases asymptotically faster than the inverse of any polynomial:*

$$\forall c \in \mathbb{N}. \exists n_c \in \mathbb{N}. \forall n \in \mathbb{N}. n \geq n_c \implies |\nu(n)| \leq n^{-c}$$

Note that in order to satisfy the above definition, an encryption scheme must necessarily be probabilistic, otherwise an adversary could trivially detect to which message corresponds the challenge ciphertext by simply encrypting one of the messages it has chosen and comparing the resulting ciphertext with the challenge ciphertext.

#### 1.4.1 The ElGamal Encryption Scheme

Let  $\{G_\eta\}$  be a family of cyclic prime-order groups indexed by a security parameter  $\eta \in \mathbb{N}$ . For a specific value of the security parameter, which we leave implicit, let  $q$  denote the order of the corresponding group in the family and let  $g$  be a generator. ElGamal encryption is defined by the following triple of algorithms:

$$\begin{aligned} \mathcal{KG}(\eta) &\stackrel{\text{def}}{=} x \xleftarrow{\$} \mathbb{Z}_q; \text{ return } (g^x, x) \\ \mathcal{E}(\alpha, m) &\stackrel{\text{def}}{=} y \xleftarrow{\$} \mathbb{Z}_q; \text{ return } (g^y, \alpha^y \times m) \\ \mathcal{D}(x, (\beta, \zeta)) &\stackrel{\text{def}}{=} \text{ return } (\zeta \times \beta^{-x}) \end{aligned}$$

We prove the IND-CPA security of ElGamal encryption under the assumption that the Decision Diffie-Hellman (DDH) problem is hard. Intuitively, for a family of finite cyclic groups, the DDH problem consists in distinguishing between triples of the form  $(g^x, g^y, g^{xy})$  and triples of the form  $(g^x, g^y, g^z)$ , where the exponents  $x, y, z$  are uniformly sampled from  $\mathbb{Z}_q$ . One characteristic of game-based proofs is to formulate computational assumptions using games; the assumption that the DDH problem is hard is no exception and can be formulated as follows:

**Definition 1.4 (Decision Diffie-Hellman assumption).** *Consider the following games*

$$\begin{aligned} \text{Game } DDH_0 &: x, y \xleftarrow{\$} \mathbb{Z}_q; \quad d \leftarrow \mathcal{B}(g^x, g^y, g^{xy}) \\ \text{Game } DDH_1 &: x, y, z \xleftarrow{\$} \mathbb{Z}_q; \quad d \leftarrow \mathcal{B}(g^x, g^y, g^z) \end{aligned}$$

and define the DDH-advantage of an adversary  $\mathcal{B}$  as follows

$$\text{Adv}_{\mathcal{B}}^{DDH} \stackrel{\text{def}}{=} |\Pr[DDH_0 : d = 1] - \Pr[DDH_1 : d = 1]|$$

We say that the DDH assumption holds for the family of groups  $\{G_\eta\}$  when the advantage of any efficient adversary  $\mathcal{B}$  in the above experiment is a negligible function of the security parameter. Note that the semantics of the games (and in particular the order  $q$  of the group) depends on the security parameter  $\eta$ .

ElGamal is an emblematic example of game-based proofs. The proof of its security, which follows the proof by Shoup [2004], embodies many of the techniques described in the next chapters. The proof is done by reduction and shows that every adversary  $\mathcal{A}$  against the chosen-plaintext security of ElGamal that achieves a given advantage can be used to construct a distinguisher  $\mathcal{B}$  that solves DDH with the same advantage and in roughly the same amount of time. We exhibit a concrete construction of this distinguisher:

**Adversary**  $\mathcal{B}(\alpha, \beta, \gamma)$  :

```

( $m_0, m_1$ )  $\leftarrow$   $\mathcal{A}_1(\alpha)$ ;
 $b \xleftarrow{\$}$   $\{0, 1\}$ ;
 $\tilde{b} \leftarrow \mathcal{A}_2(\beta, \gamma \times m_b)$ ;
return  $b = \tilde{b}$ 

```

We prove that  $\mathbf{Adv}_{\mathcal{B}}^{\text{DDH}} = \mathbf{Adv}_{\mathcal{A}}^{\text{IND-CPA}}$  for any given adversary  $\mathcal{A}$ . To conclude the proof (i.e. to show that the advantage of any efficient adversary  $\mathcal{A}$  is negligible), we show that the reduction is efficient: the adversary  $\mathcal{B}$  executes in probabilistic polynomial-time provided the IND-CPA adversary  $\mathcal{A}$  does—we do not show a concrete bound for the execution time of  $\mathcal{B}$ , although it is evident that it incurs only a constant overhead.

Figure 1.1 gives a high-level view of the reduction: games appear inside white background boxes, whereas gray background boxes contain the actual proof scripts used to prove observational equivalence between consecutive games. A proof script is simply a sequence of tactics, each intermediate tactic transforms the current goal into a simpler one, whereas the last tactic in the script ultimately solves the goal. The tactics that appear in the figure hopefully have self-explanatory names, but are explained cursorily below and in more detail in Chapter 3.

The proof proceeds by constructing an adversary  $\mathcal{B}$  against DDH such that the distribution of  $b = \tilde{b}$  (equivalently,  $d$ ) after running the IND-CPA game for ElGamal is exactly the same as the distribution obtained by running game DDH<sub>0</sub>. In addition, we show that the probability of  $d$  being true in DDH<sub>1</sub> is exactly 1/2 for the same adversary  $\mathcal{B}$ . The remaining gap between DDH<sub>0</sub> and DDH<sub>1</sub> is the DDH-advantage of  $\mathcal{B}$ . The reduction is summarized by the following equations:

$$\left| \Pr [\text{IND-CPA} : b = \tilde{b}] - 1/2 \right| = |\Pr [G_1 : d] - 1/2| \quad (1.1)$$

$$= |\Pr [\text{DDH}_0 : d] - 1/2| \quad (1.2)$$

$$= |\Pr [\text{DDH}_0 : d] - \Pr [G_3 : d]| \quad (1.3)$$

$$= |\Pr [\text{DDH}_0 : d] - \Pr [G_2 : d]| \quad (1.4)$$

$$= |\Pr [\text{DDH}_0 : d] - \Pr [\text{DDH}_1 : d]| \quad (1.5)$$

Equation (1.1) holds because games IND-CPA and  $G_1$  induce the same distribution on  $d$ . We specify this as an observational equivalence judgment as  $\text{IND-CPA} \simeq_{\{d\}} G_1$ , and prove it using certified program transformations and decision procedures. A graphical representation of the sequence of tactics used to prove this judgment is shown in Figure 1.2. We first inline the procedure calls to  $\mathcal{KG}$  and  $\mathcal{E}$  in the IND-CPA game and simplify the resulting games by propagating assignments and eliminating dead code (tactics `ep`, `deadcode`). At this point we are left with two games almost identical, except that  $y$  is sampled later in one game than in the other. The tactic `swap` hoists instructions in one game whenever is possible in order to obtain a maximal common prefix with another game, and allows us to hoist the sampling of  $y$  in the program on the left hand side. We conclude the proof by applying the tactic `eqobs_in` that decides observational equivalence of a program with itself.

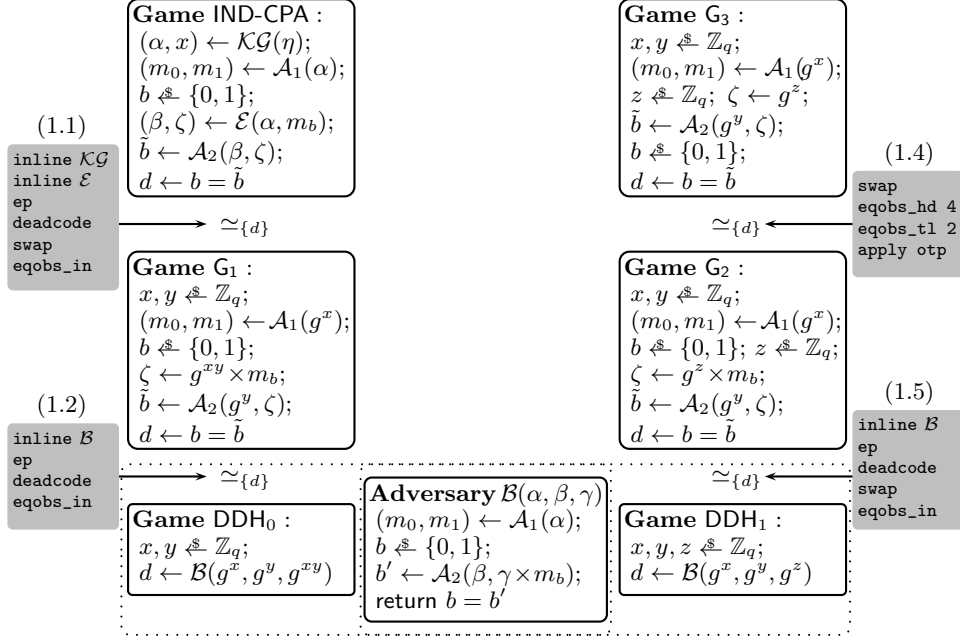



Fig. 1.1. Code-based proof of ElGamal semantic security.

Equations (1.2) and (1.5) are obtained similarly, while (1.3) is established by simple probabilistic reasoning: because in game  $G_3$  the bit  $\tilde{b}$  is independent from  $b$ , the probability of both bits being equal is exactly  $1/2$ . Finally, to prove (1.4) we begin by removing the part the two games have in common with the exception of the instruction  $z \leftarrow \mathbb{Z}_q$  (`swap`, `eqobs_hd 4`, `eqobs_t1 2`) and then apply an algebraic property of cyclic groups that we have proved as a lemma (`otp`): if one applies the group operation to a uniformly distributed element of the group and some other constant element, the result is uniformly distributed—a random element acts as a one-time pad. This allows to prove that  $z \leftarrow \mathbb{Z}_q$ ;  $\zeta \leftarrow g^z \times m_b$  and  $z \leftarrow \mathbb{Z}_q$ ;  $\zeta \leftarrow g^z$  induce the same distribution on  $\zeta$ , and thus remove the dependence of  $\tilde{b}$  on  $b$ .

The proof concludes by applying the DDH assumption to show that the IND-CPA advantage of  $\mathcal{A}$  is negligible. For this, and in view that  $\text{Adv}_{\mathcal{A}}^{\text{IND-CPA}} = \text{Adv}_{\mathcal{B}}^{\text{DDH}}$ , it suffices to prove that the adversary  $\mathcal{B}$  is probabilistic polynomial-time (under the assumption that the procedures  $\mathcal{A}_1$  and  $\mathcal{A}_2$  are so); the proof of this latter fact is entirely automated in CertiCrypt.

### 1.4.2 The Hashed ElGamal Encryption Scheme

Hashed ElGamal is a variant of the ElGamal public-key encryption scheme that does not require plaintexts to be members of the underlying group  $G$ . Instead, plaintexts in Hashed ElGamal are just bitstrings of a certain length  $\ell$  and group elements are



**Fig. 1.2.** Sequence of transformations in the proof of  $\text{IND-CPA} \simeq_{\{d\}} \mathbf{G}_1$ .

mapped into bitstrings using a hash function  $H : G \rightarrow \{0, 1\}^\ell$ . Formally, the scheme is defined by the following triple of algorithms:

$$\begin{aligned}
 \mathcal{KG}(\eta) &\stackrel{\text{def}}{=} x \xleftarrow{\$} \mathbb{Z}_q; \text{ return } (g^x, x) \\
 \mathcal{E}(\alpha, m) &\stackrel{\text{def}}{=} y \xleftarrow{\$} \mathbb{Z}_q; h \leftarrow H(\alpha^y); \text{ return } (g^y, h \oplus m) \\
 \mathcal{D}(x, (\beta, \zeta)) &\stackrel{\text{def}}{=} h \leftarrow H(\beta^x); \text{ return } (\zeta \oplus h)
 \end{aligned}$$

Hashed ElGamal encryption is semantically secure in the random oracle model under the Computational Diffie-Hellman (CDH) assumption on the underlying group family  $\{G_\eta\}$ . This is the assumption that it is hard to compute  $g^{xy}$  given only  $g^x$  and  $g^y$  where  $x$  and  $y$  are uniformly sampled from  $\mathbb{Z}_q$ . Clearly, the DDH assumption implies the CDH assumption, but the converse need not necessarily hold.<sup>1</sup>

**Definition 1.5 (Computational Diffie-Hellman assumption).** *Consider the following game*

<sup>1</sup> Groups where DDH is easy and CDH is believed to be hard are of practical importance in cryptography and are called Diffie-Hellman *gap* groups [Okamoto and Pointcheval 2001a].

**Game CDH:**  $x, y \xleftarrow{\$} \mathbb{Z}_q; \gamma \leftarrow \mathcal{B}(g^x, g^y)$

and define the success probability of  $\mathcal{B}$  against CDH as follows

$$\mathbf{Adv}_{\mathcal{B}}^{\text{CDH}} \stackrel{\text{def}}{=} \Pr[\text{CDH} : \gamma = g^{xy}]$$

We say that the CDH assumption holds for the family of groups  $\{G_\eta\}$  when the success probability of any probabilistic polynomial-time adversary  $\mathcal{B}$  is a negligible function of the security parameter.

We show that any adversary  $\mathcal{A}$  against the IND-CPA security of Hashed ElGamal that makes at most  $q_H$  queries to the hash oracle  $H$  can be used to construct an adversary  $\mathcal{B}$  that achieves a success probability of  $q_H^{-1} \mathbf{Adv}_{\mathcal{A}}^{\text{IND-CPA}}$  in solving the CDH problem. The reduction is done in the random oracle model, where hash functions are modeled as truly random functions. We represent random oracles using stateful procedures; queries are answered consistently: if some value is queried twice, the same response is given. For instance, we code the hash function  $H$  as follows:

```
Oracle  $H(\lambda)$  :
  if  $\lambda \notin \text{dom}(\mathbf{L})$  then
     $h \xleftarrow{\$} \{0, 1\}^\ell$ ;
     $\mathbf{L} \leftarrow (\lambda, h) :: \mathbf{L}$ 
  else  $h \leftarrow \mathbf{L}[\lambda]$ 
  return  $h$ 
```


The proof is sketched in Figure 1.3. We follow the convention of typesetting global variables in boldface. The figure shows the sequence of games used to relate the success of the IND-CPA adversary in the original attack game to the success of the CDH adversary  $\mathcal{B}$ ; the definition of the hash oracle is shown alongside each game. As in the proof of the semantic security of ElGamal, we begin by inlining the calls to  $\mathcal{KG}$  and  $\mathcal{E}$  in the IND-CPA game to obtain an observationally equivalent game  $G_1$  such that

$$\Pr[\text{IND-CPA} : b = \tilde{b}] = \Pr[G_1 : b = \tilde{b}] \quad (1.6)$$

We then fix the value  $\hat{h}$  that the hash oracle gives in response to  $g^{xy}$ . This is an instance of the *lazy sampling* transformation: any value that is randomly sampled at some point in a program can be sampled in advance, somewhere earlier in the program. This transformation is automated in CertiCrypt and is described in greater detail in Section 3.2.3. We get

$$\Pr[G_1 : b = \tilde{b}] = \Pr[G_2 : b = \tilde{b}] \quad (1.7)$$

We can then modify the hash oracle so that it does not store in  $\mathbf{L}$  the response given to a  $g^{xy}$  query; this will later let us remove  $\hat{h}$  altogether from the code of the hash oracle. We prove that games  $G_2$  and  $G_3$  are equivalent by considering the following relational invariant:



**Fig. 1.3.** Game-based proof of semantic security of Hashed ElGamal encryption in the Random Oracle Model.

$$\phi_{23} \stackrel{\text{def}}{=} (\mathbf{A} \in \text{dom}(\mathbf{L}) \implies \mathbf{L}[\mathbf{A}] = \hat{\mathbf{h}}\langle 1 \rangle \wedge \forall \lambda. \lambda \neq \mathbf{A}\langle 1 \rangle \implies \mathbf{L}[\lambda]\langle 1 \rangle = \mathbf{L}[\lambda]\langle 2 \rangle)$$

where  $e\langle 1 \rangle$  (resp.  $e\langle 2 \rangle$ ) denotes the value that expression  $e$  takes in the left hand side (resp. right hand side) program. Intuitively, this invariant shows that the association list  $\mathbf{L}$ , which represents the memory of the hash oracle, coincides in both programs, except perhaps on the element  $\mathbf{A}$ , which the list in the program on the left hand side ( $\mathbf{G}_2$ ) necessarily maps to  $\hat{\mathbf{h}}$ . It is easy to prove that the implementations of oracle  $H$  in games  $\mathbf{G}_2$  and  $\mathbf{G}_3$  are semantically equivalent under this invariant and preserve it. Since  $\phi_{23}$  is established just before calling  $\mathcal{A}$  and is preserved throughout the games, we can prove by inlining the call to  $H$  in game  $\mathbf{G}_2$  that

$$\Pr [\mathbf{G}_2 : b = \tilde{b}] = \Pr [\mathbf{G}_3 : b = \tilde{b}] \quad (1.8)$$

We then *undo* the previous modification to revert to the previous implementation of the hash oracle and prove that games  $\mathbf{G}_3$  and  $\mathbf{G}_4$  are observationally equivalent, from which we obtain

$$\Pr [\mathbf{G}_3 : b = \tilde{b}] = \Pr [\mathbf{G}_4 : b = \tilde{b}] \quad (1.9)$$

Let us now introduce a Boolean flag **bad** that is set at program points where the code of  $\mathbf{G}_4$  and  $\mathbf{G}_5$  differ. We argue that the difference in the probability of any event in those games is bounded by the probability of **bad** being set in  $\mathbf{G}_5$ , and therefore

$$\left| \Pr [\mathbf{G}_4 : b = \tilde{b}] - \Pr [\mathbf{G}_5 : b = \tilde{b}] \right| \leq \Pr [\mathbf{G}_5 : \mathbf{bad}] \quad (1.10)$$

This form of reasoning is pervasive in game-based cryptographic proofs and is an instance of the so-called Fundamental Lemma that we discuss in detail in Section 3.3. In addition, we establish that **bad**  $\implies \mathbf{A} \in \text{dom}(\mathbf{L})$  is a post-condition of game  $\mathbf{G}_5$  and thus

$$\Pr [\mathbf{G}_5 : \mathbf{bad}] \leq \Pr [\mathbf{G}_5 : \mathbf{A} \in \text{dom}(\mathbf{L})] \quad (1.11)$$

Since now both branches in the innermost conditional of the hash oracle are equivalent, we coalesce them to recover the original random oracle implementation of  $H$  in  $\mathbf{G}_6$ . We can now use the **swap** tactic to defer the sampling of  $\hat{\mathbf{h}}$  to the point just before computing  $v$ , and substitute

$$v \stackrel{\$}{\leftarrow} \{0, 1\}^\ell; \hat{\mathbf{h}} \leftarrow v \oplus m_b \quad \text{for} \quad \hat{\mathbf{h}} \stackrel{\$}{\leftarrow} \{0, 1\}^\ell; v \leftarrow \hat{\mathbf{h}} \oplus m_b$$

The semantic equivalence of these two program fragments can be proved using the probabilistic relational Hoare logic presented in Section 3.1—a proof is given in Section 3.2.2. Hence,

$$\Pr [\mathbf{G}_5 : b = \tilde{b}] = \Pr [\mathbf{G}_6 : b = \tilde{b}] \quad (1.12)$$

and

$$\Pr [\mathbf{G}_5 : \mathbf{A} \in \text{dom}(\mathbf{L})] = \Pr [\mathbf{G}_6 : \mathbf{A} \in \text{dom}(\mathbf{L})] \quad (1.13)$$

Observe that  $\tilde{b}$  does not depend anymore on  $b$  in  $\mathbf{G}_6$  ( $\hat{\mathbf{h}} \leftarrow v \oplus m_b$  is dead code), so



$$\Pr [\mathbf{G}_6 : b = \tilde{b}] = \frac{1}{2} \quad (1.14)$$

We finally construct an adversary  $\mathcal{B}$  against CDH that interacts with the adversary  $\mathcal{A}$  playing the role of an IND-CPA challenger. It returns a random element sampled from the list of queries that adversary  $\mathcal{A}$  made to the hash oracle. Observe that  $\mathcal{B}$  does not need to know  $x$  or  $y$  because it gets  $g^x$  and  $g^y$  as parameters. If the correct answer  $\mathbf{A} = g^{xy}$  to the CDH challenge appears in the list of queries  $\mathbf{L}$  when the experiment terminates, adversary  $\mathcal{B}$  has probability  $|\mathbf{L}|^{-1}$  of returning it as an answer. Since we know that  $\mathcal{A}$  does not make more than  $q_H$  queries to the hash oracle, we finally have that

$$\Pr [\mathbf{G}_6 : \mathbf{A} \in \text{dom}(\mathbf{L})] = \Pr [\mathbf{G}_6 : g^{xy} \in \text{dom}(\mathbf{L})] \leq q_H \Pr [\text{CDH} : \gamma = g^{xy}] \quad (1.15)$$

To summarize, from (1.6)—(1.15) we obtain

$$\begin{aligned} \left| \Pr [\text{IND-CPA} : b = \tilde{b}] - 1/2 \right| &= \left| \Pr [\mathbf{G}_4 : b = \tilde{b}] - 1/2 \right| \\ &= \left| \Pr [\mathbf{G}_4 : b = \tilde{b}] - \Pr [\mathbf{G}_6 : b = \tilde{b}] \right| \\ &= \left| \Pr [\mathbf{G}_4 : b = \tilde{b}] - \Pr [\mathbf{G}_5 : b = \tilde{b}] \right| \\ &\leq \Pr [\mathbf{G}_5 : \text{bad}] \\ &\leq \Pr [\mathbf{G}_6 : \mathbf{A} \in \text{dom}(\mathbf{L})] \\ &\leq q_H \Pr [\text{CDH} : \gamma = g^{xy}] \end{aligned}$$

For any adversary  $\mathcal{A}$  that executes in polynomial time, we can assume that the bound  $q_H$  on the number of queries is polynomial on the security parameter. Under the CDH assumption, the IND-CPA advantage of adversary  $\mathcal{A}$  must then be negligible. Otherwise, the adversary  $\mathcal{B}$  that we constructed would solve CDH with non-negligible probability, contradicting our computational assumption. To see this, we need to verify that adversary  $\mathcal{B}$  runs in probabilistic polynomial time, but this is the case because procedures  $\mathcal{A}_1, \mathcal{A}_2$  do, and  $\mathcal{B}$  does not perform any additional costly computations. As in the previous example, the proof of this latter fact is completely automated in CertiCrypt.

We note that Hashed ElGamal can also be proved semantically secure in the standard model, but under the stronger DDH assumption. A game-based proof appears in [Barthe et al. 2009a]. The security reduction can be made under the hypothesis that the family of hash functions  $H$  is *entropy smoothing*—such a family of hash functions can be built without additional assumptions using the Leftover Hash Lemma [Håstad et al. 1999].

# 2

## A Language for Describing Games

S

WE have tried so far to be as rigorous as possible in our treatment of cryptographic proofs. We argued that a game-based approach can lead to tidier, more understandable proofs that help eliminate conspicuous errors and make clear the reasoning behind each step in a proof. We moved one step forward, taking a language-based approach and regarding games as programs. But still, our understanding of what a game means remains purely intuitive. In this chapter we will make this intuition precise by defining formally the probabilistic language we use to describe games and its semantics. This semanticist perspective allows a precise specification of the interaction between an adversary and the challenger in a game, and to readily answer questions that often arise in proofs, such as: Which oracles does the adversary have access to? Can the adversary read/write this variable? How many queries the adversary can make to a given oracle? What is the type/length of a value returned by the adversary? Can the adversary repeat a query? Furthermore, the framework enables us to give very precise definitions of fundamental notions such as probabilistic polynomial-time complexity or termination which are of paramount importance in the specification of security definitions and computational hardness assumptions.

### 2.1 Mathematical Preliminaries

#### 2.1.1 The Unit Interval

The starting point of our formalization is the ALEA Coq library, developed by Paulin-Mohring and described in [Audebaud and Paulin-Mohring 2009]. It provides an axiomatization of the unit interval  $[0, 1]$ , with the following operations:

*Addition:*  $(x, y) \mapsto \min(x + y, 1)$ , where  $+$  denotes addition over reals;  
*Inversion:*  $x \mapsto 1 - x$ , where  $-$  denotes subtraction over reals;

*Multiplication:*  $(x, y) \mapsto x \times y$ , where  $\times$  denotes multiplication over reals;

*Division:*  $(x, y \neq 0) \mapsto \min(x/y, 1)$ , where  $/$  denotes division over reals; moreover, if  $y = 0$ , for convenience division is defined to be 0.

Other useful operations can be derived from these basic operations; for instance the absolute value of the difference of two values  $x, y \in [0, 1]$  can be obtained by computing  $(x - y) + (y - x)$  and their maximum by computing  $(x - y) + y$ .

The unit interval can be given an  $\omega$ -complete partial order (cpo) structure. Recall that an  $\omega$ -cpo consists of a partially ordered set such that any monotonic sequence has a least upper bound. The unit interval  $[0, 1]$  can be given the structure of a  $\omega$ -cpo by taking as order the usual  $\leq$  relation and by defining an operator  $\text{sup}$  that computes the least upper bound of a monotonic sequence  $f : \mathbb{N} \rightarrow [0, 1]$  as follows:

$$\text{sup } f = \max_{n \in \mathbb{N}} f(n)$$

More generally, for any complete partially ordered set  $D$ , we use  $\text{sup } f$  to denote the least upper bound of a monotonic sequence  $f : \mathbb{N} \rightarrow D$ . Note that a cpo structure on  $D$  induces a cpo structure in the function space  $A \rightarrow D$  by taking

$$\begin{aligned} f \leq_{A \rightarrow D} g &\stackrel{\text{def}}{=} \forall x : A. f(x) \leq_D g(x) \\ 0_{A \rightarrow D} &\stackrel{\text{def}}{=} \lambda x. 0_D \\ \text{sup}_{A \rightarrow D} f &\stackrel{\text{def}}{=} \lambda x. \text{sup}_D(f(x)) \end{aligned}$$

### 2.1.2 The Measure Monad

Programs are interpreted as functions from initial memories to sub-probability distributions over final memories. To give semantics to most programs used in cryptographic proofs, it would be sufficient to consider sub-distributions with a countable support, which admit a very direct formalization as functions of the form

$$\mu : A \rightarrow [0, 1] \quad \text{such that} \quad \sum_{x \in A} \mu(x) \leq 1$$

However, it is convenient to take a more general approach and represent instead a distribution over a set  $A$  as a probability measure, by giving a function that maps a  $[0, 1]$ -valued random variable (a function in  $A \rightarrow [0, 1]$ ) to its expected value, i.e. the integral of the random variable with respect to the probability measure. This view of distributions eliminates the need of cluttered definitions and proofs involving summations, and allows us to give a continuation-passing style semantics to programs by defining a suitable monadic structure on distributions. Formally, we represent a distribution on  $A$  as a function  $\mu$  of type

$$\mathcal{D}(A) \stackrel{\text{def}}{=} (A \rightarrow [0, 1]) \rightarrow [0, 1]$$

satisfying the following properties:

*Monotonicity:*  $f \leq g \implies \mu f \leq \mu g$ ;

*Compatibility with inverse:*  $\mu (\mathbb{1} - f) \leq 1 - \mu f$ ;  
*Additive linearity:*  $f \leq \mathbb{1} - g \implies \mu (f + g) = \mu f + \mu g$ ;  
*Multiplicative linearity:*  $\mu (k \times f) = k \times \mu f$ ;  
*Continuity:* if  $f : \mathbb{N} \rightarrow (A \rightarrow [0, 1])$  is monotonic and for all  $n \in \mathbb{N}$   $f(n)$  is monotonic, then  $\mu (\sup f) \leq \sup (\mu \circ f)$

We do not restrict our attention only to distributions with probability mass of 1, but we consider as well sub-probability distributions, that may have a total mass strictly less than 1. As we will see, this is key to give semantics to *non-terminating* programs (i.e. programs that do not terminate with probability 1).

Distributions can be interpreted as a monad whose unit and bind operators are defined as follows:

$$\begin{aligned}
 \text{unit} : A &\rightarrow \mathcal{D}(A) && \stackrel{\text{def}}{=} \lambda x. \lambda f. f \ x \\
 \text{bind} : \mathcal{D}(A) &\rightarrow (A \rightarrow \mathcal{D}(B)) \rightarrow \mathcal{D}(B) && \stackrel{\text{def}}{=} \lambda \mu. \lambda F. \lambda f. \mu (\lambda x. (F \ x) \ f)
 \end{aligned}$$

These operators satisfy the usual monadic laws

$$\begin{aligned}
 \text{bind} (\text{unit } x) F &= F \ x \\
 \text{bind } \mu \ \text{unit} &= \mu \\
 \text{bind} (\text{bind } \mu \ F) \ G &= \text{bind } \mu \ (\lambda x. \text{bind} (F \ x) \ G)
 \end{aligned}$$

The monad  $\mathcal{D}$  was proposed by Audebaud and Paulin-Mohring [2009] as a variant of the expectation monad used by Ramsey and Pfeffer [2002], and builds on earlier work by Kozen [1981]. It is, in turn, a specialization of the continuation monad  $(A \rightarrow B) \rightarrow B$ , with result type  $B = [0, 1]$ .

### 2.1.3 Lifting Predicates and Relations to Distributions

For a distribution  $\mu : \mathcal{D}(A)$  over a countable set  $A$ , we let  $\text{support}(\mu)$  denote the set of values in  $A$  with positive probability, i.e. its support:

$$\text{support}(\mu) \stackrel{\text{def}}{=} \{x \in A \mid 0 < \mu \mathbb{I}_{\{x\}}\}$$

where  $\mathbb{I}_X$  denotes the indicator function of set  $X$ ,

$$\mathbb{I}_X \stackrel{\text{def}}{=} \begin{cases} 1 & \text{if } x \in X \\ 0 & \text{otherwise} \end{cases}$$

To lift relations to probability distributions we follow the early work of Jonsson et al. [2001] on probabilistic bisimulations.

**Definition 2.1 (Lifting predicates to distributions).** *Let  $\mu$  be a distribution on a set  $A$ , and  $P$  be a predicate on  $A$ . We define the lifting of  $P$  to  $\mu$  as follows:*

$$\text{range } P \ \mu \stackrel{\text{def}}{=} \forall f. (\forall x. P \ x \implies f \ x = 0) \implies \mu \ f = 0$$

This rather contrived definition is necessary because we consider sub-probability distributions whose total measure may be less than 1; equivalently we could have stated it as:

$$\text{range } P \mu \stackrel{\text{def}}{=} \forall f. (\forall x. P x \implies f x = 1) \implies \mu f = \mu \mathbb{1}$$

Note also that due to the way distributions are formalized, the above definition is strictly stronger than the following, more intuitive definition:

$$\text{range } P \mu \stackrel{\text{def}}{=} \forall x \in \text{support}(\mu). 0 < \mu(\mathbb{I}_{\{x\}}) \implies P x$$

This latter definition makes sense only for distributions with countable support, for which it can be proved equivalent to the above definitions.

**Definition 2.2 (Lifting relations to distributions).** *Let  $\mu_1$  be a probability distribution on a set  $A$  and  $\mu_2$  a probability distribution on a set  $B$ . We define the lifting of a relation  $R \subseteq A \times B$  to  $\mu_1$  and  $\mu_2$  as follows:*

$$\mu_1 R^\# \mu_2 \stackrel{\text{def}}{=} \exists \mu : \mathcal{D}(A \times B). \pi_1(\mu) = \mu_1 \wedge \pi_2(\mu) = \mu_2 \wedge \text{range } R \mu \quad (2.1)$$

where  $\text{range } R \mu$  stands for the lifting of  $R$ , seen as a predicate on pairs in  $A \times B$ , to distribution  $\mu$ , and the projections  $\pi_1(\mu)$ ,  $\pi_2(\mu)$  of  $\mu$  are given by

$$\pi_1(\mu) \stackrel{\text{def}}{=} \text{bind } \mu (\text{unit} \circ \text{fst}) \quad \pi_2(\mu) \stackrel{\text{def}}{=} \text{bind } \mu (\text{unit} \circ \text{snd})$$

In contrast to the definition given by Jonsson et al. [2001], the definition above makes sense even when the distributions do not have a countable support. When they do, both definitions coincide; in this case,  $\mu_1 R^\# \mu_2$  amounts to saying that the probability of each element  $a$  in the support of  $\mu_1$  can be divided among the elements related to it in such a way that when summing up over these probabilities for an element  $b \in B$ , one obtains  $\mu_2 \mathbb{I}_{\{b\}}$ .

Let us give an example that conveys a better intuition; suppose one wants to prove  $\mathcal{U}_A R^\# \mathcal{U}_B$ , where  $\mathcal{U}_X$  stands for the uniform probability distribution on a finite set  $X$ . When  $A$  and  $B$  have the same size, proving this is equivalent to exhibiting a bijection  $f : A \rightarrow B$  such that for every  $a \in A$ ,  $R(a, f(a))$  holds. Indeed, using such  $f$  it is easy to build a distribution  $\mu$  on  $A \times B$  that satisfies the condition in (2.1):

$$\mu \stackrel{\text{def}}{=} \text{bind } \mathcal{U}_A (\lambda a. \text{unit } (a, f(a)))$$

This example, as trivial as it may seem, shows that probabilistic reasoning can sometimes be replaced by simpler forms of reasoning. In typical cryptographic proofs, purely probabilistic reasoning is seldom necessary and most *mundane* steps in proofs can be either entirely automated or reduced to verifying simpler conditions, much like in the above example, e.g. showing the existence of a bijection with particular properties.

The way we chose to lift relations over memories to relations over distributions is a generalization to arbitrary relations of the definition of Sabelfeld and Sands [2001]

that applies only to equivalence relations. Indeed, there is a simpler but equivalent (see [Jonsson et al. 2001]) way of lifting an equivalence relation to distributions: if  $R$  is an equivalence relation on  $A$ , then  $\mu_1 R^\# \mu_2$  holds if and only if for all equivalence classes  $[a] \subseteq A$ ,  $\mu_1 \mathbb{I}_{[a]} = \mu_2 \mathbb{I}_{[a]}$ .

Define two functions  $f$  and  $g$  to be equal modulo a relation  $\Phi$  iff

$$f =_{\Phi} g \stackrel{\text{def}}{=} \forall x y. x \Phi y \implies f(x) = g(y)$$

It can be easily shown that the above general definition of lifting satisfies

$$\mu_1 \Phi^\# \mu_2 \wedge f =_{\Phi} g \implies \mu_1 f = \mu_2 g$$

and analogously.

$$\mu_1 \Phi^\# \mu_2 \wedge f \leq_{\Phi} g \implies \mu_1 f \leq \mu_2 g$$

We use these properties to prove rules relating observational equivalence to probability in Section 3.1.

It can be shown that lifting preserves the reflexivity and symmetry of the lifted relation, but proving that it preserves transitivity is not as straightforward. Ideally, one would like to have for probability measures  $\mu_1 : \mathcal{D}(A)$ ,  $\mu_2 : \mathcal{D}(B)$ ,  $\mu_3 : \mathcal{D}(C)$  and relations  $\Psi \subseteq A \times B$ ,  $\Phi \subseteq B \times C$

$$\mu_1 \Psi^\# \mu_2 \wedge \mu_2 \Phi^\# \mu_3 \implies \mu_1 (\Psi \circ \Phi)^\# \mu_3$$

Proving this for arbitrary distributions requires proving Fubini's theorem for probability measures, which allows to compute integrals with respect to a product measure in terms of iterated integrals with respect to the original measures. Since in practice we consider distributions with a countable support, we do not need the full generality of this result, and we prove it under the assumption that the distribution  $\mu_2$  has a countable support, i.e. there exist coefficients  $c_i : [0, 1]$  and points  $b_i : B$  such that, for any  $f$ ,

$$\mu_2 f = \sum_{i=0}^{\infty} c_i f(b_i)$$

**Lemma 2.3.** *Consider  $d_1 : \mathcal{D}(A)$ ,  $d_2 : \mathcal{D}(B)$ ,  $d_3 : \mathcal{D}(C)$  such that  $d_2$  has countable support. Suppose there exist distributions  $\mu_{12} : \mathcal{D}(A \times B)$  and  $\mu_{23} : \mathcal{D}(B \times C)$  that make  $\mu_1 \Psi^\# \mu_2$  and  $\mu_2 \Phi^\# \mu_3$  hold. Then, the following distribution over  $A \times C$  is a witness for the existential in  $\mu_1 (\Psi \circ \Phi)^\# \mu_3$ :*

$$\mu_{13} f \stackrel{\text{def}}{=} \mu_2 \left( \lambda b. \mu_{12} \left( \lambda p. \left( \mathbb{1}_{\text{snd}(p)=b} / \mu_2 \mathbb{I}_{\{b\}} \right) \mu_{23} \left( \lambda q. \left( \mathbb{1}_{\text{fst}(q)=b} / \mu_2 \mathbb{I}_{\{b\}} \right) f(\text{fst } p, \text{snd } q) \right) \right) \right)$$

*Proof.* The difficult part of the proof is to show that the projections of this distribution coincide with  $\mu_1$  and  $\mu_2$ . For this, we use the fact that  $\mu_2$  is discrete to prove that iterative integration with respect to  $\mu_2$  and another measure *commutes*. This is the case because we can write integration with respect to  $\mu_2$  as a summation and

we only consider measures that are continue and linear. For instance, to see that the first projection of  $\mu_{13}$  coincides with  $\mu_1$ :

$$\begin{aligned}
\pi_1(\mu_{13}) f &= \sum_{i=0}^{\infty} c_i \mu_{12} (\lambda p. (\mathbb{1}_{\text{snd}(p)=b_i}/c_i) \mu_{23} (\lambda q. (\mathbb{1}_{\text{fst}(q)=b_i}/c_i) f(\text{fst } p))) \\
&= \mu_{12} \left( \lambda p. \sum_{i=0}^{\infty} \mathbb{1}_{\text{snd}(p)=b_i} \mu_{23} (\lambda q. (\mathbb{1}_{\text{fst}(q)=b_i}/c_i) f(\text{fst } p)) \right) \\
&= \mu_{12} \left( \lambda p. f(\text{fst } p) \sum_{i=0}^{\infty} \mathbb{1}_{\text{snd}(p)=b_i} \right) \\
&= \mu_{12} (\lambda p. f(\text{fst } p)) \\
&= \mu_1 f
\end{aligned}$$

□

## 2.2 The pWHILE Language

We describe games as programs in the pWHILE language, a probabilistic extension of an imperative language with procedure calls. This language can be regarded as a mild generalization of the language proposed by Bellare and Rogaway [2006], in that our language allows while loops whereas theirs only allow bounded for loops. The formalization of pWHILE is carefully crafted to exploit key features of Coq: it uses modules to support an extensible expression language that can be adapted according to the verification goal, dependent types to ensure that programs are well-typed and have a total semantics, and monads to give semantics to probabilistic programs and capture the cost of executing them.

We formalize programs in a deep-embedding style, i.e. the syntax of the language is encoded within the proof assistant. Deep embeddings offer one tremendous advantage over shallow embeddings, in which the language used to represent programs is the same as the underlying language of the proof assistant. Namely, a deep embedding allows to manipulate programs as syntactic objects. This permits to achieve a high level of automation in reasoning about programs through certified tactics that implement syntactic program transformations. Additionally, a deep embedding allows to formalize complexity issues neatly and to reason about programs by induction on the structure of their code.

The semantics of programs is given by an interpretation function that takes a program  $p$ —an element of the type of programs—and an initial state  $s$ , and returns the result of executing  $p$  starting from  $s$ . In a deterministic case, both  $s$  and the result of executing  $p$  starting from  $s$  would be deterministic states, i.e. memories mapping variables to values. In the case of pWHILE programs, the denotation of a program is instead a function mapping an initial state to a (sub)-probability measure over final states. We use the measure monad described in 2.1 to define the denotation of programs.

## 2.2.1 Syntax

Given a set  $\mathcal{V}$  of variable identifiers, a set  $\mathcal{P}$  of procedure identifiers, a set  $\mathcal{E}$  of deterministic expressions, and a set  $\mathcal{DE}$  of *distribution expressions*, the instructions  $\mathcal{I}$  and commands  $\mathcal{C}$  of the language can be defined inductively by the clauses:

$\mathcal{I} ::= \mathcal{V} \leftarrow \mathcal{E}$	deterministic assignment
$\mathcal{V} \xleftarrow{\mathcal{DE}}$	probabilistic assignment
<b>if</b> $\mathcal{E}$ <b>then</b> $\mathcal{C}$ <b>else</b> $\mathcal{C}$	conditional
<b>while</b> $\mathcal{E}$ <b>do</b> $\mathcal{C}$	while loop
$\mathcal{V} \leftarrow \mathcal{P}(\mathcal{E}, \dots, \mathcal{E})$	procedure call
$\mathcal{C} ::= \text{skip}$	nop
$\mathcal{I}; \mathcal{C}$	sequence

The inductive definition of the language suffices to understand the rest of the presentation and the reader may prefer to retain it for further reference. In practice, however, variable and procedure identifiers are annotated with their types, and the syntax of programs is dependently-typed. Thus,  $x \leftarrow e$  is well-formed only if the types of  $x$  and  $e$  coincide, and **if**  $e$  **then**  $c_1$  **else**  $c_2$  is well-formed only if  $e$  is a Boolean expression and  $c_1$  and  $c_2$  are themselves well-formed. An immediate benefit of using dependent types is that the type system of **Coq** ensures for free the well-typedness of expressions and commands.

In the remainder of this section we describe in detail the formalization of the syntax and semantics of the language. Most readers, particularly those not familiar with **Coq**, can skim through this section without hindering the understanding of the rest of the dissertation.

**Background on the Coq proof assistant**

We built our framework on top of **Coq**, a general purpose proof assistant that has been used for over two decades to formalize results in mathematics and computer science [The Coq development team 2009]. **Coq** provides an expressive specification language based on the Calculus of Inductive Constructions, a higher-order dependently-typed  $\lambda$ -calculus in which mathematical notions can be formalized conveniently. The **Coq** logic distinguishes between types, of type **Type**, which represent sets, and propositions, of type **Prop**, which represent formulae: thus,  $a : A$  is interpreted as “ $a$  is an element of type  $A$ ” if  $A$  is a set, and as “ $a$  is a proof of  $A$ ” if  $A$  is a proposition. In the latter case, we say that  $a$  is a proof object. Types can either be introduced by standard constructions, e.g. (generalized) function space and products, or by inductive definitions. Most common inductive types are predefined, including the type  $\mathbb{N}$  of natural numbers, the type  $\mathbb{B}$  of Boolean values, and sum and product types. We will also use the inductively defined types of homogeneous and heterogeneous lists. Homogeneous lists are composed of elements of the same type. The polymorphic inductive type of homogeneous lists is defined as follows:



**Inductive** list  $A$  : Type :=  
 | nil : nil  
 | cons :  $A \rightarrow \text{list } A \rightarrow \text{list } A$

The list constructor `cons` is usually represented using an infix notation as the operator “`::`”. Thus, the list composed of the natural numbers 1, 2 and 3, in that order, has type `list ℕ` and could be represented as `(1 :: 2 :: 3 :: nil)`. Heterogeneous lists are composed of elements whose type may depend on a value. Given a type  $A$  and a type-valued function  $P : A \rightarrow \text{Type}$ , the inductive type of heterogeneous lists is defined as follows:

**Inductive** hlist  $A$  ( $P : A \rightarrow \text{Type}$ ) : list  $A \rightarrow \text{Type}$  :=  
 | dnil : hlist nil  
 | dcons :  $\forall a l, P a \rightarrow \text{hlist } l \rightarrow \text{hlist } (a :: l)$

We will use  $A^*$  to denote the type of  $A$ -lists (i.e. `list  $A$` ), and  $P_l^*$  to denote the type of heterogeneous  $P$ -lists over a list of values  $l$  (i.e. `hlist  $P l$` ).

## Types

We formalize a dependently-typed syntax, and use the underlying type system of `Coq` to ensure for free that expressions and commands are well-formed. In our experience, the typed syntax provides particularly useful feedback when debugging proofs and makes proofs easier by restricting reasoning about programs to reasoning about their meaningful behaviors.

The types and expressions of the language are defined on top of a module that contains the declaration of user-defined types and operators. Formally, the set  $\mathcal{T}$  of types is defined as:

**Inductive**  $\mathcal{T}$  : Type :=  
 | User :  $\mathcal{T}_{\text{user}} \rightarrow \mathcal{T}$   
 | Nat :  $\mathcal{T}$   
 | Bool :  $\mathcal{T}$   
 | List :  $\mathcal{T} \rightarrow \mathcal{T}$   
 | Pair :  $\mathcal{T} \rightarrow \mathcal{T} \rightarrow \mathcal{T}$   
 | Sum :  $\mathcal{T} \rightarrow \mathcal{T} \rightarrow \mathcal{T}$   
 | Option :  $\mathcal{T} \rightarrow \mathcal{T}$

where  $\mathcal{T}_{\text{user}}$  denotes the set of user-defined types. This set can be given different definitions according to the cryptographic system under verification.

The interpretation of types—and of programs in general—depends on a security parameter  $\eta$  (a natural number),

**Definition**  $\text{interp } (\eta : \mathbb{N}) (t : \mathcal{T}) :=$   
 match  $t$  with  
 | User  $ut$   $\Rightarrow$   $\text{interp}_{\text{user}} \eta ut$   
 | Nat  $\Rightarrow$   $\mathbb{N}$   
 | Bool  $\Rightarrow$   $\mathbb{B}$   
 | List  $t$   $\Rightarrow$   $\text{list } (\text{interp } \eta t)$   
 | Pair  $t_1 t_2 \Rightarrow$   $(\text{interp } \eta t_1) \times (\text{interp } \eta t_2)$   
 | Sum  $t_1 t_2 \Rightarrow$   $(\text{interp } \eta t_1) + (\text{interp } \eta t_2)$   
 end

For instance, to introduce a type of bitstrings of a certain length  $\ell(\eta)$  depending on the security parameter, one would define  $\mathcal{T}_{\text{user}}$  as follows,

**Inductive**  $\mathcal{T}_{\text{user}} : \text{Type} :=$   
 | Bitstring :  $\mathcal{T}_{\text{user}}$

and let the interpretation of **Bitstring** for a value  $\eta$  of the security parameter be some representation of the set  $\{0, 1\}^{\ell(\eta)}$  in **Coq** (the type of bitvectors defined in the standard library of **Coq** provides a convenient representation).

### Expressions

Expressions are built from a set of  $\mathcal{T}$ -indexed variable names  $\mathcal{V}$ , using operators from the core language, such as constructors for pairs and lists, and user-defined operators. All operators are declared with typing information, as specified by the functions **targs** and **tres**, that return for each operator the list of types of its arguments, and the type of its result, respectively. The  $\mathcal{T}$ -indexed family  $\mathcal{E}$  of expressions is then defined as:

**Inductive**  $\mathcal{E} : \mathcal{T} \rightarrow \text{Type} :=$   
 | Enat  $:\>$   $\mathbb{N} \rightarrow \mathcal{E}_{\text{Nat}}$   
 | Ebool  $:\>$   $\mathbb{B} \rightarrow \mathcal{E}_{\text{Bool}}$   
 | Evar  $:\>$   $\forall t, \mathcal{V}_t \rightarrow \mathcal{E}_t$   
 | Eop  $:$   $\forall op, \mathcal{E}_{(\text{targs } op)}^* \rightarrow \mathcal{E}_{(\text{tres } op)}$   
 | Eforall  $:$   $\forall t, \mathcal{V}_t \rightarrow \mathcal{E}_{\text{Bool}} \rightarrow \mathcal{E}_{(\text{List } t)} \rightarrow \mathcal{E}_{\text{Bool}}$   
 | Eexists  $:$   $\forall t, \mathcal{V}_t \rightarrow \mathcal{E}_{\text{Bool}} \rightarrow \mathcal{E}_{(\text{List } t)} \rightarrow \mathcal{E}_{\text{Bool}}$   
 | Efind  $:$   $\forall t, \mathcal{V}_t \rightarrow \mathcal{E}_{\text{Bool}} \rightarrow \mathcal{E}_{(\text{List } t)} \rightarrow \mathcal{E}_t$

The first three clauses declare constructors as coercions; thanks to this mechanism it is possible to view an element of their domain as an element of their codomain, e.g. a natural number as an expression of type **Nat** and a variable of type  $t$  as an expression of type  $t$ . The fourth clause corresponds to the standard rule for operators; the rule requires that the types of the arguments be compatible with the declaration of the operator, as enforced by the type  $\mathcal{E}_{(\text{targs } op)}^*$  of heterogeneous lists of expressions. In this clause,  $op$  is universally quantified over an inductive type that contains a fixed set of operators for base types and user-defined operators.

The last three clauses introduce operations on lists that are commonly used in cryptographic proofs: they take as parameters a variable  $x$  of type  $t$ , a Boolean valued expression  $e$  that may depend on  $x$ , and an expression  $l$  of type  $\text{List } t$ , and respectively

- (Eforall  $x \ e \ l$ ) checks whether every element  $a$  in  $l$  verifies  $e$  when substituting  $a$  for  $x$ . We note it as  $(\forall x \in l. e)$ ;
- (Eexists  $x \ e \ l$ ) checks whether some element  $a$  in  $l$  verifies  $e$  when substituting  $a$  for  $x$ . We note it as  $(\exists x \in l. e)$ ;
- (Efind  $x \ e \ l$ ) evaluates to the first element  $a$  in the list  $l$  that verifies  $e$  when substituting  $a$  for  $x$ , or to a designated default element of type  $t$  if no such element is found. We usually do not write this operator explicitly, instead we assume that an expression  $(\exists x \in l, e)$  implicitly assigns to the variable  $x$  the value of  $(\text{Efind } x \ e \ l)$ .

It is worth noting that dependent types allow for rich specifications of operators. For instance, one can define a type for bitstrings of fixed length  $\{0, 1\}^k$ , and a concatenation operator that keeps track of bitstring lengths with type

$$\forall m \ n, \{0, 1\}^m \rightarrow \{0, 1\}^n \rightarrow \{0, 1\}^{m+n}$$

In addition to the set of deterministic expressions defined above, to encode random assignments we use a set of type-indexed distribution expressions  $\mathcal{DE}$ . An element of  $\mathcal{DE}_t$  denotes some discrete distribution over values of type  $t$ . The core language includes expressions denoting the uniform distribution on natural intervals of the form  $[0..n]$ , and on Boolean values. Again, the set of distribution expressions of the core language can be extended by the user,

$$\begin{aligned} \text{Inductive } \mathcal{DE} : \mathcal{T} \rightarrow \text{Type} := \\ & | \text{Dnat} : \mathcal{E}_{\text{Nat}} \rightarrow \mathcal{DE}_{\text{Nat}} \\ & | \text{Dbool} : \mathcal{DE}_{\text{Bool}} \\ & | \text{Duser} : \forall t, \mathcal{DE}_{\text{user}} \ t \rightarrow \mathcal{DE}_t \end{aligned}$$

## Programs

Commands are built from a set of procedure names  $\mathcal{P}$  indexed by the type of their arguments and return value. Formally, the sets  $\mathcal{I}$  of instructions and  $\mathcal{C}$  of commands are defined as follows:

$$\begin{aligned} \text{Inductive } \mathcal{I} : \text{Type} := \\ & | \text{Assign} : \forall t, \mathcal{V}_t \rightarrow \mathcal{E}_t \rightarrow \mathcal{I} \\ & | \text{Rand} : \forall t, \mathcal{V}_t \rightarrow \mathcal{DE}_t \rightarrow \mathcal{I} \\ & | \text{Cond} : \mathcal{E}_{\text{Bool}} \rightarrow \mathcal{C} \rightarrow \mathcal{C} \rightarrow \mathcal{I} \\ & | \text{While} : \mathcal{E}_{\text{Bool}} \rightarrow \mathcal{C} \rightarrow \mathcal{I} \\ & | \text{Call} : \forall l \ t, \mathcal{P}_{(l,t)} \rightarrow \mathcal{V}_t \rightarrow \mathcal{E}_l^* \rightarrow \mathcal{I} \\ \text{where } \mathcal{C} := \mathcal{I}^* \end{aligned}$$

For instance,  $b \leftarrow \{0,1\}$  (a shorthand for `Rand b Dbool`) is an instruction that samples a random bit with uniform probability and assigns it to variable  $b$ . Note that the above syntax lacks a construct for sequential composition; instead, we use lists to represent sequences of instructions.

**Definition 2.4 (Program).** *A program is a pair consisting of a command  $c \in \mathcal{C}$  and an environment  $E : \forall l t, \mathcal{P}_{(l,t)} \rightarrow \text{decl}_{(l,t)}$ , which maps procedure identifiers to their declaration. The declaration of a procedure  $p \in \mathcal{P}_{(l,t)}$  consists of its formal parameters, its body, and a return expression,*

$$\mathbf{Record} \text{ decl}_{(l,t)} \stackrel{\text{def}}{=} \{\text{args} : \mathcal{V}_l^*; \text{body} : \mathcal{C}; \text{re} : \mathcal{E}_t\}$$

An environment specifies the type of the parameters and the return expression of procedures, so that procedure calls are always well-typed. In a typical formalization, the environment will map procedures to closed commands, with the exception of adversaries whose code is unknown, and thus modeled by variables of type  $\mathcal{C}$ . This is a standard trick to deal with uninterpreted functions in a deep embedding.

We frequently make no distinction between a game  $G = (c, E)$  and its main command  $c$  when the environment either has no relevance or is clear from the context. In the remainder, we revert to a more natural notation to specify games: we rely on standard notation as in [Barthe et al. 2009a,c]. In particular, we write procedures that might have multiple exit points and use explicit `return` instructions instead of specifying a single return expression.

### 2.2.2 Semantics

The semantics of commands and expressions depends on a natural number representing the security parameter. As we have seen, the interpretation of types and operators may depend on this parameter, but for the sake of readability we omit it most of the time. The denotation of a command is defined relative to an initial memory, mapping variables to values of their respective types. Since variables are partitioned into local and global variables, we will sometimes represent a memory  $m$  as a pair of mappings  $(m.\text{loc}, m.\text{glob})$  for local and global variables, respectively. We let  $\mathcal{M}$  denote the type of memories and  $\emptyset$  denote a mapping associating variables to default values of their corresponding types.

Expressions are deterministic; their semantics is standard and given by a function

$$\llbracket e : \mathcal{E}_t \rrbracket_{\mathcal{E}} : \mathcal{M} \rightarrow \text{interp } t$$

that evaluates an expression in a given memory and returns a value of the right type. The semantics of distribution expressions is given by a function a

$$\llbracket d : \mathcal{D}\mathcal{E}_t \rrbracket_{\mathcal{D}\mathcal{E}} : \mathcal{M} \rightarrow \mathcal{D}(\text{interp } t)$$

that given a distribution expression  $d$  of type  $t$  and a memory  $m$ , returns a measure over values of type  $t$ . For instance, in Section 1.4.2 we have used  $\{0,1\}^\ell$  to denote the uniform distribution on bitstrings of a certain length  $\ell$ ; formally, we have

$$\llbracket \{0,1\}^\ell \rrbracket_{\mathcal{D}\mathcal{E}} m : \mathcal{D}(\{0,1\}^\ell) \stackrel{\text{def}}{=} \lambda f. \sum_{x \in \{0,1\}^\ell} 2^{-\ell} f(x)$$

Observe that distribution expressions are not restricted to constant distributions. Indeed, for any expression of type  $e \in \mathcal{E}_{\text{Nat}}$ , the semantics of the uniform distribution on the natural interval  $[0..e]$  depends on the evaluation of  $e$ ,

$$\llbracket [0..e] \rrbracket_{\mathcal{D}\mathcal{E}} m : \mathcal{D}(\mathbb{N}) \stackrel{\text{def}}{=} \lambda f. \sum_{i=0}^n \frac{1}{n+1} f(i) \quad \text{where } n = \llbracket e \rrbracket_{\mathcal{E}} m$$

Thanks to dependent types, the semantics of expressions and distribution expressions is total. In the following, and whenever there is no confusion, we will drop the subscripts in  $\llbracket \cdot \rrbracket_{\mathcal{E}}$  and  $\llbracket \cdot \rrbracket_{\mathcal{D}\mathcal{E}}$ .

The (small-step) semantics of commands relates an initial deterministic state to a sub-probability distribution over final deterministic states. It uses a frame stack to deal with procedure calls. Formally, a deterministic state is a triple consisting of the current command  $c : \mathcal{C}$ , a memory  $m : \mathcal{M}$ , and a frame stack  $F : \text{frame}^*$ . We let  $\mathcal{S}$  denote the set of deterministic states,

$$\mathcal{S} \stackrel{\text{def}}{=} \mathcal{C} \times \mathcal{M} \times \text{frame}^*$$

One step execution  $\llbracket \cdot \rrbracket_1 : \mathcal{S} \rightarrow \mathcal{D}(\mathcal{S})$  is defined by the rules of Figure 2.1; in the figure, we use  $a \rightsquigarrow b$  as a notation for  $\llbracket a \rrbracket_1 = b$ .

We briefly comment on the transition rules for calling a procedure (5th rule) and returning from a call (2nd rule). Upon a call, a new frame is appended to the stack, containing the destination variable, the return expression of the called procedure, the continuation to the call, and the local memory of the caller. The state resulting from the call contains the body of the called procedure, the global part of the memory, a local memory initialized to map the formal parameters to the value of the actual parameters just before the call, and the updated stack. When returning from a call with a non-empty stack, the top frame is popped, the return expression is evaluated and the resulting value is assigned to the destination variable after previously restoring the local memory of the caller; the continuation taken from the frame becomes the current command. If the stack is empty when returning from a call, the execution of the program terminates and the final state is embedded into the monad using the unit operator.

We then define an  $n$ -step execution function  $\llbracket \cdot \rrbracket^n$  as follows:

$$\begin{aligned} \llbracket s \rrbracket^0 &\stackrel{\text{def}}{=} \text{unit } s \\ \llbracket s \rrbracket^{n+1} &\stackrel{\text{def}}{=} \text{bind } \llbracket s \rrbracket^n \llbracket \cdot \rrbracket_1 \end{aligned}$$

Finally, the denotation of a command  $c$  in an initial memory  $m$  is defined to be the (limit) distribution of reachable final memories:

$$\llbracket c \rrbracket m : \mathcal{D}(\mathcal{M}) \stackrel{\text{def}}{=} \lambda f. \sup \{ \llbracket (c, m, \text{nil}) \rrbracket^n f_{|\text{final}} \mid n \in \mathbb{N} \}$$

where  $f_{|\text{final}}$  is a function that when applied to a state  $(c, m, F)$  equals  $f(m)$  if the state is a final state and 0 otherwise, i.e.

---


$$\begin{aligned}
(\text{skip}, m, \text{nil}) &\rightsquigarrow \text{unit}(\text{skip}, m, \text{nil}) \\
(\text{skip}, m, (x, e, c, l) :: F) &\rightsquigarrow \text{unit}(c, (l, m.\text{glob})\{\llbracket e \rrbracket_{\mathcal{E}} m/x\}, F) \\
(x \leftarrow e; c, m, F) &\rightsquigarrow \text{unit}(c, m\{\llbracket e \rrbracket_{\mathcal{E}} m/x\}, F) \\
(x \stackrel{\$}{\leftarrow} d; c, m, F) &\rightsquigarrow \text{bind}(\llbracket d \rrbracket_{\mathcal{DE}} m) (\lambda v. \text{unit}(c, m\{v/x\}, F)) \\
(x \leftarrow p(e); c, m, F) &\rightsquigarrow \text{unit}(p.\text{body}, (\emptyset\{\llbracket e \rrbracket_{\mathcal{E}} m/p.\text{args}\}, m.\text{glob}), \\
&\quad (x, p.\text{re}, c, m.\text{loc}) :: F) \\
(\text{if } e \text{ then } c_1 \text{ else } c_2; c, m, F) &\rightsquigarrow \begin{cases} \text{unit}(c_1; c, m, F) & \text{if } \llbracket e \rrbracket_{\mathcal{E}} m = \text{true} \\ \text{unit}(c_2; c, m, F) & \text{if } \llbracket e \rrbracket_{\mathcal{E}} m = \text{false} \end{cases} \\
(\text{while } e \text{ do } c; c', m, F) &\rightsquigarrow \begin{cases} \text{unit}(c; \text{while } e \text{ do } c; c', m, F) & \text{if } \llbracket e \rrbracket_{\mathcal{E}} m = \text{true} \\ \text{unit}(c', m, F) & \text{if } \llbracket e \rrbracket_{\mathcal{E}} m = \text{false} \end{cases}
\end{aligned}$$


---

**Fig. 2.1.** Probabilistic one-step semantics of pWHILE programs.

$$\begin{aligned}
f_{|\text{final}}(c, m, F) &: \mathcal{S} \rightarrow [0, 1] \\
f_{|\text{final}}(c, m, F) &\stackrel{\text{def}}{=} \begin{cases} f(m) & \text{if } c = \text{skip} \wedge F = \text{nil} \\ 0 & \text{otherwise} \end{cases}
\end{aligned}$$

The set of final states grows monotonically as the number of execution steps increases, which implies that the sequence  $\llbracket (c, m, \text{nil}) \rrbracket^n f_{|\text{final}}$  is increasing because  $f$  is non-negative. Because, in addition, this sequence is upper bounded by 1, the least upper bound in the definition of the denotation of a command always exists and corresponds to the limit of the sequence.

Figure 2.2 summarizes the denotational semantics of commands as equations following from the above limit construction. The denotation of a program relates an initial memory to a (sub-)probability distribution over memories using the measure monad presented in the previous section:

$$\llbracket c \rrbracket : \mathcal{M} \rightarrow \mathcal{D}(\mathcal{M})$$

Note that the function  $\llbracket \cdot \rrbracket$  maps  $\mathcal{M}$  to  $\mathcal{D}(\mathcal{M})$ , but it is trivial—although less convenient—to define a semantic function  $\llbracket \cdot \rrbracket^{\#}$  from  $\mathcal{D}(\mathcal{M})$  to  $\mathcal{D}(\mathcal{M})$  using the bind operator of the monad:

$$\begin{aligned}
\llbracket c \rrbracket^{\#} &: \mathcal{D}(\mathcal{M}) \rightarrow \mathcal{D}(\mathcal{M}) \\
\llbracket c \rrbracket^{\#} &\stackrel{\text{def}}{=} \lambda \mu. \text{bind } \mu \llbracket c \rrbracket
\end{aligned}$$

We have shown that the semantics of programs maps memories to discrete distributions, provided expressions in  $\mathcal{DE}$  evaluate to distributions with countable support. We use this together with Lemma 2.3 to prove the soundness of some relational Hoare logic rules (namely, [Comp] and [Trans]) in Section 3.1.

---

$\llbracket \text{skip} \rrbracket m$	$= \text{unit } m$
$\llbracket i; c \rrbracket m$	$= \text{bind } (\llbracket i \rrbracket m) \llbracket c \rrbracket$
$\llbracket x \leftarrow e \rrbracket m$	$= \text{unit } m\{\llbracket e \rrbracket_{\mathcal{E}} m/x\}$
$\llbracket x \xleftarrow{s} d \rrbracket m$	$= \text{bind } (\llbracket d \rrbracket_{\mathcal{D}\mathcal{E}} m) (\lambda v. \text{unit } m\{v/x\})$
$\llbracket x \leftarrow p(\mathbf{e}) \rrbracket m$	$= \text{bind } (\llbracket p.\text{body} \rrbracket (\emptyset\{\llbracket \mathbf{e} \rrbracket_{\mathcal{E}} m/p.\text{args}\}, m.\text{glob}))$ $(\lambda m'. \text{unit } (m.\text{loc}, m'.\text{glob})\{\llbracket p.\text{re} \rrbracket_{\mathcal{E}} m'/x\})$
$\llbracket \text{if } e \text{ then } c_1 \text{ else } c_2 \rrbracket m$	$= \begin{cases} \llbracket c_1 \rrbracket m & \text{if } \llbracket e \rrbracket_{\mathcal{E}} m = \text{true} \\ \llbracket c_2 \rrbracket m & \text{if } \llbracket e \rrbracket_{\mathcal{E}} m = \text{false} \end{cases}$
$\llbracket \text{while } e \text{ do } c \rrbracket m$	$= \lambda f. \text{sup } (\lambda n. \llbracket \text{while } e \text{ do } c \rrbracket_n m f)$ where $\llbracket \text{while } e \text{ do } c \rrbracket_0 = \text{skip}$ $\llbracket \text{while } e \text{ do } c \rrbracket_{n+1} = \text{if } e \text{ then } c; \llbracket \text{while } e \text{ do } c \rrbracket_n$

---

**Fig. 2.2.** Denotational semantics of pWHILE programs.

### Computing probabilities

The advantage of using this monadic semantics is that, if we use an arbitrary function as a continuation to the denotation of a program, what we get (for free) as a result is its expected value w.r.t. the distribution of final memories. In particular, we can compute the probability of an event  $A$  (represented as a function in  $\mathcal{M} \rightarrow \mathbb{B}$ ) in the distribution obtained after executing a command  $c$  in an initial memory  $m$  by measuring its characteristic function  $\mathbf{1}_A$ :

$$\Pr[c, m : A] \stackrel{\text{def}}{=} \llbracket c \rrbracket m \mathbf{1}_A$$

For instance, one can verify that the denotation of  $x \xleftarrow{s} \{0, 1\}; y \xleftarrow{s} \{0, 1\}$  in an initial memory  $m$  is

$$\lambda f. \frac{1}{4} (f(m\{0, 0/x, y\}) + f(m\{0, 1/x, y\}) + f(m\{1, 0/x, y\}) + f(m\{1, 1/x, y\}))$$

and conclude that the probability of the event  $(x \Rightarrow y)$  after executing the command above is  $3/4$ .

In what follows, when writing probabilities we sometimes omit the initial memory  $m$ ; in that case one may safely assume that the memory is initially  $\emptyset$ , which maps variables to default values of the right type.

## 2.3 Probabilistic Polynomial-Time Programs

In general, cryptographic proofs reason about effective adversaries, consuming polynomially bounded resources. The complexity notion that captures this intuition, and which is pervasive in cryptographic proofs, is that of *strict probabilistic polynomial-time* [Goldreich 2001]. Concretely, a program is said to be strict

probabilistic polynomial-time (PPT) whenever there exists a polynomial bound (in some security parameter  $\eta$ ) on the cost of each possible execution, regardless of the outcome of its random choices. Said otherwise, a probabilistic program is PPT whenever the same program seen as a non-deterministic program terminates and the cost of each possible run is bounded by a polynomial.

Termination and efficiency are orthogonal. Consider, for instance, the following two programs:

$$\begin{aligned} c_1 &\stackrel{\text{def}}{=} b \leftarrow \text{true}; \text{ while } b \text{ do } b \stackrel{\$}{\leftarrow} \{0, 1\} \\ c_2 &\stackrel{\text{def}}{=} b \stackrel{\$}{\leftarrow} \{0, 1\}; \text{ if } b \text{ then while true do skip} \end{aligned}$$

The former terminates with probability 1 (it terminates within  $n$  iterations with probability  $1 - 2^{-n}$ ), but may take an arbitrarily large number of iterations to terminate. The latter terminates with probability  $1/2$ , but when it does, it takes only a constant time. We deal with termination and efficiency separately.

**Definition 2.5 (Termination).** *The probability that a program  $c$  terminates starting from an initial memory  $m$  is  $\Pr[c, m : \text{true}] = \llbracket c \rrbracket m \mathbf{1}$ . We say that a program  $c$  is absolutely terminating, and note it  $\text{lossless}(c)$ , iff it terminates with probability 1 in any initial memory,*

$$\text{lossless}(c) \stackrel{\text{def}}{=} \forall m. \Pr[c, m : \text{true}] = 1$$

To deal with efficiency, we non-intrusively instrument the semantics of our language to compute the cost of running a program. The instrumented semantics ranges over  $\mathcal{D}(\mathcal{M} \times \mathbb{N})$  instead of  $\mathcal{D}(\mathcal{M})$ . We recall that our semantics is implicitly parametrized by a security parameter  $\eta$ , on which we base our notion of complexity. Our characterization of PPT programs relies on an axiomatization of the execution time and memory usage of expressions:

- We postulate the execution time of each operator, in the form of a function that depends on the inputs of the operator—which corresponds to the so-called functional time model;
- We postulate for each datatype a size measure, in the form of a function that assigns to each value its memory footprint.

We stress that making complexity assumptions on operators is perfectly legitimate. It is a well-known feature of dependent type theories (as is the case of the calculus of Coq) that they cannot express the cost of the computations they purport without using computational reflection, i.e. formalizing an execution model, such as probabilistic Turing machines, within the theory itself and proving that functions in type theory denote machines that execute in polynomial time. In our opinion, such a step is overkill. A simpler solution to the problem is to restrict in as much as possible the set of primitive operators, so as to minimize the set of assumptions upon which the complexity proofs rely.

**Definition 2.6 (Polynomially bounded distribution).** *We say that a family of distributions  $\{\mu_\eta : \mathcal{D}(\mathcal{M} \times \mathbb{N})\}$  is  $(p, q)$ -bounded, where  $p$  and  $q$  are polynomials,*



whenever for every value of the security parameter  $\eta$  and any pair  $(m, n)$  occurring with non-zero probability in  $\mu_\eta$ , the size of values in  $m$  is bounded by  $p(\eta)$  and the cost  $n$  is bounded by  $q(\eta)$ . This notion can be formally defined by means of the range predicate introduced in Section 2.1.3:

$$\text{bounded}(p, q, \mu) \stackrel{\text{def}}{=} \forall \eta. \text{range}(\lambda(m, n). \forall x \in \mathcal{V}. |m(x)| \leq p(\eta) \wedge n \leq q(\eta)) \mu_\eta$$

**Definition 2.7 (Strict probabilistic polynomial-time program).** We say that a program  $c$  is strict probabilistic polynomial-time (PPT) iff it terminates absolutely, and there exist polynomial transformers  $F, G$  such that for every  $(p, q)$ -bounded distribution family  $\mu_\eta$ ,  $(\text{bind } \mu_\eta \llbracket c \rrbracket)$  is  $(F(p), q + G(p))$ -bounded.

We can recover some intuition by taking  $\mu = \text{unit}(m, 0)$  in the above definition. In this case, we may paraphrase the condition as follows: if the size of values in  $m$  is bounded by some polynomial  $p$ , and an execution of the program in  $m$  terminates with non-zero probability in memory  $m'$ , then the size of values in  $m'$  is bounded by the polynomial  $F(p)$ , and the cost of the execution is bounded by the polynomial  $G(p)$ . It is in this latter polynomial that bounds the cost of executing the program that we are ultimately interested. The increased complexity in the definition is needed for proving compositionality results, such as the fact that PPT programs are closed under sequential composition.

Although our formalization of termination and efficiency relies on semantic definitions, it is not necessary for users to reason directly about the semantics of a program to prove it meets those definitions. CertiCrypt implements a certified algorithm showing that every program without loops and recursive calls terminates absolutely.<sup>1</sup> We also provide another algorithm that, together with the first, establishes that a program is PPT provided that, additionally, the program does not contain expressions that might generate values of super-polynomial size or take a super-polynomial time when evaluated in a polynomially bounded memory.

#### *Exact bounds on execution time*

Extracting an exact security result from a reductionist game-based proof requires to lower bound the success probability of the reduction and to upper bound the overhead incurred in execution time. Computing a bound on the success probability is what takes most of the effort since it requires examining the whole sequence of games and a careful bookkeeping of the probability of events. On the other hand, bounding the overhead of a reduction only requires examining the last game in the sequence. While we have put a great effort in automating the computation of probability bounds and we developed an automated method to obtain asymptotic polynomial bounds on the execution time of reductions, we did not bother to provide a method to compute exact time bounds. To do so, we would need an alternative

<sup>1</sup> It is of course a weak result in terms of termination of probabilistic programs, but nevertheless sufficient as regards cryptographic applications. Extending our formalization to a certified termination analysis for loops is interesting, but orthogonal to our main goals.

cost-instrumented semantics that does not take into account the time spent in evaluating calls to oracles, but instead just records the number of queries that have been made. Assume that an adversary  $\mathcal{A}$  executes within time  $t$  (without taking into account oracle calls) and makes at most  $q_{\mathcal{O}_i}$  queries to oracle  $\mathcal{O}_i$ . Suppose we have a reduction where an adversary  $\mathcal{B}$  uses  $\mathcal{A}$  as a sub-procedure; assume wlog that  $\mathcal{B}$  only calls  $\mathcal{A}$  once and does not make any additional oracle calls. Then, we can argue that if  $\mathcal{B}$  executes within time  $t'$  without taking into account the cost of evaluating calls to  $\mathcal{A}$  (this could easily be computed by considering  $\mathcal{A}$  as an oracle for  $\mathcal{B}$ ), then  $\mathcal{B}$  executes within time

$$t + t' + \sum_i q_{\mathcal{O}_i} t_{\mathcal{O}_i}$$

where  $t_{\mathcal{O}_i}$  upper bounds the cost one query to oracle  $\mathcal{O}_i$ .

## 2.4 Adversaries

In order to reason about games in the presence of unknown adversaries, we must specify an interface for adversaries and construct proofs under the assumption that adversaries are well-formed against their specification. Assuming that adversaries respect their interface provides us with an induction principle to reason over all (well-formed) adversaries. We make an extensive use of this induction principle: each time a proof system is introduced, the principle allows us to establish proof rules for adversaries. Likewise, each time we implement a program transformation, the induction principle allows us to prove the correctness of the transformation for programs that contain procedure calls to adversaries.

Formally, the interface of an adversary consists of a triple  $(\mathcal{O}, \mathcal{RW}, \mathcal{R})$ , where  $\mathcal{O}$  is the set of procedures that the adversary may call,  $\mathcal{RW}$  the set of variables that it may read and write, and  $\mathcal{R}$  the set of variables that it may only read. We say that an adversary  $\mathcal{A}$  with interface  $(\mathcal{O}, \mathcal{RW}, \mathcal{R})$  is well-formed if the judgment  $\vdash_{\text{wf}} \mathcal{A}$  can be derived from the rules in Figure 2.3. Note that the rules are generic, only making sure that the adversary makes a correct use of variables and procedures. These rules guarantee that a well-formed adversary always initializes local variables before using them, only writes global variables in  $\mathcal{RW}$ , and only reads global variables in  $\mathcal{RW} \cup \mathcal{R}$ . For convenience, we allow adversaries to call procedures outside  $\mathcal{O}$ , but these procedures must themselves respect the same interface.

Additional constraints may be imposed on adversaries. For example, exact security proofs usually impose an upper bound to the number of calls adversaries can make to a given oracle, while some properties, such as IND-CCA2 (see §2.5.2 below), restrict the parameters with which oracles may be called at different stages in an experiment. Likewise, some proofs impose extra conditions such as forbidding repeated or malformed queries. These kinds of properties can be formalized using global variables that record calls to oracles and verifying as post-condition that all calls were legitimate.

---


$$\begin{array}{c}
I \vdash \text{skip} : I \quad \frac{I \vdash i : I' \quad I' \vdash c : O}{I \vdash i; c : O} \quad \frac{\text{writable}(x) \quad \text{fv}(e) \subseteq I}{I \vdash x \leftarrow e : I \cup \{x\}} \quad \frac{\text{writable}(x) \quad \text{fv}(d) \subseteq I}{I \vdash x \leftarrow^* d : I \cup \{x\}} \\
\frac{\text{fv}(e) \subseteq I \quad I \vdash c_1 : O_1 \quad I \vdash c_2 : O_2}{I \vdash \text{if } e \text{ then } c_1 \text{ else } c_2 : O_1 \cap O_2} \quad \frac{\text{fv}(e) \subseteq I \quad I \vdash c : I}{I \vdash \text{while } e \text{ do } c : I} \\
\frac{\text{fv}(\mathbf{e}) \subseteq I \quad \text{writable}(x) \quad p \in \mathcal{O}}{I \vdash x \leftarrow p(\mathbf{e}) : I \cup \{x\}} \quad \frac{\text{fv}(\mathbf{e}) \subseteq I \quad \text{writable}(x) \quad \vdash_{\text{wf}} \mathcal{B}}{I \vdash x \leftarrow \mathcal{B}(\mathbf{e}) : I \cup \{x\}} \\
\frac{\mathcal{RW} \cup \mathcal{R} \cup \mathcal{A}.\text{args} \vdash \mathcal{A}.\text{body} : \mathcal{O} \quad \text{fv}(\mathcal{A}.\text{re}) \subseteq \mathcal{O}}{\vdash_{\text{wf}} \mathcal{A}} \\
\text{writable}(x) \stackrel{\text{def}}{=} \text{local}(x) \vee x \in \mathcal{RW}
\end{array}$$


---

**Fig. 2.3.** Rules for well-formedness of an adversary against interface  $(\mathcal{O}, \mathcal{RW}, \mathcal{R})$ . A judgment of the form  $I \vdash c : O$  can be interpreted as follows: assuming variables in  $I$  may be read, the adversarial code fragment  $c$  respects the interface and after its execution variables in  $O$  may be read. Thus, if  $I \vdash c : O$ , then  $I \subseteq O$ .

## 2.5 Making Security Properties Precise

Before going any further in the formalization of cryptographic proofs, we need to be sure that the results that we prove are meaningful. Security definitions in cryptography have so many subtleties that it is not clear that the whole cryptographic community agrees even on the most fundamental of these definitions. To illustrate this point, let us analyze in detail two pervasive definitions that we use in subsequent chapters: the security of a signature scheme against existential forgery under adaptive chosen-message attacks (EF-CMA security), and the indistinguishability under adaptive chosen-ciphertext attacks (IND-CCA2 security) of an encryption scheme.

### 2.5.1 EF-CMA Security

We start by recalling the definition of digital signature schemes.

**Definition 2.8 (Digital signature scheme).** *A digital signature scheme is composed of a triple of algorithms:*

- Key generation: *Given a security parameter  $\eta$ , the key generation algorithm  $\text{KG}(\eta)$  returns a public/secret key pair  $(pk, sk)$ ;*
- Signing: *Given a secret key  $sk$  and a message  $m$ , the signing algorithm  $\text{Sign}(sk, m)$  produces a signature of  $m$  under  $sk$ ;*
- Verification: *Given a public key  $pk$ , a message  $m$ , and a purported signature  $\sigma$  for  $m$ , the verification algorithm  $\text{Verify}(pk, m, \sigma)$  returns a Boolean value indicating whether the signature is valid or not.*

*Key generation and signing may be probabilistic, while the verification algorithm is deterministic. We require that verification always succeeds for authentic signatures:*

for every pair of keys  $(pk, sk)$  that can be output by the key generation algorithm, and every message  $m$ , it must be the case that  $\text{Verify}(pk, m, \text{Sign}(sk, m)) = \text{true}$ .

We informally describe the way in which EF-CMA security is typically defined. The experiment begins by choosing a public verification key  $pk$  and a secret signing key  $sk$ , using the key-generation algorithm of the signature scheme. The public key is given to the forger, who can ask for the signature of messages of its choice to a signing oracle and eventually halts and outputs a message  $m$  together with a purported signature  $\sigma$ . The forger wins when the signature  $\sigma$  verifies. We say that the scheme is secure when the winning probability of any probabilistic polynomial-time forger is negligible. Since the forger could trivially win by asking for the signature of  $m$ , the forger is not allowed to query  $m$  to the signing oracle. Figure 2.4 depicts this experiment as a game.

<b>Game <math>G_{\text{EF}}^A</math> :</b> $S \leftarrow \text{nil};$ $(pk, sk) \leftarrow \mathcal{KG}(\eta);$ $(m, \sigma) \leftarrow \mathcal{A}(pk)$	<b>Oracle <math>\text{Sign}_{\mathcal{A}}(m)</math> :</b> $S \leftarrow m :: S;$ $\sigma \leftarrow \text{Sign}(sk, m);$ return $\sigma$
---	---

**Fig. 2.4.** The EF-CMA experiment; the signing oracle is instrumented to record the queries made by the forger.

The above definition seems unambiguous at first sight. There are however two ways of forbidding the adversary from querying  $m$  to the signing oracle. The first is simply to reject adversaries that may query  $m$  with non-zero probability; this amounts to restrict the quantification over adversaries—instead of considering all efficient forgers one considers only those forgers that do not query the message they output to the signing oracle. The second way is to test a posteriori whether the adversary queried  $m$  to its oracle, and to declare that it lost in this case. Following [Bellare et al. 2009], we call the former definitional style the *exclusion* (“E”) style and the latter the *penalty* (“P”) style. Both styles are common and used interchangeably in the literature; for instance [Bellare and Rogaway 1996; Katz and Wang 2003] use the penalty style while [Bellare and Rogaway 1993] uses the exclusion style. The question is whether the two styles result in equivalent definitions or not.

It should be clear that security in the penalty style implies security in the exclusion style. To see this, consider any efficient forger  $\mathcal{A}$  valid in the exclusion style definition. The forger is also valid in the penalty style definition and it achieves the same success probability. Since the probability of  $\mathcal{A}$  ever querying the message  $m$  whose signature it forges is 0, we have

$$\Pr [G_{\text{EF}}^A : \text{Verify}(pk, m, \sigma)] = \Pr [G_{\text{EF}}^A : \text{Verify}(pk, m, \sigma) \wedge m \notin S]$$

Implication in the other direction (i.e., that security in the exclusion style implies security in the penalty style definition) is not as evident. Given an adversary  $\mathcal{A}$

that forges a signature for a fresh message  $m$  with non-negligible probability, but that may ask the signature of the message  $m$  to the signing oracle with non-zero probability, can we construct an adversary that never queries  $m$  to the signing oracle and achieves comparable success in forging a signature? The theorem we present next proves that we can and, what is more, without any probability loss.

**Theorem 2.9.** *If a signature scheme is EF-CMA secure according to the exclusion style definition, then it is secure according to the penalty style definition.*

*Proof.* Let  $\mathcal{A}$  be an adversary against the EF-CMA security of the scheme in the penalty style definition. We exhibit an adversary  $\mathcal{B}$  that is valid in the exclusion style definition and outputs a successful forgery with at least the same probability as  $\mathcal{A}$ :

<b>Adversary <math>\mathcal{B}(pk)</math> :</b> $\overline{\mathcal{S}} \leftarrow \text{nil};$ $(m, \sigma) \leftarrow \mathcal{A}(pk);$ if $m \in \overline{\mathcal{S}}$ then $m \xleftarrow{\$} \{0, 1\}^k \setminus \overline{\mathcal{S}}$ return $(m, \sigma)$	<b>Oracle <math>\text{Sign}_{\mathcal{A}}(m)</math> :</b> $\overline{\mathcal{S}} \leftarrow m :: \overline{\mathcal{S}};$ $\sigma \leftarrow \text{Sign}_{\mathcal{B}}(m);$ return $\sigma$
---	---

The forger  $\mathcal{B}$  uses  $\mathcal{A}$  as a subroutine; it intercepts the signing queries that  $\mathcal{A}$  makes and answers them using oracle  $\text{Sign}_{\mathcal{A}}$ . This oracle just records the message queried and forwards it to the original signing oracle. When  $\mathcal{A}$  outputs a purported forgery  $(m, \sigma)$ ,  $\mathcal{B}$  checks if  $m$  ever appeared in a signing query and if it is the case, replaces  $m$  with a fresh message. We have that  $m \notin \overline{\mathcal{S}}$  and  $\overline{\mathcal{S}} = \mathcal{S}$  are post-conditions of game  $\text{G}_{\text{EF}}^{\mathcal{B}}$ , which implies that  $\mathcal{B}$  is a valid adversary according to the exclusion style definition. In addition,

$$\Pr [\text{G}_{\text{EF}}^{\mathcal{A}} : \text{Verify}(pk, m, \sigma) \wedge m \notin \mathcal{S}] \leq \Pr [\text{G}_{\text{EF}}^{\mathcal{B}} : \text{Verify}(pk, m, \sigma)]$$

□

For a matter of taste and definitional clarity, we define EF-CMA security using the penalty style.

**Definition 2.10 (EF-CMA security).** *A signature scheme  $(\mathcal{KG}, \text{Sign}, \text{Verify})$  is secure against existential forgeries under chosen-message attacks if the probability*

$$\Pr [\text{G}_{\text{EF}}^{\mathcal{A}} : \text{Verify}(pk, m, \sigma) \wedge m \notin \mathcal{S}]$$

*is negligible for any probabilistic polynomial-time adversary  $\mathcal{A}$ .*

The penalty and exclusion style definitions of EF-CMA security turned out to be perfectly equivalent. Indeed, this equivalence can be regarded as folklore. This may lead us to think that there is no point in analyzing this kind of subtle differences in security definitions. But such a way of thinking is perilous. We have been lucky that both formulations of EF-CMA security are equivalent. We will see in the next section that being sloppy can sometimes lead to consider as equivalent definitions that in reality are not.

## 2.5.2 IND-CCA2 Security


The notion of IND-CCA2 security for an encryption scheme is defined relative to a two-phase experiment where the adversary has access to a decryption oracle. The experiment begins by generating a pair of keys  $(pk, sk)$  and giving the public key  $pk$  to the adversary. In the first phase the adversary chooses two messages  $m_0$  and  $m_1$ . The challenger then tosses a fair coin  $b$ , encrypts  $m_b$  under  $pk$  and gives the resulting ciphertext  $\hat{c}$  back to the adversary. The adversary ends the second phase by outputting a guess  $\tilde{b}$  for the hidden bit  $b$ . Figure 2.5 depicts this experiment as a game. We say that the scheme is IND-CCA2 secure if no probabilistic polynomial-time adversary  $\mathcal{A}$  guesses  $b$  correctly with a probability non-negligibly greater than  $1/2$ .

<b>Game <math>G_{\text{INDCCA}}</math> :</b> $L_{\mathcal{D}} \leftarrow \text{nil};$ $(pk, sk) \leftarrow \mathcal{KG}(\eta);$ $(m_0, m_1) \leftarrow \mathcal{A}_1(pk);$ $b \xleftarrow{\$} \{0, 1\};$ $\hat{c} \leftarrow \mathcal{E}(pk, m_b);$ $\hat{c}_{\text{def}} \leftarrow \text{true};$ $\tilde{b} \leftarrow \mathcal{A}_2(\hat{c})$	<b>Oracle <math>\mathcal{D}_{\mathcal{A}}(c)</math> :</b> $L_{\mathcal{D}} \leftarrow (\hat{c}_{\text{def}}, c) :: L_{\mathcal{D}};$ $m \leftarrow \mathcal{D}(sk, c);$ <b>return</b> $m$
---	--

**Fig. 2.5.** The IND-CCA2 experiment; the decryption oracle is instrumented to record the queries made in each phase.

Observe that the adversary could trivially win by asking the decryption oracle to decrypt  $\hat{c}$ . Consequently, the definition forbids the adversary from querying  $\hat{c}$  to the decryption oracle. As in the definition of EF-CMA security in the previous section, there are two ways of enforcing this restriction, in a penalty style or in an exclusion style. In addition, we now face another dilemma: should we allow the adversary to query  $\hat{c}$  to its oracle in the first phase of the experiment or should we forbid such type of queries altogether? These two dimensions give rise to four different ways of formally defining IND-CCA2 security. Namely, in a penalty style, restricting the oracle queries only in the second phase of the experiment (IND-CCA2-SP) or in both (IND-CCA2-BP), and in an exclusion style, restricting the queries only in the second phase (IND-CCA2-SE) or in both phases (IND-CCA2-BE). There are some obvious relations between these definitions. As in the case of EF-CMA security, security in a penalty style definition implies security in the corresponding exclusion style definition. In a similar manner, security in the version of the definitions where the adversary is forbidden to query the challenge ciphertext  $\hat{c}$  just in the second phase implies security when this prohibition is extended to the first phase of the experiment. Figure 2.6 summarizes these and the remaining non-trivial relations between the different formulations of IND-CCA2 security.

Surprisingly, neither of the “B” style definitions implies security in the corresponding “S” variant. What is more, the penalty and the exclusion style definitions



**Fig. 2.6.** Relations between the different formulations of IND-CCA2 security. An implication  $X \rightarrow Y$  means that security according to definition  $X$  implies security according to  $Y$ . A negated implication is a separation result.

are not equivalent if the adversary is forbidden from querying  $\hat{c}$  to its oracle in both phases of the IND-CCA2 experiment. We will give a proof of the implication  $\text{IND-CCA2-SE} \rightarrow \text{IND-CCA2-SP}$  and a rough idea of how to construct pathological schemes that justify the separation results  $\text{IND-CCA2-BP} \not\rightarrow \text{IND-CCA2-SP}$  and  $\text{IND-CCA2-BE} \not\rightarrow \text{IND-CCA2-BP}$ ; the separation of  $\text{IND-CCA2-SE}$  and  $\text{IND-CCA2-BE}$  follows from the diagram. For further details the reader may refer to [Bellare et al. 2009], where these results were first reported.

**Theorem 2.11.** *If an encryption scheme  $(\mathcal{KG}, \mathcal{E}, \mathcal{D})$  is IND-CCA2-SE secure, then it is IND-CCA2-SP secure as well.*

*Proof.* We show that for any adversary  $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$  against the IND-CCA2-SP security of the scheme, there exists an IND-CCA2-SE adversary  $\mathcal{B}$  that guesses the hidden bit  $b$  with at least the same probability and does not query the challenge ciphertext in the second phase of the experiment.

In the first phase,  $\mathcal{B}$  behaves exactly as  $\mathcal{A}$ . When  $\mathcal{B}$  gets the challenge ciphertext  $\hat{c}$ , it calls  $\mathcal{A}_2(\hat{c})$  in a simulated environment where it replaces the decryption oracle with an oracle of its own;  $\mathcal{B}_2$  returns whatever  $\mathcal{A}_2$  returns. When  $\mathcal{A}_2$  makes a decryption query  $c$ , if  $c \neq \hat{c}$ , the simulated oracle responds by forwarding the query to the original oracle, otherwise it returns some fixed message  $\perp$ . It is easy to see that  $\mathcal{B}_2$  never queries the challenge  $\hat{c}$  to the decryption oracle. Moreover, we have

$$\begin{aligned} \Pr \left[ G_{\text{IND-CCA}}^{\mathcal{A}} : \tilde{b} = b \wedge (\text{true}, \hat{c}) \notin L_{\mathcal{D}} \right] &= \Pr \left[ G_{\text{IND-CCA}}^{\mathcal{B}} : \tilde{b} = b \wedge (\text{true}, \hat{c}) \notin L_{\mathcal{D}} \right] \\ &\leq \Pr \left[ G_{\text{IND-CCA}}^{\mathcal{B}} : \tilde{b} = b \right] \end{aligned}$$

which concludes the proof.  $\square$

**Theorem 2.12.** *The three separation results from Figure 2.6 hold.*

*Proof.*

*IND-CCA2-BP  $\not\rightarrow$  IND-CCA2-SP :*

Let  $(\mathcal{KG}_f, f, f^{-1})$  be a family of trapdoor one-way permutations and  $(\mathcal{KG}, \mathcal{E}, \mathcal{D})$  an IND-CCA2-BP secure encryption scheme. We show how to construct an encryption scheme  $(\overline{\mathcal{KG}}, \overline{\mathcal{E}}, \overline{\mathcal{D}})$  which is IND-CCA2-BP secure but not IND-CCA2-SP secure.

$\overline{\mathcal{KG}}(\eta) :$ $(pk, sk) \leftarrow \mathcal{KG}(\eta);$ $(pk_f, sk_f) \leftarrow \mathcal{KG}_f(\eta);$ $\hat{x} \xleftarrow{\$} \{0, 1\}^k;$ $\hat{y} \leftarrow f(pk_f, \hat{x});$ $\overline{pk} \leftarrow (pk, pk_f, \hat{y});$ $\overline{sk} \leftarrow (sk, \hat{x});$ $\text{return } (\overline{pk}, \overline{sk})$	$\overline{\mathcal{E}}((pk, pk_f, \hat{y}), m) :$ $\text{if } f(pk_f, m) = \hat{y} \text{ then}$ $\quad \text{return } 1 \parallel 1^k$ $\text{else}$ $\quad c \leftarrow \mathcal{E}(pk, m);$ $\quad \text{return } 0 \parallel c$	$\overline{\mathcal{D}}((sk, \hat{x}), s \parallel c) :$ $\text{if } s = 0 \text{ then}$ $\quad \text{return } \mathcal{D}(sk, c)$ $\text{else}$ $\quad \text{if } c = 1^k \text{ then}$ $\quad \quad \text{return } \hat{x}$ $\quad \text{else return } \perp$
--	--	---

The above scheme is devised in such a way that the ability of an adversary to query the challenge ciphertext in the first phase leads to an attack, but this attack is no longer possible if such a query is disallowed. The intuition is to introduce a weak message  $\hat{x}$  with a single ciphertext  $(1 \parallel 1^k)$ . This message should be hard to compute without the secret key of the scheme, but it can be trivially obtained by asking for the decryption of  $(1 \parallel 1^k)$ . In the other hand, the encryption algorithm of the scheme should be able to efficiently test if a given message equals the weak message  $\hat{x}$ . We include the message  $\hat{x}$  in plain as part of the secret key of the scheme, but conceal its value in the public key using a one-way permutation.

To show that the above scheme is not IND-CCA2-SP secure, consider the following adversary  $(\mathcal{A}_1, \mathcal{A}_2)$ :

<b>Adversary <math>\mathcal{A}_1(\overline{pk}) :</math></b> $m_0 \leftarrow \mathcal{D}(1 \parallel 1^k);$ $m_1 \xleftarrow{\$} \{0, 1\}^k \setminus \{m_0\};$ $\text{return } (m_0, m_1)$	<b>Adversary <math>\mathcal{A}_2(c) :</math></b> $\text{if } c = 1 \parallel 1^k \text{ then return } 0$ $\text{else return } 1$
--	--

This adversary guesses the hidden bit in the IND-CCA2 experiment with probability 1. However, it queries the challenge ciphertext to the decryption oracle with probability 1/2 during its first phase, and therefore this adversary does not do any better than a random guess according to the winning condition of the IND-CCA2-BP variant.

To see why the scheme is IND-CCA2-BP secure, observe that the only way to guess the hidden bit  $b$  with probability significantly greater than 1/2 is to either break the security of the original encryption scheme, or to somehow obtain the value of the weak message  $\hat{x}$ . Indeed, given an adversary  $\mathcal{B}$  that breaks the IND-CCA2-BP security of the scheme with non-negligible probability, one can construct an adversary  $\mathcal{A}$  against the IND-CCA2-BP security of the original scheme and an inverter  $\mathcal{I}$  for the one-way trapdoor permutation such that at least one of them succeeds with non-negligible probability [Bellare et al. 2009, Theorem 3.1].

*IND-CCA2-BE*  $\leftrightarrow$  *IND-CCA2-BP* :

Let  $(\mathcal{KG}, \mathcal{E}, \mathcal{D})$  be an encryption scheme IND-CCA2-BE secure. Again, using a trapdoor one-way permutation  $(\mathcal{KG}_f, f, f^{-1})$  we construct an encryption scheme



$(\overline{\mathcal{KG}}, \overline{\mathcal{E}}, \overline{\mathcal{D}})$  which is IND-CCA2-BE secure but not IND-CCA2-BP secure.

$$\begin{array}{l}
 \overline{\mathcal{KG}}(\eta) : \\
 (pk, sk) \leftarrow \mathcal{KG}(\eta); \\
 (pk_f, sk_f) \leftarrow \mathcal{KG}_f(\eta); \\
 \hat{x} \xleftarrow{\$} \{0, 1\}^k; \\
 \hat{y} \leftarrow f(pk_f, \hat{x}); \\
 \overline{pk} \leftarrow (pk, pk_f, \hat{y}); \\
 sk \leftarrow (sk, \hat{x}); \\
 \text{return } (\overline{pk}, sk)
 \end{array}
 \left|
 \begin{array}{l}
 \overline{\mathcal{E}}((pk, pk_f, \hat{y}), m) : \\
 \text{if } f(pk_f, m) = \hat{y} \text{ then} \\
 \quad w \xleftarrow{\$} \{0, 1\}^k; \\
 \quad \text{return } 1 \parallel w \\
 \text{else} \\
 \quad c \leftarrow \mathcal{E}(pk, m); \\
 \quad \text{return } 0 \parallel c
 \end{array}
 \left|
 \begin{array}{l}
 \overline{\mathcal{D}}((sk, \hat{x}), s \parallel c) : \\
 \text{if } s = 0 \text{ then} \\
 \quad \text{return } \mathcal{D}(sk, c) \\
 \text{else} \\
 \quad \text{if } |c| = k \text{ then} \\
 \quad \quad \text{return } \hat{x} \\
 \quad \text{else return } \perp
 \end{array}
 \right.$$

To show that the above scheme is not IND-CCA2-BP secure, consider the following adversary  $(\mathcal{A}_1, \mathcal{A}_2)$ :

$$\begin{array}{l}
 \text{Adversary } \mathcal{A}_1(\overline{pk}) : \\
 m_0 \leftarrow \mathcal{D}(1 \parallel 1^k); \\
 m_1 \xleftarrow{\$} \{0, 1\}^k \setminus \{m_0\}; \\
 \text{return } (m_0, m_1)
 \end{array}
 \left|
 \begin{array}{l}
 \text{Adversary } \mathcal{A}_2(s \parallel c) : \\
 \text{if } s = 1 \text{ then return } 0 \\
 \text{else return } 1
 \end{array}
 \right.$$

This adversary guesses the hidden bit in the IND-CCA2 game with probability 1 and queries the challenge ciphertext to the decryption oracle only with probability  $2^{-k}/2$ . Therefore, its winning probability according to IND-CCA2-BP is  $1 - 2^{-k}/2$ . Observe, however, that this adversary is not valid according to the IND-CCA2-BE variant because it queries the challenge ciphertext to the decryption oracle with non-zero probability.

To see why the scheme is IND-CCA2-BE secure, observe that the only way to guess the hidden bit  $b$  with probability significantly greater than  $1/2$  is to either break the IND-CCA2-BP security of the original encryption scheme, or to somehow obtain the value of the weak message  $\hat{x}$ . But a valid adversary cannot obtain the value  $\hat{x}$  from the decryption oracle because any ciphertext of the form  $(1 \parallel c)$  might be the challenge ciphertext with probability  $2^{-k}/2$ . Therefore,  $\hat{x}$  is concealed by the one-way permutation and any adversary that succeeds in obtaining it can be used to invert the permutation with non-negligible probability.

*IND-CCA2-BE*  $\not\Rightarrow$  *IND-CCA2-SE* :

Follows from the above separation results and the diagram in the figure.  $\square$

In the remainder whenever we talk about IND-CCA2 security we will be referring to the IND-CCA2-SE variant of the definition, which is together with IND-CCA2-SP the strongest variant according to the taxonomy in Figure 2.6.

**Definition 2.13 (IND-CCA2 security).** *An encryption scheme  $(\mathcal{KG}, \mathcal{E}, \mathcal{D})$  is IND-CCA2 secure if the advantage*

$$\left| \Pr \left[ G_{\text{INDCCA}}^{\mathcal{A}} : \tilde{b} = b \right] - \frac{1}{2} \right|$$

is negligible for any probabilistic polynomial-time adversary  $\mathcal{A}$  that does not query the decryption oracle with the challenge ciphertext during the second phase of the IND-CCA2 experiment.



# 3

## Reasoning about Games

ACCORDING to Shoup [2004], steps in game-based cryptographic proofs can be classified into three broad categories:

1. Transitions based on indistinguishability, which are typically justified by appealing to a decisional assumption (e.g. the DDH assumption);
2. Transitions based on failure events, where it is argued that two games behave identically unless a failure event occurs;
3. Bridging steps, which correspond to refactoring the code of games in a way that is not observable by adversaries. This is in general done to prepare the ground for applying a *lossy* transition of one of the above two classes.

A bridging step from a game  $G_1$  to a game  $G_2$  typically replaces a program fragment  $c_1$  by an observationally equivalent fragment  $c_2$ . In general, however,  $c_1$  and  $c_2$  are observationally equivalent only in the particular context where the substitution is done. We justify such transitions through a relational Hoare logic that generalizes observational equivalence through pre- and post-conditions that characterize the context where the substitution is valid. This relational Hoare logic may as well be used to establish (in)equalities between the probability of events in two games (as shown by the rules [PrEq] and [PrLe] below) and to establish program invariants that serve to justify other program transformations or more complex probabilistic reasoning.

### 3.1 Probabilistic Relational Hoare Logic (pRHL)

The relational Hoare logic that we propose elaborates on and extends to probabilistic programs Benton's 2004 relational Hoare logic. Benton's logic uses judgments of the form  $\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi$ , that relate two programs,  $c_1$  and  $c_2$ , w.r.t. a pre-condition  $\Psi$  and a post-condition  $\Phi$ , both defined as relations on deterministic states. Such a judgment states that for every pair of initial memories  $m_1, m_2$  satisfying the pre-condition  $\Psi$ , if the evaluations of  $c_1$  in  $m_1$  and  $c_2$  in  $m_2$  terminate

with final memories  $m'_1$  and  $m'_2$  respectively, then  $m'_1 \Phi m'_2$  holds. In a probabilistic setting, the evaluation of a program in an initial memory yields instead a (sub-)probability distribution over program memories. In order to give a meaning to a judgment like the above one, we therefore need to lift relations over memories to relations over distributions.<sup>1</sup> We use the mechanism presented in Section 2.1.

**Definition 3.1 (pRHL judgment).** *We say that two programs  $c_1$  and  $c_2$  are equivalent with respect to pre-condition  $\Psi$  and post-condition  $\Phi$  iff*

$$\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi \stackrel{\text{def}}{=} \forall m_1 m_2. m_1 \Psi m_2 \implies (\llbracket c_1 \rrbracket m_1) \Phi^\# (\llbracket c_2 \rrbracket m_2)$$

**Definition 3.2 (Semantic equivalence).** *We say that two programs  $c_1$  and  $c_2$  are semantically equivalent, and note it as  $\vdash c_1 \equiv c_2$ , if they are equivalent w.r.t equality on memories as pre- and post-condition.*

Rather than defining the rules for pRHL and proving them sound in terms of the meaning of judgments, we place ourselves in a semantic setting and derive the rules as lemmas. This allows to easily extend the system by deriving extra rules, or even to resort to the semantic definition if the system turns out to be insufficient. Figure 3.1 gathers some representative derived rules. To improve readability, we define for a Boolean expression  $e$  the relations

$$e\langle 1 \rangle \stackrel{\text{def}}{=} \lambda m_1 m_2. \llbracket e \rrbracket m_1 = \text{true} \quad e\langle 2 \rangle \stackrel{\text{def}}{=} \lambda m_1 m_2. \llbracket e \rrbracket m_2 = \text{true}$$

As pRHL allows for arbitrary relations, we freely use higher-order logic; in particular, PER and SYM are predicates over relations that stand for *partial equivalence relation* and *symmetric relation* respectively.

Most rules admit, in addition to their symmetrical version of Figure 3.1, one-sided (left and right) variants, e.g. for assignments

$$\frac{m_1 \Psi m_2 = (m_1 \{ \llbracket e_1 \rrbracket m_1 / x_1 \}) \Phi m_2}{\vdash x_1 \leftarrow e_1 \sim \text{skip} : \Psi \Rightarrow \Phi} [\text{Assn1}]$$

The rule [Case] allows to reason by case analysis on the evaluation of an arbitrary relation in the initial memories. Together with simple rules in the spirit of

$$\frac{\vdash c_1 \sim c : \Psi \wedge e\langle 1 \rangle \Rightarrow \Phi}{\vdash \text{if } e \text{ then } c_1 \text{ else } c_2 \sim c : \Psi \wedge e\langle 1 \rangle \Rightarrow \Phi} [\text{Cond1T}]$$

it subsumes [Cond] and allows to prove judgments that otherwise would not be derivable, such as the semantic equivalence of the programs (if  $e$  then  $c_1$  else  $c_2$ ) and (if  $\neg e$  then  $c_2$  else  $c_1$ ):

<sup>1</sup> An alternative would be to develop a logic in which  $\Psi$  and  $\Phi$  are relations over distributions. However, we do not believe such a logic would allow a similar level of proof automation.

$$\begin{array}{c}
\vdash \text{skip} \sim \text{skip} : \Phi \Rightarrow \Phi \text{ [Skip]} \quad \frac{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Theta \quad \vdash c'_1 \sim c'_2 : \Theta \Rightarrow \Phi}{\vdash c_1; c'_1 \sim c_2; c'_2 : \Psi \Rightarrow \Phi} \text{[Seq]} \\
\\
\frac{m_1 \Psi \ m_2 = (m_1 \{ \llbracket e_1 \rrbracket m_1 / x_1 \}) \ \Phi \ (m_2 \{ \llbracket e_2 \rrbracket m_2 / x_2 \})}{\vdash x_1 \leftarrow e_1 \sim x_2 \leftarrow e_2 : \Psi \Rightarrow \Phi} \text{[Assn]} \\
\\
\frac{m_1 \Psi \ m_2 \Rightarrow (\llbracket d_1 \rrbracket m_1) \ \Theta^\# (\llbracket d_2 \rrbracket m_2) \quad \text{where } v_1 \ \Theta \ v_2 = (m_1 \{ v_1 / x_1 \}) \ \Phi \ (m_2 \{ v_2 / x_2 \})}{\vdash x_1 \stackrel{\#}{\sim} d_1 \sim x_2 \stackrel{\#}{\sim} d_2 : \Psi \Rightarrow \Phi} \text{[Rnd]} \\
\\
\frac{m_1 \Psi \ m_2 \Rightarrow \llbracket e_1 \rrbracket m_1 = \llbracket e_2 \rrbracket m_2 \quad \vdash c_1 \sim c_2 : \Psi \wedge e_1 \langle 1 \rangle \Rightarrow \Phi \quad \vdash c'_1 \sim c'_2 : \Psi \wedge \neg e_1 \langle 1 \rangle \Rightarrow \Phi}{\vdash \text{if } e_1 \text{ then } c_1 \text{ else } c'_1 \sim \text{if } e_2 \text{ then } c_2 \text{ else } c'_2 : \Psi \Rightarrow \Phi} \text{[Cond]} \\
\\
\frac{m_1 \Phi \ m_2 \Rightarrow \llbracket e_1 \rrbracket m_1 = \llbracket e_2 \rrbracket m_2 \quad \vdash c_1 \sim c_2 : \Phi \wedge e_1 \langle 1 \rangle \Rightarrow \Phi}{\vdash \text{while } e_1 \text{ do } c_1 \sim \text{while } e_2 \text{ do } c_2 : \Phi \Rightarrow \Phi \wedge \neg e_1 \langle 1 \rangle} \text{[While]} \\
\\
\frac{\Psi \subseteq \Psi' \quad \vdash c_1 \sim c_2 : \Psi' \Rightarrow \Phi' \quad \Phi' \subseteq \Phi}{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi} \text{[Sub]} \\
\\
\frac{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi \quad \text{SYM}(\Psi) \quad \text{SYM}(\Phi)}{\vdash c_2 \sim c_1 : \Psi \Rightarrow \Phi} \text{[Sym]} \\
\\
\vdash c \equiv c \text{ [Ref]} \quad \frac{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi \quad \vdash c_2 \sim c_3 : \Psi \Rightarrow \Phi \quad \text{PER}(\Psi) \quad \text{PER}(\Phi)}{\vdash c_1 \sim c_3 : \Psi \Rightarrow \Phi} \text{[Trans]} \\
\\
\frac{\vdash c_1 \sim c_2 : \Psi \wedge \Psi' \Rightarrow \Phi \quad \vdash c_1 \sim c_2 : \Psi \wedge \neg \Psi' \Rightarrow \Phi}{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi} \text{[Case]}
\end{array}$$

Fig. 3.1. Selection of derived rules of pRHL.

$$\begin{array}{c}
\frac{}{\vdash c_1 \sim c_1 : = \wedge \neg \neg e \langle 2 \rangle \Rightarrow =} \text{[Sub, Ref]} \\
\frac{}{\vdash c_1 \sim \text{if } \neg e \text{ then } c_2 \text{ else } c_1 : = \wedge \neg \neg e \langle 2 \rangle \Rightarrow =} \text{[Cond2F]} \\
\frac{}{\vdash c_1 \sim \text{if } \neg e \text{ then } c_2 \text{ else } c_1 : = \wedge e \langle 1 \rangle \Rightarrow =} \text{[Sub]} \\
\frac{}{\vdash \text{if } e \text{ then } c_1 \text{ else } c_2 \sim \text{if } \neg e \text{ then } c_2 \text{ else } c_1 : = \wedge e \langle 1 \rangle \Rightarrow =} \text{[Cond1T]} \dots \\
\frac{}{\vdash \text{if } e \text{ then } c_1 \text{ else } c_2 \equiv \text{if } \neg e \text{ then } c_2 \text{ else } c_1} \text{[Case]}
\end{array}$$

We use [Case] as well to justify the correctness of dataflow analyses that exploit the information provided by entering branches.

The rule [Sym] can be generalized by taking the inverse of the relations instead of requiring that pre- and post-condition be symmetric:

$$\frac{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi}{\vdash c_2 \sim c_1 : \Psi^{-1} \Rightarrow \Phi^{-1}} \text{[Inv]}$$

The rule [Trans], although appealing, is of limited practical use. Consider, for instance, “independent” pre- and post-conditions of the form

$$m_1 \Psi m_2 \stackrel{\text{def}}{=} \Psi_1 m_1 \wedge \Psi_2 m_2 \quad m_1 \Phi m_2 \stackrel{\text{def}}{=} \Phi_1 m_1 \wedge \Phi_2 m_2$$

In order to apply the rule [Trans], we are essentially forced to have

$$\Psi_1 = \Psi_2 \quad \text{and} \quad \Phi_1 = \Phi_2$$

and we must also choose the same pre- and post-condition for the intermediate game  $c_2$ . This constraints make the rule [Trans] impractical in some cases; we use instead the rule [Comp] to introduce intermediate games in those cases:

$$\frac{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi \quad \vdash c_2 \sim c_3 : \Psi' \Rightarrow \Phi'}{\vdash c_1 \sim c_3 : \Psi \circ \Psi' \Rightarrow \Phi \circ \Phi'} [\text{Comp}]$$

The soundness of this rule relies on Lemma 2.3 and on the fact that the denotation of a program maps an initial memory to a distribution with countable support. This is true if we only allow values to be sampled from distributions with countable support, a reasonable restriction that does not affect our application to cryptographic proofs.

We can specialize rule [Rnd] when the distributions from where random values are sampled have countable support. In this case, there is a simpler condition that makes the hypothesis of the rule hold. We say that two distributions  $\mu_1 : \mathcal{D}(A)$  and  $\mu_2 : \mathcal{D}(B)$  with countable support are equivalent modulo a relation  $R \subseteq A \times B$ , and note it  $\mu_1 \simeq_R \mu_2$ , when there exists a bijection  $f : \text{support}(\mu_1) \rightarrow \text{support}(\mu_2)$  such that

$$\forall a \in \text{support}(\mu_1). \mu_1 \mathbb{I}_{\{a\}} = \mu_2 \mathbb{I}_{\{f(a)\}} \wedge R(a, f(a))$$

We can then prove that the following rule is sound:

$$\frac{m_1 \Psi m_2 \implies \llbracket d_1 \rrbracket m_1 \simeq_{\Theta} \llbracket d_2 \rrbracket m_2 \quad v_1 \Theta v_2 = (m_1 \{v_1/x_1\}) \Phi (m_2 \{v_2/x_2\})}{\vdash x_1 \stackrel{\#}{\sim} d_1 \sim x_2 \stackrel{\#}{\sim} d_2 : \Psi \Rightarrow \Phi} [\text{Perm}]$$

If  $d_1$  and  $d_2$  are both interpreted as uniform distributions over some set of values, the premise of the rule boils down to exhibiting a bijection  $f$  between the supports of  $(\llbracket d_1 \rrbracket m_1)$  and  $(\llbracket d_2 \rrbracket m_2)$  such that  $\Theta(v, f(v))$  holds for any  $v$  in the support of  $\llbracket d_1 \rrbracket m_1$ . To see that the rule is sound, note that  $\mu_1 \simeq_R \mu_2$  implies  $\mu_1 R^{\#} \mu_2$ ; it suffices to take the following distribution as a witness for the existential:

$$\mu \stackrel{\text{def}}{=} \text{bind } \mu_1 (\lambda v. \text{unit}(v, f(v)))$$

Hence, the soundness of the above rule is immediate from the soundness of rule [Rnd]. Section 3.2.2 shows that rule [Perm] is enough to prove several program equivalences appearing in cryptographic proofs. However, observe that rule [Perm] is far from being complete as shown by the following program equivalence that cannot be derived using just this rule:

$$\vdash a \stackrel{\#}{\sim} [0..1] \sim b \stackrel{\#}{\sim} [0..3]; a \leftarrow b \bmod 2 : \text{true} \Rightarrow =_{\{a\}}$$

One cannot use the above rule to prove such an equivalence because the supports of the distributions from where random values are sampled in the programs do not

have the same size and hence it is not possible to find a bijection relating them. We can further generalize the rule to prove the above equivalence by requiring instead the existence of a bijection between the support of one distribution and a partition of the support of the other, as in the following rule:

$$\frac{m_1 \Psi m_2 \implies \text{let } S_1 = \text{support}(\llbracket d_1 \rrbracket m_1), S_2 = \text{support}(\llbracket d_2 \rrbracket m_2) \text{ in} \\ \exists f : S_1 \rightarrow \mathcal{P}(S_2). \bigcup_{v \in S_1} f(v) = S_2 \wedge (\forall v_1 \neq v_2 \in S_1. f(v_1) \cap f(v_2) = \emptyset) \wedge \\ (\forall v \in S_1. \mu_1 \mathbb{I}_{\{v\}} = \mu_2 \mathbb{I}_{f(v)} \wedge \forall w \in f(v). (m_1\{v/x_1\}) \Phi (m_2\{w/x_2\}))}{\vdash x_1 \stackrel{\text{ex}}{\sim} d_1 \sim x_2 \stackrel{\text{ex}}{\sim} d_2 : \Psi \Rightarrow \Phi}$$

The following two rules allow to fall back from the world of pRHL into the world of probabilities, in which security statements are expressed:

$$\frac{m_1 \Psi m_2 \quad \vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi \quad \Phi \implies (A\langle 1 \rangle \iff B\langle 2 \rangle)}{\text{Pr}[c_1, m_1 : A] = \text{Pr}[c_2, m_2 : B]} [\text{PrEq}]$$

and analogously,

$$\frac{m_1 \Psi m_2 \quad \vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi \quad \Phi \implies (A\langle 1 \rangle \implies B\langle 2 \rangle)}{\text{Pr}[c_1, m_1 : A] \leq \text{Pr}[c_2, m_2 : B]} [\text{PrLe}]$$

By taking  $A = B = \text{true}$  we can observe that observational equivalence enjoys some form of termination sensitivity:

$$(\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi) \wedge m_1 \Psi m_2 \implies \llbracket c_1 \rrbracket m_1 \mathbf{1} = \llbracket c_2 \rrbracket m_2 \mathbf{1}$$

We conclude with an example that nicely illustrates some of the intricacies of pRHL. Let  $c = b \stackrel{\text{ex}}{\sim} \{0, 1\}$  and  $\Phi = (b\langle 1 \rangle = b\langle 2 \rangle)$ . We have for any pair of initial memories  $(\llbracket c \rrbracket m_1) \Phi^\# (\llbracket c \rrbracket m_2)$ . Indeed, the following distribution is a witness for the existential of the lifting:

$$\mu f = \frac{1}{2} f(m_1\{0/b\}, m_2\{0/b\}) + \frac{1}{2} f(m_1\{1/b\}, m_2\{1/b\})$$

Perhaps more surprisingly, we also have  $(\llbracket c \rrbracket m_1) \neg\Phi^\# (\llbracket c \rrbracket m_2)$ , for which it suffices to take the following distribution as a witness for the existential:

$$\mu' f = \frac{1}{2} f(m_1\{0/b\}, m_2\{1/b\}) + \frac{1}{2} f(m_1\{1/b\}, m_2\{0/b\})$$

Thus, we have at the same time  $\vdash c \sim c : \text{true} \Rightarrow \Phi$  and  $\vdash c \sim c : \text{true} \Rightarrow \neg\Phi$  (but of course not  $\vdash c \sim c : \text{true} \Rightarrow \text{false}$ ) and as a consequence the ‘‘obvious’’ rule

$$\frac{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi \quad \vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi'}{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi \wedge \Phi'}$$

is unsound. While this example may seem unintuitive or even inconsistent if one reasons in terms of deterministic states, its intuitive significance in a probabilistic



setting is that observing either  $\Phi$  or  $\neg\Phi$  is not enough to tell apart the distributions resulting from two executions of  $c$ . This example shows why lifting a relation to distributions involves an existential quantification, and why it is not possible to always use the product distribution as a witness (one cannot establish neither of the above judgments using the product distribution). This interpretation of pRHL judgments is strongly connected to the relation between relational logics and information flow [Amtoft et al. 2006; Benton 2004]—formally characterized for instance by Benton’s embedding of a type system for secure information flow into RHL.

As an additional example, observe that we have

$$\begin{aligned} \vdash x \stackrel{\mathbb{S}}{\sim} \{0, 1\}; y \stackrel{\mathbb{S}}{\sim} \{0, 1\} &\sim x \stackrel{\mathbb{S}}{\sim} \{0, 1\}; y \leftarrow x : \text{true} \Rightarrow =_{\{x\}} \\ \vdash x \stackrel{\mathbb{S}}{\sim} \{0, 1\}; y \stackrel{\mathbb{S}}{\sim} \{0, 1\} &\sim x \stackrel{\mathbb{S}}{\sim} \{0, 1\}; y \leftarrow x : \text{true} \Rightarrow =_{\{y\}} \end{aligned}$$

but clearly the following judgment does not hold

$$\vdash x \stackrel{\mathbb{S}}{\sim} \{0, 1\}; y \stackrel{\mathbb{S}}{\sim} \{0, 1\} \sim x \stackrel{\mathbb{S}}{\sim} \{0, 1\}; y \leftarrow x : \text{true} \Rightarrow =_{\{x,y\}}$$

since after executing the program on the right-hand side the values of  $x$  and  $y$  always coincide while this happens only with probability  $1/2$  for the program on the left-hand side.

### 3.1.1 Observational Equivalence

Observational equivalence is derived as an instance of relational Hoare judgments in which pre- and post-conditions are restricted to equality over a subset of program variables. Observational equivalence of programs  $c_1, c_2$  w.r.t. an input set of variables  $I$  and an output set of variables  $O$  is defined as

$$\vdash c_1 \simeq_O^I c_2 \stackrel{\text{def}}{=} \vdash c_1 \sim c_2 :=_I \Rightarrow =_O$$

The rules of pRHL can be specialized to the case of observational equivalence. For example, for conditional statements we have

$$\frac{m_1 =_I m_2 \implies \llbracket e_1 \rrbracket m_1 = \llbracket e_2 \rrbracket m_2 \quad \vdash c_1 \simeq_O^I c_2 \quad \vdash c'_1 \simeq_O^I c'_2}{\vdash \text{if } e_1 \text{ then } c_1 \text{ else } c'_1 \simeq_O^I \text{if } e_2 \text{ then } c_2 \text{ else } c'_2}$$

It follows that observational equivalence is symmetric and transitive, although it is not reflexive. Indeed, observational equivalence can be seen as a generalization of probabilistic non-interference: if we take  $I = O = L$ , the set of *low* variables, then  $c$  is non-interferent iff  $\vdash c \simeq_L^L c$ .

Observational equivalence is more amenable to mechanization than full-fledged pRHL. To support automation, CertiCrypt implements a calculus of variable dependencies and provides a function `eqobs_in`, that given a program  $c$  and a set of output variables  $O$ , computes a set of input variables  $I$  such that  $\vdash c \simeq_O^I c$ . Analogously, it provides a function `eqobs_out`, that given a set of input variables

$I$ , computes a set of output variables  $O$  such that  $\vdash c \simeq_O^I c$ . This suggests a simple procedure to establish a self-equivalence of the form  $\vdash c \simeq_O^I c$ : just compute a set  $I'$  such that  $\vdash c \simeq_{O'}^{I'} c$  using `eqobs_in` and check whether  $I' \subseteq I$ , or equivalently, compute a set  $O'$  such that  $\vdash c \simeq_{O'}^I c$  using `eqobs_out` and check whether  $O \subseteq O'$ .

CertiCrypt provides as well a (sound, but incomplete) relational weakest precondition calculus that can be used to automate proofs of program invariants; it deals with judgments of the form

$$\vdash c_1 \sim c_2 : \Psi \Rightarrow =_O \wedge \Phi$$

and requires that the programs have (almost) the same control-flow structure.

## 3.2 Bridging Steps

CertiCrypt provides a powerful set of tactics and algebraic equivalences to automate bridging steps in proofs. Most tactics rely on an implementation of a certified optimizer for `pWHILE`. Algebraic equivalences are provided as lemmas that follow from algebraic properties of the interpretation of language constructs.

### 3.2.1 Certified Program Transformations

We automate several transformations that consist in applying compiler optimizations. More precisely, we provide support for a rich set of transformations based on dependency and dataflow analyses, and for inlining procedure calls in programs. Each transformation is implemented as a function in CertiCrypt that performs the transformation itself, together with a rule that proves its correctness and a tactic that applies the rule backwards.

*Transformations based on dependencies*

The functions `eqobs_in` and `eqobs_out` and the relational Hoare logic presented in Section 3.1 provide the foundations to support transformations such as dead code elimination and code reordering.

We write and prove the correctness of a function `context` that strips off two programs  $c_1$  and  $c_2$  their maximal common context relative to sets  $I$  and  $O$  of input and output variables. The correctness of `context` is expressed by the following rule

$$\frac{\text{context}(I, c_1, c_2, O) = (I', c'_1, c'_2, O') \quad \vdash c'_1 \simeq_{O'}^{I'} c'_2}{\vdash c_1 \simeq_O^I c_2}$$

The tactic `eqobs_ctxt` applies this rule backwards. Using the same idea, we implement tactics that strip off two programs only their common prefix (`eqobs_hd`) or suffix (`eqobs_tl`).

We provide a tactic (**swap**) that given two programs tries to hoist their common instructions to obtain a maximal common prefix<sup>2</sup>, which can then be eliminated using the above tactics. Its correctness is based on the rule

$$\frac{\begin{array}{c} \vdash c_1 \simeq_{O_1}^{I_1} c_1 \quad \vdash c_2 \simeq_{O_2}^{I_2} c_2 \quad \text{modifies}(c_1, O_1) \quad \text{modifies}(c_2, O_2) \\ O_1 \cap O_2 = \emptyset \quad I_1 \cap O_2 = \emptyset \quad I_2 \cap O_1 = \emptyset \end{array}}{\vdash c_1; c_2 \equiv c_2; c_1}$$

where  $\text{modifies}(c, X)$  is a semantic predicate expressing that program  $c$  only modifies variables in  $X$ . This is formally expressed by

$$\forall m. \text{range}(\lambda m'. m =_{\mathcal{V} \setminus X} m') (\llbracket c \rrbracket m)$$

which ensures that reachable final memories coincide with the initial memory except maybe on variables in  $X$ . The tactic **swap** uses an algorithm that over-approximates the set of modified variables to decide whether two instructions can be swapped.

We provide a tactic (**deadcode**) that performs dead code elimination relative to a set  $O$  of output variables. The corresponding transformation behaves more like an aggressive slicing algorithm: it removes portions of code that do not affect variables in  $O$  and performs at the same time branch prediction (substituting  $c_1$  for if **true then**  $c_1$  **else**  $c_2$ ), branch coalescing (substituting  $c$  for if  $e$  **then**  $c$  **else**  $c$ ), and self-assignment elimination. Its correctness relies on the rule

$$\frac{\text{modifies}(c, X) \quad \text{lossless}(c) \quad \text{fv}(\Phi) \cap X = \emptyset}{\vdash c \sim \text{skip} : \Phi \Rightarrow \Phi}$$

#### *Optimizations based on dataflow analyses*

CertiCrypt has built-in, generic, support for such optimizations: given an abstract domain  $D$  (a semi-lattice) for the analysis, transfer functions for assignment and branching instructions, and an operator that optimizes expressions in the language, we construct a certified optimization function  $\text{optimize} : \mathcal{C} \rightarrow D \rightarrow \mathcal{C} \times D$ . When given a command  $c$  and an element  $\delta \in D$ , this function transforms  $c$  into its optimized version  $c'$  assuming the validity of  $\delta$ . In addition, it returns an abstract post-condition  $\delta' \in D$ , valid after executing  $c$  (or  $c'$ ). We use these abstract post-conditions to state the correctness of the optimization and to apply it recursively. The correctness of **optimize** is proved using a mixture of the techniques of [Benton 2004] and [Bertot et al. 2006; Leroy 2006]: we express the validity of the information contained in the analysis domain using a predicate  $\text{valid}(\delta, m)$  that states the agreement between the compile time abstract values in  $\delta$  and the run time memory  $m$ . Correctness is expressed in terms of a pRHL judgment:

$$\mathbf{let} (c', \delta') := \text{optimize}(c, \delta) \mathbf{in} \vdash c \sim c' : \simeq_\delta \Rightarrow \simeq_{\delta'}$$

<sup>2</sup> One could also provide a complementary tactic that hoists instructions to obtain a maximal common suffix.

where  $m_1 \succ_\delta m_2 \stackrel{\text{def}}{=} m_1 = m_2 \wedge \text{valid}(\delta, m_1)$ . The following useful rule is derived using [Comp]:

$$\frac{m_1 \Psi m_2 \implies \text{valid}(\delta, m_1) \quad \text{optimize}(c_1, \delta) = (c'_1, \delta') \quad \vdash c'_1 \sim c_2 : \Psi \Rightarrow \Phi}{\vdash c_1 \sim c_2 : \Psi \Rightarrow \Phi} [\text{Opt}]$$

Our case studies extensively use instantiations of [Opt] to perform expression propagation (tactic `ep`). In contrast, we found that common subexpression elimination is seldom used.

### 3.2.2 Algebraic Equivalences

Bridging steps frequently make use of algebraic properties of language constructs. The proof of semantic security of ElGamal uses the fact that in a cyclic multiplicative group, multiplication by a uniformly sampled element acts as a one-time pad:

$$\vdash x \stackrel{\$}{\leftarrow} \mathbb{Z}_q; \alpha \leftarrow g^x \times \beta \simeq_{\{\alpha\}} y \stackrel{\$}{\leftarrow} \mathbb{Z}_q; \alpha \leftarrow g^y$$

In the proof of IND-CCA2 security of OAEP described in Section 6.1 we use the equivalences

$$\vdash x \stackrel{\$}{\leftarrow} \{0, 1\}^k; y \leftarrow x \oplus z \simeq_{\{x, y, z\}}^{\{z\}} y \stackrel{\$}{\leftarrow} \{0, 1\}^k; x \leftarrow y \oplus z$$

and (for a permutation  $f$ ):

$$\vdash x \stackrel{\$}{\leftarrow} \{0, 1\}^{k-\rho}; y \stackrel{\$}{\leftarrow} \{0, 1\}^\rho; z \leftarrow f(x||y) \simeq_{\{z\}} z \stackrel{\$}{\leftarrow} \{0, 1\}^k$$

We show the usefulness of rule [Perm] by proving the first of these two equivalences, known as *optimistic sampling*, that we also used in §1.4.2. Define

$$\begin{aligned} \Psi &\stackrel{\text{def}}{=} z\langle 1 \rangle = z\langle 2 \rangle \\ \Phi &\stackrel{\text{def}}{=} x\langle 1 \rangle = x\langle 2 \rangle \wedge y\langle 1 \rangle = y\langle 2 \rangle \wedge z\langle 1 \rangle = z\langle 2 \rangle \\ \Theta &\stackrel{\text{def}}{=} m_1\{x\langle 1 \rangle \oplus z\langle 1 \rangle / y\} \Phi m_2\{y\langle 2 \rangle \oplus z\langle 2 \rangle / x\} \\ &= x\langle 1 \rangle = y\langle 2 \rangle \oplus z\langle 2 \rangle \wedge z\langle 1 \rangle = z\langle 2 \rangle \end{aligned}$$

By rule [Assn] we have

$$\vdash y \leftarrow x \oplus z \sim x \leftarrow y \oplus z : \Theta \Rightarrow \Phi \tag{3.1}$$

We apply rule [Perm] to prove

$$\vdash x \stackrel{\$}{\leftarrow} \{0, 1\}^k \sim y \stackrel{\$}{\leftarrow} \{0, 1\}^k : \Psi \Rightarrow \Theta \tag{3.2}$$

For doing so we must show that for any pair of memories  $m_1, m_2$  that coincide on  $z$  there exists a permutation  $f$  on  $\{0, 1\}^k$  such that

$$\forall v \in \{0, 1\}^k. v = f(v) \oplus m_2(z) \wedge m_1(z) = m_2(z)$$

Take  $f(v) \stackrel{\text{def}}{=} v \oplus m_2(z)$  to be such a permutation. Conclude from (3.1) and (3.2) by a final application of rule [Seq].

### 3.2.3 Inter-procedural Code Motion

Game-based proofs commonly include bridging steps consisting in a semantics-preserving reordering of instructions. When the reordering is intra-procedural, the tactic `swap` presented in the previous section generally suffices to justify the transformation. However, proofs in the random oracle model (see §1.4.2 for an example of a random oracle) often include transformations where random values used inside oracles are sampled beforehand, or conversely, where sampling a random value at some point in a game is deferred to a later point, possibly in a different procedure. The former type of transformation, called eager sampling, is useful for moving random choices upfront: a systematic application of eager sampling transforms a probabilistic game  $G$  that samples a fixed number of values into a semantically equivalent game  $S; G'$ , where  $S$  samples the values that might be needed in  $G$ , and  $G'$  is a completely deterministic program to the exception of adversaries that may still make their own random choices.<sup>3</sup> The dual transformation, called lazy sampling, can be used to postpone sampling random values until they are actually used for the first time—thus, one readily knows the exact distribution of these values by reasoning locally, without the need to maintain and reason about probabilistic invariants. In this section, we present a general method to prove the correctness of inter-procedural code motion. The method is based on a logic for swapping statements that generalizes the earlier lemma reported in [Barthe et al. 2009c].

#### A logic for swapping statements

The primary tool for performing eager/lazy sampling is an extension of the relational Hoare logic with rules for swapping statements. As the goal is to move code across procedures, it is essential that the logic considers two potentially different environments  $E$  and  $E'$ . The logic deals with judgments of the form

$$\vdash E, (c; S) \sim E', (S; c') : \Psi \Rightarrow \Phi$$

In most cases, the logic will be applied with  $S$  being a sequence of (guarded) sampling statements; however, we do not constrain  $S$  and merely require that it satisfies three basic properties for some sets of variables  $X$  and  $I$ :

$$\text{modifies}(E, S, X) \quad \text{modifies}(E', S, X) \quad \vdash E, S \simeq_X^{I \cup X} E', S$$

Some rules of the logic are given in Figure 3.2; for the sake of readability all rules are specialized to  $\equiv$ , although we formalized more general versions of the rules, e.g. for conditional statements,

$$\frac{\begin{array}{l} \vdash E, (c_1; S) \sim E', (S; c'_1) : P \wedge e\langle 1 \rangle \Rightarrow Q \quad P \Longrightarrow e\langle 1 \rangle = e'\langle 2 \rangle \\ \vdash E, (c_2; S) \sim E', (S; c'_2) : P \wedge \neg e\langle 1 \rangle \Rightarrow Q \quad \text{fv}(e') \cap X = \emptyset \end{array}}{\vdash E, (\text{if } e \text{ then } c_1 \text{ else } c_2; S) \sim E', (S; \text{if } e' \text{ then } c'_1 \text{ else } c'_2) : P \Rightarrow Q} \text{ [S-Cond]}$$

<sup>3</sup> Making adversaries deterministic is the goal of the *coin fixing* technique, as described by Bellare and Rogaway [2006].

$$\begin{array}{c}
\frac{x \notin I \cup X \quad \text{fv}(e) \cap X = \emptyset}{\vdash E, (x \leftarrow e; S) \equiv E', (S; x \leftarrow e)} \text{[S-Assn]} \quad \frac{x \notin I \cup X \quad \text{fv}(d) \cap X = \emptyset}{\vdash E, (x \leftarrow d; S) \equiv E', (S; x \leftarrow d)} \text{[S-Rnd]} \\
\frac{\vdash E, (c_1; S) \equiv E', (S; c'_1) \quad \vdash E, (c_2; S) \equiv E', (S; c'_2)}{\vdash E, (c_1; c_2; S) \equiv E', (S; c'_1; c'_2)} \text{[S-Seq]} \\
\frac{\vdash E, (c_1; S) \equiv E', (S; c'_1) \quad \vdash E, (c_2; S) \equiv E', (S; c'_2) \quad \text{fv}(e) \cap X = \emptyset}{\vdash E, (\text{if } e \text{ then } c_1 \text{ else } c_2; S) \equiv E', (S; \text{if } e \text{ then } c'_1 \text{ else } c'_2)} \text{[S-Cond]} \\
\frac{\vdash E, (c; S) \equiv E', (S; c') \quad \text{fv}(e) \cap X = \emptyset}{\vdash E, (\text{while } e \text{ do } c; S) \equiv E', (S; \text{while } e \text{ do } c')} \text{[S-While]} \\
\frac{\vdash E, (f.\text{body}; S) \equiv E', (S; f.\text{body}) \quad E(f).\text{args} = E'(f).\text{args} \quad E(f).\text{re} = E'(f).\text{re} \quad \text{fv}(E(f).\text{re}) \cap X = \emptyset \quad x \notin I \cup X \quad \text{fv}(e) \cap X = \emptyset}{\vdash E, (x \leftarrow f(e); S) \equiv E', (S; x \leftarrow f(e))} \text{[S-Call]} \\
\frac{\vdash_{\text{wf}} \mathcal{A} \quad X \cap (\mathcal{RW} \cup \mathcal{R}) = \emptyset \quad I \cap \mathcal{RW} = \emptyset \quad \forall f \notin \mathcal{O}. E(f) = E'(f) \quad \forall f \in \mathcal{O}. E(f).\text{args} = E'(f).\text{args} \wedge E(f).\text{re} = E'(f).\text{re} \wedge \vdash E, (f.\text{body}; S) \equiv E', (S; f.\text{body})}{\vdash E, (x \leftarrow \mathcal{A}(e); S) \equiv E', (S; x \leftarrow \mathcal{A}(e))} \text{[S-Adv]}
\end{array}$$

**Fig. 3.2.** Selected rules of a logic for swapping statements.

### An application

Consider the games  $G_{\text{lazy}}$  and  $G_{\text{eager}}$  in Figure 3.3. Both games define an oracle

<p><b>Game <math>G_{\text{lazy}}</math> :</b>  <math>L \leftarrow \text{nil}; b \leftarrow \mathcal{A}()</math></p> <p><b>Oracle <math>\mathcal{O}_{\text{lazy}}(x)</math> :</b>  if <math>x \notin \text{dom}(L)</math> then  <math>y \leftarrow \{0, 1\}^\ell</math>;  <math>L \leftarrow (x, y) :: L</math>  else <math>y \leftarrow L[x]</math>  return <math>y</math></p>	<p><b>Game <math>G_{\text{eager}}</math> :</b>  <math>L \leftarrow \text{nil}; \hat{y} \leftarrow \{0, 1\}^\ell; b \leftarrow \mathcal{A}()</math></p> <p><b>Oracle <math>\mathcal{O}_{\text{eager}}(x)</math> :</b>  if <math>x \notin \text{dom}(L)</math> then  if <math>x = 0^k</math> then <math>y \leftarrow \hat{y}</math> else <math>y \leftarrow \{0, 1\}^\ell</math>;  <math>L \leftarrow (x, y) :: L</math>  else <math>y \leftarrow L[x]</math>  return <math>y</math></p>
--	--

**Fig. 3.3.** An example of eager sampling justified by inter-procedural code motion.

$\mathcal{O} : \{0, 1\}^k \rightarrow \{0, 1\}^\ell$ . While in game  $G_{\text{lazy}}$  the oracle is implemented as a typical random oracle that chooses its answers on demand, in  $G_{\text{eager}}$  we use a fresh variable  $\hat{y}$  to fix in advance the response to a query of the form  $0^k$ . We can prove that both games are perfectly indistinguishable from the point of view of an adversary  $\mathcal{A}$  (who cannot write  $L$ ). Define

$$c \stackrel{\text{def}}{=} b \leftarrow \mathcal{A}() \quad S \stackrel{\text{def}}{=} \text{if } 0^k \notin \text{dom}(L) \text{ then } \hat{y} \leftarrow \{0, 1\}^\ell \text{ else } \hat{y} \leftarrow L[0^k]$$

and take  $I = \{\mathbf{L}\}$ ,  $X = \{\hat{y}\}$ . We introduce an intermediate game using rule [Trans],

$$\frac{\vdash \mathbf{G}_{\text{lazy}} \simeq_{\{b\}}^{\mathcal{V}} E_{\text{lazy}}, (\mathbf{L} \leftarrow \text{nil}; c; S) \quad \vdash E_{\text{lazy}}, (\mathbf{L} \leftarrow \text{nil}; c; S) \simeq_{\{b\}}^{\mathcal{V}} \mathbf{G}_{\text{eager}}}{\vdash \mathbf{G}_{\text{lazy}} \simeq_{\{b\}}^{\mathcal{V}} \mathbf{G}_{\text{eager}}} \text{[Trans]}$$

We prove the premise on the left by eliminating  $S$  as dead code, since it does not modify variable  $b$ . To prove the other premise, we introduce an intermediate game ( $E_{\text{eager}}, (\mathbf{L} \leftarrow \text{nil}; S; c)$ ). Its equivalence to  $\mathbf{G}_{\text{eager}}$  is direct by propagating the initial assignment to  $\mathbf{L}$  to the condition in  $S$  and then simplifying the conditional to its first branch. Its equivalence to ( $E_{\text{lazy}}, (\mathbf{L} \leftarrow \text{nil}; c; S)$ ) is justified by appealing to rule [S-Adv],

$$\frac{\frac{\vdash \mathbf{L} \leftarrow \text{nil} \equiv \mathbf{L} \leftarrow \text{nil}}{\vdash E_{\text{lazy}}, (\mathcal{O}_{\text{lazy}}; S) \equiv E_{\text{eager}}, (S; \mathcal{O}_{\text{eager}})} \text{[Refl]} \quad \frac{\vdash E_{\text{eager}}, (S; \mathcal{O}_{\text{eager}})}{\vdash E_{\text{eager}}, (S; c)} \text{[S-Adv]}}{\vdash E_{\text{lazy}}, (\mathbf{L} \leftarrow \text{nil}; c; S) \simeq_{\{b\}}^{\mathcal{V}} E_{\text{eager}}, (\mathbf{L} \leftarrow \text{nil}; S; c)} \text{[Seq]}$$

We are thus left to show

$$\vdash E_{\text{lazy}}, (\mathcal{O}_{\text{lazy}}.\text{body}; S) \equiv E_{\text{eager}}, (S; \mathcal{O}_{\text{eager}}.\text{body})$$

The proof of this latter judgment starts by an application of the generalized rule for conditionals of the logic for swapping statements. Let

$$\begin{aligned} e &= e' = x \notin \text{dom}(\mathbf{L}) \\ c_1 &= y \stackrel{\#}{\leftarrow} \{0, 1\}^\ell; \mathbf{L} \leftarrow (x, y) :: \mathbf{L} \\ c'_1 &= (\text{if } x = 0^k \text{ then } y \leftarrow \hat{y} \text{ else } y \stackrel{\#}{\leftarrow} \{0, 1\}^\ell); \mathbf{L} \leftarrow (x, y) :: \mathbf{L} \\ c_2 &= c'_2 = y \leftarrow \mathbf{L}[x] \end{aligned}$$

There are two non-trivial proof obligations:

$$1. \vdash c_2; S \sim S; c'_2 : =_{\mathcal{V}} \wedge (x \in \text{dom}(\mathbf{L})) \langle 1 \rangle \Rightarrow =_{\mathcal{V}}$$

This corresponds to showing that the code in the *else* branch in the conditional of each implementation of  $\mathcal{O}$  commutes with  $S$ , and follows from [S-Assn];

$$2. \vdash c_1; S \sim S; c'_1 : =_{\mathcal{V}} \wedge (x \notin \text{dom}(\mathbf{L})) \langle 1 \rangle \Rightarrow =_{\mathcal{V}}$$

By case analysis on  $x = 0^k$ :

- a) If  $x = 0^k$ , we can invoke certified program transformations—using the precondition that  $x \notin \text{dom}(\mathbf{L})$ —to simplify the goal to the following easily provable form:

$$\vdash y \stackrel{\#}{\leftarrow} \{0, 1\}^\ell; \mathbf{L} \leftarrow (x, y) :: \mathbf{L}; \hat{y} \leftarrow y \equiv \hat{y} \stackrel{\#}{\leftarrow} \{0, 1\}^\ell; y \leftarrow \hat{y}; \mathbf{L} \leftarrow (x, y) :: \mathbf{L}$$

- b) Otherwise, we do a further case analysis on  $0^k \in \text{dom}(\mathbf{L})$

- i. If  $0^k \in \text{dom}(\mathbf{L})$ , we have to prove that  $\vdash c_1; \hat{y} \leftarrow \mathbf{L}[0^k] \equiv \hat{y} \leftarrow \mathbf{L}[0^k]; c_1$  which is trivial;
- ii. Otherwise, the goal simplifies to  $\vdash c_1; \hat{y} \stackrel{\#}{\leftarrow} \{0, 1\}^\ell \equiv \hat{y} \stackrel{\#}{\leftarrow} \{0, 1\}^\ell; c_1$  which is also trivial.  $\square$

### 3.3 Reasoning about Failure Events

One common technique to justify a *lossy* transformation  $G, A \rightarrow G', A$ , where  $\Pr[G : A] \neq \Pr[G' : A]$  is based on what cryptographers call *failure events*. This technique relies on a *fundamental lemma* that allows to bound the difference in the probability of an event in two games: one identifies a failure event and argues that both games behave identically until failure occurs. One can then bound the difference in probability of another event by the probability of failure in either game. Consider for example the following two program snippets and their instrumented versions:

$$\begin{array}{ll} s \stackrel{\text{def}}{=} & \text{if } e \text{ then } c_1; c \text{ else } c_2 & s_{\mathbf{bad}} \stackrel{\text{def}}{=} & \text{if } e \text{ then } c_1; \mathbf{bad} \leftarrow \text{true}; c \text{ else } c_2 \\ s' \stackrel{\text{def}}{=} & \text{if } e \text{ then } c_1; c' \text{ else } c_2 & s'_{\mathbf{bad}} \stackrel{\text{def}}{=} & \text{if } e \text{ then } c_1; \mathbf{bad} \leftarrow \text{true}; c' \text{ else } c_2 \end{array}$$

If we ignore variable  $\mathbf{bad}$ ,  $s$  and  $s_{\mathbf{bad}}$ , and  $s'$  and  $s'_{\mathbf{bad}}$ , respectively, are observationally equivalent. Moreover,  $s_{\mathbf{bad}}$  and  $s'_{\mathbf{bad}}$  behave identically unless  $\mathbf{bad}$  is set. Thus, the difference of the probability of an event  $A$  in a game  $G$  containing the program fragment  $s$  and a game  $G'$  containing instead  $s'$  can be bounded by the probability of  $\mathbf{bad}$  being set in either  $G$  or  $G'$ .

**Lemma 3.3 (Fundamental Lemma).** *Let  $G_1, G_2$  be two games and let  $A, B$ , and  $F$  be events. If  $\Pr[G_1 : A \wedge \neg F] = \Pr[G_2 : B \wedge \neg F]$ , then*

$$|\Pr[G_1 : A] - \Pr[G_2 : B]| \leq \max(\Pr[G_1 : F], \Pr[G_2 : F])$$

*Proof.*

$$\begin{aligned} & |\Pr[G_1 : A] - \Pr[G_2 : B]| \\ &= |\Pr[G_1 : A \wedge F] + \Pr[G_1 : A \wedge \neg F] - \Pr[G_2 : B \wedge F] - \Pr[G_2 : B \wedge \neg F]| \\ &= |\Pr[G_1 : A \wedge F] - \Pr[G_2 : B \wedge F]| \\ &\leq \max(\Pr[G_1 : A \wedge F], \Pr[G_2 : B \wedge F]) \\ &\leq \max(\Pr[G_1 : F], \Pr[G_2 : F]) \end{aligned}$$

□

To apply this lemma, we developed a syntactic criterion to discharge its hypothesis for the case where  $A = B$  and  $F = \mathbf{bad}$ . The hypothesis can be automatically established by inspecting the code of both games: it holds if their code differs only after program points setting the flag  $\mathbf{bad}$  to true and  $\mathbf{bad}$  is never reset to false afterwards. Note also that if both games terminate with probability 1, then  $\Pr[G_1 : \mathbf{bad}] = \Pr[G_2 : \mathbf{bad}]$ , and that if, for instance, only game  $G_2$  terminates with probability 1, it must be the case that  $\Pr[G_1 : \mathbf{bad}] \leq \Pr[G_2 : \mathbf{bad}]$ .

#### 3.3.1 A Logic for Bounding the Probability of Events

Many steps in game-based proofs require to provide an upper bound for the measure of some function  $g$  after the execution of a command  $c$  (throughout this section,



we assume a fixed environment  $E$  that we omit from the presentation). This is typically the case when applying the Fundamental Lemma presented in the previous section: we need to bound the probability of the failure event **bad** (equivalently, the expected value of its characteristic function  $\mathbb{1}_{\text{bad}}$ ). A function  $f$  is an upper bound of  $(\lambda m. \llbracket c \rrbracket m g)$  when

$$\models \llbracket c \rrbracket g \preceq f \stackrel{\text{def}}{=} \forall m. \llbracket c \rrbracket m g \leq f m$$

Figure 3.4 gathers some rules for proving the validity of such triples. The rule for adversary calls assumes that  $f$  depends only on variables that the adversary cannot modify directly (but may modify indirectly through oracle calls, of course). The correctness of this rule is proved using the induction principle for well-formed adversaries together with the rest of the rules of the logic.

---


$$\begin{array}{c}
\vdash \llbracket \text{skip} \rrbracket f \preceq f \quad \frac{f = \lambda m. g(m\{\llbracket e \rrbracket m/x\})}{\vdash \llbracket x \leftarrow e \rrbracket g \preceq f} \quad \frac{f = \lambda m. \llbracket d \rrbracket m (\lambda v. g(m\{v/x\}))}{\vdash \llbracket x \leftarrow d \rrbracket g \preceq f} \\
\frac{\vdash \llbracket c_1 \rrbracket g \preceq f \quad \llbracket c_2 \rrbracket h \preceq g}{\vdash \llbracket c_1; c_2 \rrbracket h \preceq f} \quad \frac{\vdash \llbracket c_1 \rrbracket g \preceq f \quad \llbracket c_2 \rrbracket g \preceq f}{\vdash \llbracket \text{if } e \text{ then } c_1 \text{ else } c_2 \rrbracket g \preceq f} \quad \frac{\vdash \llbracket c \rrbracket f \preceq f}{\vdash \llbracket \text{while } e \text{ do } c \rrbracket f \preceq f} \\
\frac{\vdash g \leq g' \quad \llbracket c \rrbracket g' \preceq f' \quad f' \leq f}{\vdash \llbracket c \rrbracket g \preceq f} \quad \frac{\vdash \llbracket p.\text{body} \rrbracket g \preceq f \quad f =_X f \quad g =_Y g \quad x \notin (X \cup Y)}{\vdash \llbracket x \leftarrow p(\mathbf{e}) \rrbracket g \preceq f} \\
\frac{\vdash_{\text{wf}} \mathcal{A} \quad \forall p \in \mathcal{O}. \vdash \llbracket p.\text{body} \rrbracket f \preceq f \quad f =_X f \quad X \cap (\{x\} \cup \mathcal{RW}) = \emptyset}{\vdash \llbracket x \leftarrow \mathcal{A}(\mathbf{e}) \rrbracket f \preceq f} \\
\frac{f =_I f \quad \vdash c \simeq_O^I c' \quad g =_O g \quad \vdash \llbracket c' \rrbracket g \preceq f}{\vdash \llbracket c \rrbracket g \preceq f}
\end{array}$$


---

**Fig. 3.4.** Selected rules of a logic for bounding the probability of events.

The rules bear some similarity with the rules of (standard) Hoare logic. However, there are some subtle differences. For example, the premises of the rules for branching statements do not consider guards. The rule

$$\frac{\vdash \llbracket c_1 \rrbracket g \preceq f|_e \quad \vdash \llbracket c_2 \rrbracket g \preceq f|_{\neg e}}{\vdash \llbracket \text{if } e \text{ then } c_1 \text{ else } c_2 \rrbracket g \preceq f}$$

where  $f|_e$  is defined as  $(\lambda m. \text{if } \llbracket e \rrbracket m \text{ then } f(m) \text{ else } 0)$  can be derived from the rule for conditionals in the figure by two simple applications of the “rule of consequence”. Moreover, the rule for conditional statements (and its variant above) is incomplete: consider a statement of the form  $\llbracket \text{if true then } c_1 \text{ else } c_2 \rrbracket g \preceq f$  such that  $\llbracket c_1 \rrbracket g \preceq f$  is valid, but not  $\llbracket c_2 \rrbracket g \preceq f$ ; the triple  $\llbracket \text{if true then } c_1 \text{ else } c_2 \rrbracket g \preceq f$  is valid, but to

derive it one needs to resort to observational equivalence. More general rules exist, but we have not formalized them since we did not need them in our proofs.<sup>4</sup>

### Digression

The differences between the above triples and those of Hoare logic are inherent to their definition, which is tailored to establish upper bounds for the probability of events. Nevertheless, the validity of a Hoare triple  $\{P\} c \{Q\}$  (in which pre- and post-conditions are Boolean-valued predicates) is equivalent to the validity of the triple  $\llbracket c \rrbracket \mathbb{1}_{\neg Q} \preceq \mathbb{1}_{\neg P}$ . We can consider dual triples of the form  $\llbracket c \rrbracket g \succeq f$  whose validity is defined as:

$$\vDash \llbracket c \rrbracket g \succeq f \stackrel{\text{def}}{=} \forall m. \llbracket c \rrbracket m \ g \geq f \ m$$

This allows to express termination of a program as  $\llbracket c \rrbracket \mathbb{1} \succeq \mathbb{1}$  and admits an embedding of Hoare triples, mapping  $\{P\} c \{Q\}$  to  $\llbracket c \rrbracket \mathbb{1}_Q \succeq \mathbb{1}_P$ . However, this embedding does not preserve validity for non-terminating programs under the partial correctness interpretation. Consider a program  $c$  that never terminates: we have  $\{\text{true}\} c \{\text{false}\}$ , but clearly not  $\llbracket c \rrbracket \mathbb{1}_{\text{false}} \succeq \mathbb{1}$ .

### 3.3.2 Automation

In most applications of Lemma 3.3, failure can only be triggered by oracle calls. Typically, the flag **bad** that signals failure is set in the code of an oracle for which an upper bound for the number of queries made by the adversary is known. The following lemma provides a general method for bounding the probability of failure under such circumstances.

**Lemma 3.4 (Failure Event Lemma).** *Consider an event  $F$  and a game  $G$  that gives adversaries access to an oracle  $\mathcal{O}$ . Let  $\text{cntr} : \mathcal{E}_{\text{Nat}}, h : \mathbb{N} \rightarrow [0, 1]$  be such that  $\text{cntr}$  and  $F$  do not depend on variables that can be written outside  $\mathcal{O}$ , and for any initial memory  $m$ ,*

$$\neg F(m) \implies \Pr[\mathcal{O}.\text{body}, m : F] \leq h(\llbracket \text{cntr} \rrbracket m)$$

and

$$\begin{aligned} & \text{range}(\llbracket \mathcal{O}.\text{body} \rrbracket m) (\lambda m'. \llbracket \text{cntr} \rrbracket m < \llbracket \text{cntr} \rrbracket m') \vee \\ & \text{range}(\llbracket \mathcal{O}.\text{body} \rrbracket m) (\lambda m'. \llbracket \text{cntr} \rrbracket m = \llbracket \text{cntr} \rrbracket m' \wedge F \ m' = F \ m) \end{aligned}$$

Then, for any initial memory  $m$  satisfying  $\neg F(m)$  and  $\llbracket \text{cntr} \rrbracket m = 0$ ,

$$\Pr[G, m : F \wedge \text{cntr} \leq q] \leq \sum_{i=0}^{q-1} h(i)$$

<sup>4</sup> More generally, it seems possible to make the logic complete, at the cost of considering more complex statements with pre-conditions on memories.

*Proof.* Define  $f : \mathcal{M} \rightarrow [0, 1]$  as follows

$$f(m) \stackrel{\text{def}}{=} \begin{cases} 0 & \text{if } \llbracket \text{cntr} \rrbracket m > q \\ \mathbb{1}_F(m) + \mathbb{1}_{\neg F}(m) \sum_{i=\llbracket \text{cntr} \rrbracket m}^{q-1} h(i) & \text{if } \llbracket \text{cntr} \rrbracket m \leq q \end{cases}$$

We show  $\llbracket G \rrbracket f \preceq f$  by structural induction on the code of  $G$  using the rules of the logic presented in the previous section. We first prove that  $\mathcal{O}$  satisfies the triple  $\llbracket \mathcal{O}.\text{body} \rrbracket f \preceq f$ . We must show that for every  $m$ ,  $\llbracket \mathcal{O}.\text{body} \rrbracket m f \leq f(m)$ . This is trivial when  $\text{cntr}$  is not incremented, because we have

$$\llbracket \mathcal{O}.\text{body} \rrbracket m f = f(m) \quad (\llbracket \mathcal{O}.\text{body} \rrbracket m \mathbb{1} \leq f(m))$$

When  $\text{cntr}$  is incremented and  $\llbracket \text{cntr} \rrbracket m \geq q$ , this is trivial too, because the left hand side becomes 0. We are left with the case where  $\mathcal{O}.\text{body}$  increments  $\text{cntr}$  and  $\llbracket \text{cntr} \rrbracket m < q$ . If  $F(m)$ , the right hand side is equal to 1 and the inequality holds. Otherwise, we have from the hypotheses that

$$\begin{aligned} \llbracket \mathcal{O}.\text{body} \rrbracket m f &\leq \llbracket \mathcal{O}.\text{body} \rrbracket m \left( \lambda m'. \mathbb{1}_F(m') + \mathbb{1}_{\neg F}(m') \sum_{i=\llbracket \text{cntr} \rrbracket m'}^{q-1} h(i) \right) \\ &\leq \Pr[\mathcal{O}.\text{body}, m : F] + \Pr[\mathcal{O}.\text{body}, m : \neg F] \sum_{i=\llbracket \text{cntr} \rrbracket m+1}^{q-1} h(i) \\ &\leq h(\llbracket \text{cntr} \rrbracket m) + \sum_{i=\llbracket \text{cntr} \rrbracket m+1}^{q-1} h(i) \\ &= \sum_{i=\llbracket \text{cntr} \rrbracket m}^{q-1} h(i) = f(m) \end{aligned}$$

Using the rules in Figure 3.4, we can then extend this result to adversary calls and to the rest of the game, showing that  $\llbracket G \rrbracket f \preceq f$ .

Finally, let  $m$  be a memory such that  $\neg F(m)$  and  $\llbracket \text{cntr} \rrbracket m = 0$ . It follows immediately from  $\llbracket G \rrbracket f \preceq f$  that

$$\Pr[G, m : F \wedge \text{cntr} \leq q] \leq \llbracket G \rrbracket m f \leq f(m) = \sum_{i=0}^{q-1} h(i) \quad \square$$

When failure is defined as the probability of a flag **bad** being set by an oracle and the number of queries the adversary makes to this oracle is upper bounded by  $q$ , the above lemma can be used to bound the probability of failure by taking  $F = \mathbf{bad}$  and defining  $h$  suitably. In most practical applications the probability of an oracle call raising failure is history-independent and hence  $h$  is a constant function. The proof of Lemma 4.3 given in Section 4.3.2 is an exception for which the full generality of the lemma is needed.

# 4

## The PRP/PRF Switching Lemma

CRYPTOGRAPHIC systems are generally built incrementally by combining basic primitives with the goal of achieving a higher level security goal. Rather than designing a system for a particular choice of a primitive, one designs the system assuming a generic and simplified model of the primitive. The security of the whole system is then analyzed under the assumption that this model behaves in an ideal way. Since in practice the construction that implements the primitive will definitely deviate from this ideal behavior, the actual security of the system depends on how wide the gap between the idealized and the actual behavior is. Pseudorandom functions (PRF) and pseudorandom permutations (PRP) are two idealized primitives that are used to model blockciphers and thus play a central role in the design of symmetric-key systems. Although the most natural assumption to make about a blockcipher is that it behaves as a pseudorandom permutation, most commonly the security of a system based on a blockcipher is analyzed by replacing the blockcipher with a perfectly random function. The PRP/PRF switching lemma is used to fill the gap: given a bound for the security of a blockcipher as a pseudorandom permutation, it gives a bound for its security as a pseudorandom function.

In this Chapter we will formally define the notions of pseudorandom function and pseudorandom permutation and their security, and we will overview two different game-based proofs of the PRP/PRF switching lemma. Both use the Fundamental Lemma of game-playing (Lemma 3.3) to bound the advantage of an adversary by the probability of a failure event, but each proof bounds the probability of failure using a different technique. We first present a proof that uses the principle of eager sampling so that all random choices are done up front and the probability is directly computable. We then present a significantly more compact proof that uses Lemma 3.4 (see §3.3.2) to bound the probability of failure.

## 4.1 Pseudorandom Functions

A pseudorandom function is a key-indexed family of functions  $\{f_k \mid k \in K\}$  with the property that an instance selected at random according to some distribution on  $K$  is computationally indistinguishable from a perfectly random function. Unless otherwise said we will consider that the distribution on  $K$  that makes this property hold is the uniform distribution.

Consider an adversary who has only blackbox access to an oracle and is put in either of two scenarios: one where the oracle is a random instance of a function drawn from a family of pseudorandom functions, and other where the oracle is a perfectly random function. This adversary should only be able to tell apart both scenarios with a small probability. We can define this formally using games.

**Definition 4.1 (PRF-advantage).** Let  $\{f_k : A \rightarrow B \mid k \in K\}$  be a pseudorandom function family, and  $\mathcal{A}$  an adversary with blackbox access to an oracle  $\mathcal{O}$  as in the following two games:

**Game  $G_{\text{PRF}}$  :**  
 $k \xleftarrow{\$} K; b \leftarrow \mathcal{A}()$   
**Oracle  $\mathcal{O}(x)$  :**  
 return  $f_k(x)$

**Game  $G_{\text{RF}}$  :**  
 $L \leftarrow \text{nil}; b \leftarrow \mathcal{A}()$   
**Oracle  $\mathcal{O}(x)$  :**  
 if  $x \notin \text{dom}(L)$  then  
 $y \xleftarrow{\$} B;$   
 $L \leftarrow (x, y) :: L$   
 return  $L[x]$

The PRF-advantage of  $\mathcal{A}$  against  $f$  is defined as

$$\text{Adv}_{\text{PRF}_f}^{\mathcal{A}} \stackrel{\text{def}}{=} |\Pr[G_{\text{PRF}} : b = 1] - \Pr[G_{\text{RF}} : b = 1]|$$

The concept of pseudorandom function was first introduced by Goldreich, Goldwasser, and Micali [1986]. Rather than considering a single family of key-indexed functions, they consider a collection of families parametrized by a security parameter  $\eta$ . In this asymptotic setting, a pseudorandom function is secure if all adversaries that execute in polynomial-time on  $\eta$  have a negligible PRF-advantage (as a function of  $\eta$ ). In contrast, in the setting of exact security there is no absolute notion of security for pseudorandom functions. The above definition only associates to each adversary  $\mathcal{A}$  a real number, its PRF-advantage. In practice one considers all adversaries consuming no more than a certain amount of computational resources, and gives an upper bound for their PRF-advantage.

## 4.2 Pseudorandom Permutations

A pseudorandom permutation is key-indexed family of permutations  $\{f_k \mid k \in K\}$  on  $A$  such that a permutation randomly drawn from the family is computationally indistinguishable from a permutation drawn uniformly from the set of all permutations on  $A$ . Again, we define this notion formally using games.

**Definition 4.2 (PRP-advantage).** Let  $\{f_k : A \rightarrow A \mid k \in K\}$  be a pseudorandom permutation family, and  $\mathcal{A}$  an adversary with blackbox access to an oracle  $\mathcal{O}$  as in the following two games:

<p><b>Game <math>G_{\text{PRP}}</math> :</b>  <math>k \xleftarrow{\\$} K; b \leftarrow \mathcal{A}()</math></p> <p><b>Oracle <math>\mathcal{O}(x)</math> :</b>  return <math>f_k(x)</math></p>	<p><b>Game <math>G_{\text{RP}}</math> :</b>  <math>L \leftarrow \text{nil}; b \leftarrow \mathcal{A}()</math></p> <p><b>Oracle <math>\mathcal{O}(x)</math> :</b>  if <math>x \notin \text{dom}(L)</math> then  <math>y \xleftarrow{\\$} A \setminus \text{ran}(L);</math>  <math>L \leftarrow (x, y) :: L</math>  return <math>L[x]</math></p>
--	--

where the instruction  $y \xleftarrow{\$} A \setminus \text{ran}(L)$  samples uniformly an element of  $A$  that is not in the range of the association list  $L$ , thus ensuring that oracle  $\mathcal{O}$  in  $G_{\text{RP}}$  implements an injective—and therefore bijective—function. The PRP-advantage of  $\mathcal{A}$  against  $f$  is defined as

$$\text{Adv}_{\text{PRP}_f}^{\mathcal{A}} \stackrel{\text{def}}{=} |\Pr[G_{\text{PRP}} : b = 1] - \Pr[G_{\text{RP}} : b = 1]|$$

The notion of pseudorandom permutation is due to Luby and Rackoff [1988], who also observe that the notions of pseudorandom function and permutation are no different in an asymptotic setting, and show how to construct a pseudorandom permutation from a pseudorandom function.

### 4.3 The PRP/PRF Switching Lemma

We already observed that every pseudorandom permutation family is also a pseudorandom function family. But how well does a pseudorandom permutation perform as a pseudorandom function? Let us first consider the simpler problem of comparing a perfectly random function to a random permutation. Suppose you give to an adversary blackbox access to an oracle implementing either a random function or a random permutation, and you ask it to tell you which is the case. For the sake of concreteness let us assume the domain of the permutation (and the domain and range of the function) is  $\{0, 1\}^\ell$ . Due to the birthday problem, no matter the strategy the adversary follows, after roughly  $2^{\ell/2}$  queries to the oracle it will be able to tell in which scenario it is with a high probability. If the oracle is a random function, a collision is almost sure to occur, whereas it could not occur when the oracle is a random permutation. The birthday problem gives us a lower bound for the advantage of an adversary in distinguishing a random function from a random permutation. The following lemma gives an upper bound.

**Lemma 4.3 (PRP/PRF switching lemma).** Let  $\mathcal{A}$  be an adversary with blackbox access to an oracle  $\mathcal{O}$  implementing either a random permutation on  $\{0, 1\}^\ell$  as in game  $G_{\text{RP}}$  or a random function from  $\{0, 1\}^\ell$  to  $\{0, 1\}^\ell$  as in game  $G_{\text{RF}}$ . Suppose, in addition, that  $\mathcal{A}$  makes at most  $q > 0$  queries to oracle  $\mathcal{O}$ . Then,

$$|\Pr[\mathbf{G}_{\text{RP}} : b = 1] - \Pr[\mathbf{G}_{\text{RF}} : b = 1]| \leq \frac{q(q-1)}{2^{\ell+1}} \quad (4.1)$$

We overview two different machine-checked proofs of the PRP/PRF switching lemma that exploit the code-based techniques presented in earlier sections. Both proofs use the Fundamental Lemma to bound the advantage of the adversary by the probability of a failure event. The first proof uses the eager sampling technique of Section 3.2.3 to bound the probability of failure, whereas the second one relies on Lemma 3.4 of Section 3.3.2. We begin by introducing in Figure 4.1 annotated versions  $\mathbf{G}_{\text{RP}}^{\text{bad}}$  and  $\mathbf{G}_{\text{RF}}^{\text{bad}}$  of the games  $\mathbf{G}_{\text{RP}}$  and  $\mathbf{G}_{\text{RF}}$ . These annotated games set a flag **bad** whenever the oracle corresponding to a random function would return a value colliding with a response to a previous query, but are otherwise semantically equivalent to the original games. The annotated games are syntactically identical until the point where **bad** is set, so we can appeal to Lemma 3.3 to bound the difference in the probability of  $b$  being equal to 1 in the original games:

$$|\Pr[\mathbf{G}_{\text{RP}} : b = 1] - \Pr[\mathbf{G}_{\text{RF}} : b = 1]| \leq \Pr[\mathbf{G}_{\text{RF}}^{\text{bad}} : \text{bad}]$$

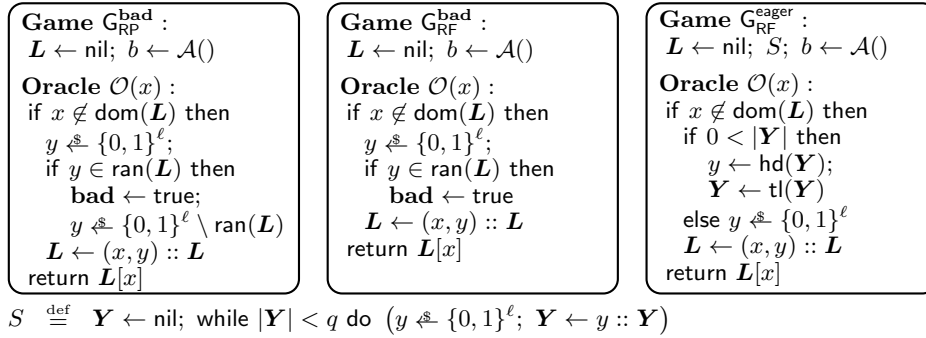


Fig. 4.1. Games used in the proofs of the PRP/PRF Switching Lemma.

### 4.3.1 A Proof Based on Eager Sampling

We make a first remark: the probability of **bad** being set in game  $\mathbf{G}_{\text{RF}}^{\text{bad}}$  is bounded by the probability of having a collision in  $\text{ran}(L)$  at the end of the game. Let us write this latter event as  $\text{col}(L)$ ,

$$\text{col}(L) \stackrel{\text{def}}{=} \exists x_1, x_2 \in \text{dom}(L). x_1 \neq x_2 \wedge L[x_1] = L[x_2]$$

We prove that

$$\vdash G_{\text{RF}}^{\text{bad}} \sim G_{\text{RF}} : \text{true} \Rightarrow \text{bad}(1) \implies \text{col}(\mathbf{L}) \langle 2 \rangle$$

Thus,

$$\Pr [G_{\text{RF}}^{\text{bad}} : \text{bad}] \leq \Pr [G_{\text{RF}} : \text{col}(\mathbf{L})] \quad (4.2)$$

Using the logic for swapping statements, we then modify the oracle in  $G_{\text{RF}}$  so that the responses to the first  $q$  queries are instead chosen at the beginning of the game and stored in a list  $\mathbf{Y}$ , thus obtaining the equivalent eager version  $G_{\text{RF}}^{\text{eager}}$  shown in Figure 4.1; each time a query is made, the oracle pops a value from list  $\mathbf{Y}$  and gives it back to the adversary as the response. Since the initialization code  $S$  terminates and does not modify  $\mathbf{L}$ , we can conclude that

$$\Pr [G_{\text{RF}} : \text{col}(\mathbf{L})] = \Pr [G_{\text{RF}}; S : \text{col}(\mathbf{L})] = \Pr [G_{\text{RF}}^{\text{eager}} : \text{col}(\mathbf{L})]$$

We prove using the relational Hoare logic that having a collision in the range of  $\mathbf{L}$  at the end of this last game is bounded by the probability of having a collision in  $\mathbf{Y}$  immediately after executing  $S$ . We conclude that the bound in (4.1) holds by analyzing the loop in  $S$ .

Observe that if there are no collisions in  $\mathbf{Y}$  in a memory  $m$ , we can prove by induction on  $(q - |\mathbf{Y}|)$  that the probability of sampling a colliding value in the remaining loop iterations is

$$\Pr [S, m : \exists i, j \in \mathbb{N}. i < j < q \wedge \mathbf{Y}[i] = \mathbf{Y}[j]] = \sum_{i=|\mathbf{Y}|}^{q-1} \frac{i}{2^\ell}$$

We thus have,

$$\begin{aligned} \Pr [G_{\text{RF}}^{\text{eager}} : \text{col}(\mathbf{L})] &\leq \Pr [S, m\{\text{nil}/\mathbf{Y}\} : \exists i, j \in \mathbb{N}. i < j < q \wedge \mathbf{Y}[i] = \mathbf{Y}[j]] \\ &= \frac{q(q-1)}{2^{\ell+1}} \end{aligned}$$

□

### 4.3.2 A Proof Based on the Failure Event Lemma

The bound in (4.1) follows from a direct application of Lemma 3.4. It suffices to take  $F = \text{bad}$ ,  $h(i) = i 2^{-\ell}$ , and  $\text{cntr} = |\mathbf{L}|$ . If  $\text{bad}$  is initially set to false in memory  $m$ , we have

$$\Pr [G_{\text{RF}}^{\text{bad}}, m : \text{bad}] = \Pr [b \leftarrow \mathcal{A}(), m\{\text{nil}/\mathbf{L}\} : \text{bad} \wedge |\mathbf{L}| \leq q] \leq \sum_{i=0}^{q-1} h(i) = \frac{q(q-1)}{2^{\ell+1}}$$

The first equation holds because  $\mathcal{A}$  does not make more than  $q$  queries to  $\mathcal{O}$ . The inequality is obtained from Lemma 3.4; we use the logic of Section 3.3.1 to bound the probability of  $\text{bad}$  being set in one call to the oracle by  $h(\text{cntr})$ .



### 4.3.3 Comparison of Both Proofs

The proof of the PRP/PRF switching lemma that bounds the probability of failure using Lemma 3.4 presented in Section 4.3.2 is considerably shorter compared to the one presented in Section 4.3.1, that uses the principle of eager sampling to reduce the problem of bounding the probability of failure to local reasoning about a loop. The former proof takes just about 100 lines of Coq, compared to the 400 lines that takes the latter. Both proofs are significantly more compact than the 900-lines proof reported in Barthe et al. [2009c]. That proof used an earlier mechanization of the eager sampling technique that only allowed to fix the value of one response of the oracle at a time. Thus, in order to fix in advance the response to all the  $q$  queries that could be made by the adversary, an induction argument was necessary.

## 4.4 Pseudorandom Permutations as Pseudorandom Functions

In view of Lemma 4.3, we can now answer our original question: how well does a pseudorandom permutation perform as a pseudorandom function?

Let  $\{f_k \mid k \in K\}$  be a pseudorandom permutation family on  $\{0, 1\}^\ell$  and let  $\mathcal{A}$  be an adversary that makes at most  $q > 0$  queries to its oracle. The PRF-advantage of  $\mathcal{A}$  is

$$\begin{aligned} \text{Adv}_{\text{PRF}_f}^{\mathcal{A}} &= |\Pr[\text{G}_{\text{PRF}} : b = 1] - \Pr[\text{G}_{\text{RF}} : b = 1]| \\ &= |\Pr[\text{G}_{\text{PRP}} : b = 1] - \Pr[\text{G}_{\text{RP}} : b = 1] + \Pr[\text{G}_{\text{RP}} : b = 1] - \Pr[\text{G}_{\text{RF}} : b = 1]| \\ &\leq |\Pr[\text{G}_{\text{PRP}} : b = 1] - \Pr[\text{G}_{\text{RP}} : b = 1]| + |\Pr[\text{G}_{\text{RP}} : b = 1] - \Pr[\text{G}_{\text{RF}} : b = 1]| \\ &\leq \text{Adv}_{\text{PRP}_f}^{\mathcal{A}} + \frac{q(q-1)}{2^{\ell+1}} \end{aligned}$$

Consider this bound in an asymptotic setting, and assume  $\ell$  is linearly proportional to the security parameter  $\eta$ . Every polynomial-time adversary can make only a polynomial number of queries to its oracle, so  $q$  is polynomial on the security parameter, and the term on the right hand side of the last inequality is negligible on  $\eta$  if  $f$  is secure as a pseudorandom permutation. Hence,  $f$  is secure as a pseudorandom function whenever  $f$  is secure as a pseudorandom permutation.

We already observed that the construction of Luby and Rackoff [1988] provides a means to build a pseudorandom permutation from a pseudorandom function; many other authors have studied variations of this construction. In contrast, the reverse direction has been historically much less studied. Although it follows from the PRP/PRF switching lemma that in a complexity-theoretical setting, a pseudorandom permutation *is* a pseudorandom function, there are constructions that achieve better security at a low efficiency cost [Hall et al. 1998].

## 4.5 Discussion

Despite the apparent simplicity of the PRP/PRF switching lemma, some purported proofs in the literature contain a subtle error in reasoning about conditional probabilities (cf. Impagliazzo and Rudich [1989]). Let us briefly report the argument in those proofs.

*Intuitive proof.*

Let *collision* be the event that adversary  $\mathcal{A}$  gets the same answer to two different queries when interacting with a random function. Since a random permutation behaves the same as a random function when no collisions are observed, we have that

$$\Pr[\mathbf{G}_{\text{RP}} : b = 1] = \Pr[\mathbf{G}_{\text{RF}} : b = 1 \mid \neg\text{collision}] \quad (4.3)$$

Let  $x = \Pr[\mathbf{G}_{\text{RF}} : b = 1 \mid \neg\text{collision}]$ ,  $y = \Pr[\mathbf{G}_{\text{RF}} : b = 1 \mid \text{collision}]$ . Then,

$$\begin{aligned} & |\Pr[\mathbf{G}_{\text{RP}} : b = 1] - \Pr[\mathbf{G}_{\text{RF}} : b = 1]| \\ &= |x - (x \Pr[\mathbf{G}_{\text{RF}} : \neg\text{collision}] + y \Pr[\mathbf{G}_{\text{RF}} : \text{collision}])| \\ &= |x(1 - \Pr[\mathbf{G}_{\text{RF}} : \neg\text{collision}]) - y \Pr[\mathbf{G}_{\text{RF}} : \text{collision}]| \\ &= |x - y| \Pr[\mathbf{G}_{\text{RF}} : \text{collision}] \end{aligned}$$

Since  $0 \leq x, y \leq 1$ ,

$$|\Pr[\mathbf{G}_{\text{RP}} : b = 1] - \Pr[\mathbf{G}_{\text{RF}} : b = 1]| \leq \Pr[\mathbf{G}_{\text{RF}} : \text{collision}]$$

Since the adversary makes at most  $q$  queries to the oracle,

$$\Pr[\mathbf{G}_{\text{RF}} : \text{collision}] \leq \frac{q(q-1)}{2^{\ell+1}}$$


and the bound (4.1) follows.  $\square$

The reader may be wondering where is the error in the above argument. The problem is that equation (4.3) might not hold, no matter how appealing the intuitive justification we gave can look. We can actually illustrate this with a counterexample for  $\ell = 1$ , by showing a particular adversary for which the equation does not hold:

**Adversary  $\mathcal{A}_1()$  :**  
 $y \leftarrow \mathcal{O}(0)$ ;  
 if  $y = 0$  then return 1  
 else  
 $y \leftarrow \mathcal{O}(1)$ ;  
 if  $y = 1$  then return 1 else return 0

Let us analyze how this adversary fares in each game with respect to equation (4.3). We can better depict the behavior of the adversary using a tree. The values

in leaves represent the bit  $b$  the adversary returns in each case whereas the label on the left of a node indicates the probability of reaching it.




It follows that

$$\frac{1}{2} = \Pr[\text{G}_{\text{RP}} : b = 1] \neq \Pr[\text{G}_{\text{RF}} : b = 1 \mid \neg\text{collision}] = \frac{2}{3} \quad (4.4)$$

The reason of the discrepancy is that the number of queries made by the adversary varies depending on the answer it receives from its first query.

The same would happen if the number of queries depended on the internal random choices of  $\mathcal{A}$ . For instance, the following adversary makes either zero or two queries depending on the result of sampling a fair coin and achieves the same probabilities in equation (4.4) as the adversary we showed previously:

**Adversary  $\mathcal{A}_2()$  :**  
 $a \xleftarrow{\$} \{0, 1\}$ ;  
 if  $a = 0$  then return 1  
 else  
 $y_0 \leftarrow \mathcal{O}(0)$ ;  $y_1 \leftarrow \mathcal{O}(1)$ ;  
 if  $y_0 \wedge y_1$  then return 1 else return 0



Obviously, the result we proved in the previous section holds for both of the above adversaries. We can compute the actual advantage they achieve and check that the bound in Equation (4.1) holds:

$$\begin{aligned} \left| \Pr \left[ \mathcal{G}_{\text{RP}}^{\mathcal{A}_1} : b = 1 \right] - \Pr \left[ \mathcal{G}_{\text{RF}}^{\mathcal{A}_1} : b = 1 \right] \right| &= \left| \frac{1}{2} - \frac{3}{4} \right| = \frac{1}{4} \leq \frac{q(q-1)}{2^{\ell+1}} = \frac{1}{2} \\ \left| \Pr \left[ \mathcal{G}_{\text{RP}}^{\mathcal{A}_2} : b = 1 \right] - \Pr \left[ \mathcal{G}_{\text{RF}}^{\mathcal{A}_2} : b = 1 \right] \right| &= \left| \frac{1}{2} - \frac{5}{8} \right| = \frac{1}{8} \leq \frac{q(q-1)}{2^{\ell+1}} = \frac{1}{2} \end{aligned}$$

Indeed, both adversaries make at most  $q = 2$  queries and the bound holds. The following adversary  $\mathcal{A}_3$  mounting a birthday attack with  $q = 2$  achieves the maximal value of  $1/2$ :

**Adversary  $\mathcal{A}_3()$  :**

$y_1 \leftarrow \mathcal{O}(0)$ ;

$y_2 \leftarrow \mathcal{O}(1)$ ;

if  $y_1 \neq y_2$  then return 1 else return 0

$$\left| \Pr \left[ \mathcal{G}_{\text{RP}}^{\mathcal{A}_3} : b = 1 \right] - \Pr \left[ \mathcal{G}_{\text{RF}}^{\mathcal{A}_3} : b = 1 \right] \right| = \left| 1 - \frac{1}{2} \right| = \frac{1}{2}$$

## 4.6 Related Work

The standard *proof* of the PRP/PRF switching lemma is due to Impagliazzo and Rudich [1989, Theorem 5.1]. The above observation about the error in the reasoning in the standard proof of the lemma and the first counterexample we showed are due to Bellare and Rogaway [2006].

Bellare and Rogaway [2006, Lemma 1] give a game-based proof of the PRP/PRF switching lemma. Their proof is similar to ours, but they make the additional assumption that the adversary never asks an oracle query twice. Just as in the proof we presented, they use the Fundamental Lemma of game-playing to bound the difference in the probability of the adversary outputting 1 when interacting with either a random permutation or a random function. However, the justification of the bound on the probability of the **bad** flag being set when the adversary interacts with a random function remains informal. The same authors give also a proof of the PRP/PRF switching lemma that does not use games, under the assumption that the adversary is deterministic and makes exactly  $q$  different queries to its oracle. We note that the assumption of the adversary being deterministic is without loss of generality only if it is computationally unbounded, and therefore the argument does not hold in an asymptotic setting where the adversary must execute in polynomial-time.

Shoup [2004, Section 5.1] gives a game-based proof of the PRP/PRF switching lemma under the assumption that the adversary makes exactly  $q$  distinct queries to its oracle. In the games he considers, the challenger acts as intermediary between the oracle and the adversary. Rather than the adversary calling the oracle at its

discretion, it is the challenger that calls the adversary to get the next query; it then asks the query itself to the oracle and gives the response back the adversary in the next call. There is probably nothing wrong with this formulation, but we feel that it imposes unnecessary restrictions on the form of the adversary that do not help to make the proof more clear.

Affeldt, Tanaka, and Marti [2007] present a formalization of a game-based proof of the PRP/PRF switching lemma in the Coq proof assistant. What they prove in reality is a simplified variant that only holds for non-adaptive and deterministic adversaries. They formalize adversaries as purely deterministic mathematical functions that take a natural number and return an element in the domain of its oracle (a query). This implies that the queries the adversary makes do not depend on the responses to previous queries or on any random choices. For instance, the two adversaries we gave in the previous section as counterexamples to the probabilistic reasoning in the original proof of the PRP/PRF switching lemma are ruled out by this formulation.

# 5

## Unforgeability of Full-Domain Hash Signatures

IN this chapter we will go through the formalization in CertiCrypt of two different proofs of security of the Full-Domain Hash (FDH) signature scheme. The FDH scheme was first proposed by Bellare and Rogaway [1996] as an efficient RSA-based signature scheme, but is in fact an instance of an earlier construction described by the same authors in 1993. Here, we will consider this latter, more general construction, which is based on a family of one-way trapdoor permutations  $f$  on a cyclic group  $G$ , and a hash function  $H : \{0, 1\}^* \rightarrow G$  whose range is the full domain of  $f$ . The RSA-based scheme is obtained by instantiating  $f$  with the RSA function, and the hash function with some cryptographic hash function, such as SHA-1 with the length of its output extended to match that of the RSA modulus.

**Definition 5.1 (Trapdoor permutation).** *A family of trapdoor permutations is a triple of algorithms  $(\mathcal{KG}, f, f^{-1})$ . For a given value of the security parameter  $\eta$ , the key generator  $\mathcal{KG}(\eta)$  randomly selects a pair of keys  $(pk, sk)$  such that  $f(pk, \cdot)$  is a permutation on its domain and  $f^{-1}(sk, \cdot)$  is its inverse. We say that a family of trapdoor permutations is one-way if it cannot be inverted in probabilistic polynomial-time on a uniformly distributed element in its domain.*

**Definition 5.2 (Full-Domain Hash signature scheme).** *Let  $(\mathcal{KG}_f, f, f^{-1})$  be a family of trapdoor permutations on cyclic groups  $G_\eta$  and let  $H$  be family of hash functions from bitstrings of arbitrary length onto the domain of the permutations. The Full-Domain Hash digital signature scheme is composed of the following triple of algorithms:*

$$\begin{aligned} \mathcal{KG}(\eta) &\stackrel{\text{def}}{=} (pk, sk) \leftarrow \mathcal{KG}_f(\eta); \text{ return } (pk, sk) \\ \text{Sign}(sk, m) &\stackrel{\text{def}}{=} \text{return } f^{-1}(sk, H(m)) \\ \text{Verify}(pk, m, \sigma) &\stackrel{\text{def}}{=} \text{return } (f(pk, \sigma) = H(m)) \end{aligned}$$

*The key generation algorithm just runs the key generation algorithm of the underlying trapdoor permutation obtaining a public key  $pk$ , that is used as the verification*

key of the scheme, and a secret key (trapdoor)  $sk$  used to compute signatures. The signature of a message  $m \in \{0, 1\}^*$  is simply  $f^{-1}(sk, H(m))$ , the preimage under  $f$  of its digest. To verify a purported signature  $\sigma$  on a message  $m$ , it suffices to check whether  $H(m)$  and  $f(pk, \sigma)$  coincide.

The FDH scheme can be proved secure in the random oracle model against existential forgery under adaptive chosen-message attacks (see Definition 2.10 in §2.5). This means that if we regard the hash function  $H$  as a truly random function, then any computationally feasible adversary with access to the public key and that can ask for the signature of messages of its choice, succeeds in forging a signature for a fresh message only with a negligible probability. This asymptotic security statement is desirable, but of limited practical utility because it does not give any hint as to how to choose the scheme parameters to attain a certain degree of security. A much more useful result would be an exact security statement, a bound that quantifies the gap between the security of the scheme and the intractability of inverting the trapdoor permutation.

Consider an adversary against the existential unforgeability of FDH that makes at most  $q_H(\eta)$  and  $q_S(\eta)$  queries to the hash and signing oracles, respectively. In a code-based setting, such an adversary is regarded as a black-box procedure  $\mathcal{A}$  run in the context of the following attack game:

<p><b>Game <math>G_{EF}</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}()</math>;  <math>L, S \leftarrow \text{nil}</math>;  <math>(m, \sigma) \leftarrow \mathcal{A}(pk)</math>;  <math>h \leftarrow H(m)</math></p>	<p><b>Oracle <math>H(m)</math> :</b>  if <math>m \notin \text{dom}(L)</math> then  <math>h \xleftarrow{\\$} G</math>; <math>L \leftarrow (m, h) :: L</math>  return <math>L[m]</math></p> <p><b>Oracle <math>\text{Sign}(m)</math> :</b>  <math>S \leftarrow m :: S</math>; <math>h \leftarrow H(m)</math>;  return <math>f^{-1}(sk, h)</math></p>
--	--

This adversary succeeds in forging a FDH signature for a fresh message with probability

$$\Pr[G_{EF} : h = f(pk, \sigma) \wedge m \notin S]$$

Note that in the above game the signing oracle makes a hash query each time the adversary asks for the signature of a message, and an additional hash query is made at the end as part of the verification of the signature returned by the adversary. Thus the number of effective hash queries made during the whole game is at most  $q_H + q_S + 1$ . All this is captured by the following post-condition of  $G_{EF}$ ,

$$\Psi \stackrel{\text{def}}{=} |L| \leq q_H + q_S + 1 \wedge |S| \leq q_S$$

This implies in particular that  $\Pr[G_{EF} : A] = \Pr[G_{EF} : A \wedge \Psi]$  for any event  $A$ .

In the remainder of this chapter we will show two different ways of constructing an inverter  $\mathcal{I}$  that uses the forger  $\mathcal{A}$  to invert  $f$ . These constructions effectively reduce the security of the signature scheme to the intractability of inverting the underlying trapdoor permutation.

## 5.1 The Original Proof

We first present a game-based proof of the original result by Bellare and Rogaway [1993], which provides a security bound that depends on the number of queries the adversary makes to both, the hash and the signing oracle.

**Theorem 5.3 (Original bound).** *Let  $\mathcal{A}$  be an adversary mounting a chosen-message existential forgery attack against FDH that makes at most  $q_H$  queries to the hash oracle  $H$  and at most  $q_S$  queries to the signing oracle  $\text{Sign}$ . Suppose  $\mathcal{A}$  succeeds in forging a signature for a fresh message within time  $t$  with probability  $\epsilon$ . Then, there exists an inverter  $\mathcal{I}$  that finds the preimage of an element uniformly drawn from the range of  $f$  with probability  $\epsilon'$  within time  $t'$ , where*

$$\epsilon' \geq (q_H + q_S + 1)^{-1} \epsilon \quad (5.1)$$

$$t' \leq t + (q_H + q_S) O(t_f) \quad (5.2)$$

and  $t_f$  is an upper bound for the time needed to compute the image of a group element under the permutation  $f$ .

*Proof.* The inverter  $\mathcal{I}$  shown in the context of game  $G_{OW}$  in Figure 5.1 achieves the probability and time bounds in the statement. It simulates an environment for  $\mathcal{A}$  where it replaces the hash and signing oracles with versions of its own.

<p><b>Game <math>G_{OW}</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>y \xleftarrow{\\$} G;</math>  <math>x \leftarrow \mathcal{I}(pk, y)</math></p> <p><b>Adversary <math>\mathcal{I}(pk, y)</math> :</b>  <math>\hat{pk} \leftarrow pk;</math>  <math>\hat{y} \leftarrow y;</math>  <math>j \xleftarrow{\\$} \{0, \dots, q\};</math>  <math>P, L \leftarrow \text{nil};</math>  <math>(m, \sigma) \leftarrow \mathcal{A}(pk);</math>  return <math>\sigma</math></p>	<p><b>Oracle <math>H(m)</math> :</b>  if <math>m \notin \text{dom}(L)</math> then    if <math> L  = j</math> then <math>h \leftarrow \hat{y}</math>    else <math>r \xleftarrow{\\$} G; h \leftarrow f(\hat{pk}, r)</math>  <math>P \leftarrow (m, r) :: P;</math>  <math>L \leftarrow (m, h) :: L;</math>  return <math>L[m]</math></p> <p><b>Oracle <math>\text{Sign}(m)</math> :</b>  <math>h \leftarrow H(m);</math>  return <math>P[m]</math></p>
--	--

**Fig. 5.1.** The inverter  $\mathcal{I}$  in the context of the one-wayness game for the family of trapdoor permutations  $(\mathcal{KG}_f, f, f^{-1})$ .

In order to stand a chance of forging a signature for a fresh message  $m$ , the adversary  $\mathcal{A}$  must ask for the hash value of  $m$  by querying oracle  $H$ . Otherwise, the hash of  $m$  would be completely random and independent of the adversary's output, and thus the purported signature  $\sigma$  would only be valid with a negligible probability. The inverter  $\mathcal{I}$  tries to guess the query where the hash value of  $m$  is asked for the first time. Let  $q \stackrel{\text{def}}{=} q_S + q_H$ . The inverter first randomly chooses an index  $j$  in  $\{0, \dots, q\}$  and then runs the forger intercepting its oracle queries. The inverter answers to the  $j$ -th hash query (we index the queries from 0) with its



own challenge  $y$ , and to the remaining hash queries with a random element in the range of  $f$  with a known preimage; it stores this preimage in a list  $\mathbf{P}$ . When the adversary makes a signing query, the inverter first makes the corresponding hash query itself and then obtains the preimage of the hash value under  $f$  from the list  $\mathbf{P}$ . The simulation is perfect provided the forger never asks the signature of the message corresponding to the  $j$ -th hash query, because in this case its preimage will not be in the list  $\mathbf{P}$ . A sufficient condition for the simulation to be correct is that the guess  $j$  be correct (i.e.  $m = M[j]$ ), because  $m$  cannot appear in a signing query (it must be fresh).

We can readily analyze the extra time the inverter spends in simulating the environment for  $\mathcal{A}$  in  $\mathsf{G}_{\text{OW}}$ . The only significant overhead is in the simulation of the hash oracle  $H$ . For all but one hash query, the simulated oracle computes the image under  $f$  of some element. The time bound (5.2) follows because the adversary makes (either directly or indirectly, through the signing oracle) at most  $q_{\text{H}} + q_{\text{S}}$  hash queries.

To prove that the probability bound in equation (5.1) holds, we will exhibit a sequence of games relating the probability

$$\epsilon' = \Pr [\mathsf{G}_{\text{OW}} : x = f^{-1}(sk, y)]$$

of  $\mathcal{I}$  successfully inverting  $f$  on a random challenge  $y$ , to the probability

$$\epsilon = \Pr [\mathsf{G}_{\text{EF}} : h = f(pk, \sigma) \wedge m \notin \mathbf{S}]$$

of adversary  $\mathcal{A}$  forging a signature for a fresh message. For each game, the main experiment is shown alongside the code of procedures in the environment; code pieces that change with respect to the previous game in the sequence appear on a gray background.

We start from the game  $\mathsf{G}_{\text{EF}}$  that encodes the existential forgery attack. In this initial game, the hash oracle  $H$  is implemented as a random oracle whereas the signing oracle is implemented as specified by the scheme. In order to be able to encode the freshness condition on the message whose signature is forged, the signing oracle is instrumented to record the queries it gets.

<b>Game <math>\mathsf{G}_{\text{EF}}</math> :</b> $(pk, sk) \leftarrow \mathcal{KG}();$ $\mathbf{L}, \mathbf{S} \leftarrow \text{nil};$ $(m, \sigma) \leftarrow \mathcal{A}(pk);$ $h \leftarrow H(m)$	<b>Oracle <math>H(m)</math> :</b> if $m \notin \text{dom}(\mathbf{L})$ then $h \xleftarrow{\$} G; \mathbf{L} \leftarrow (m, h) :: \mathbf{L}$ return $\mathbf{L}[m]$ <b>Oracle <math>\text{Sign}(m)</math> :</b> $\mathbf{S} \leftarrow m :: \mathbf{S}; h \leftarrow H(m);$ return $f^{-1}(sk, h)$
---	---

In game  $\mathsf{G}_1$  we instrument the hash oracle to keep track of the indices of queries. We use for this purpose a list  $\mathbf{M}$  where we store the messages queried to the hash oracle so far. Note that this is not really necessary because the value of  $\mathbf{M}$

can be recovered at any given moment from the value of  $\mathbf{L}$ . We only make this instrumentation for convenience and to make the proof cleaner. We also introduce the guess  $\mathbf{j}$  that will be used later by the inverter. We sample  $\mathbf{j}$  uniformly at the end of the game so that its independence from the output of the adversary is evident.

<b>Game <math>G_1</math> :</b> $(pk, sk) \leftarrow \mathcal{KG}_f();$ $\mathbf{M} \leftarrow \text{nil};$ $\mathbf{L}, \mathbf{S} \leftarrow \text{nil};$ $(m, \sigma) \leftarrow \mathcal{A}(pk);$ $h \leftarrow H(m);$ $\mathbf{j} \xleftarrow{\$} \{0, \dots, q\}$	<b>Oracle <math>H(m)</math> :</b> if $m \notin \text{dom}(\mathbf{L})$ then $h \xleftarrow{\$} G; \mathbf{L} \leftarrow (m, h) :: \mathbf{L};$ $\mathbf{M} \leftarrow m :: \mathbf{M}$ return $\mathbf{L}[m]$ <b>Oracle <math>\text{Sign}(m)</math> :</b> $\mathbf{S} \leftarrow m :: \mathbf{S}; h \leftarrow H(m);$ return $f^{-1}(sk, h)$
--	---

Consider the predicate

$$\Phi_1 \stackrel{\text{def}}{=} |\mathbf{L}| = |\mathbf{M}| \wedge (\forall m \in \text{dom}(\mathbf{L}). \exists i < |\mathbf{M}|. m = \mathbf{M}[i])$$

We prove that

$$\vdash \mathbf{G}_{\text{EF}} \sim \mathbf{G}_1 : \text{true} \Rightarrow =_{\{\mathbf{L}, \mathbf{S}, pk, m, \sigma, h\}} \wedge \Phi_1 \langle 2 \rangle \quad (5.3)$$

Using the tactics `wp` and `eqobs_in` we construct the relational procedure information for the oracles in the environment of both games. We then extend it automatically to the adversary  $\mathcal{A}$  to obtain the information  $\iota$  that we need to prove the equivalence. The script we use to prove (5.3) is just

`deadcode`  $\iota$ ; `eqobs_ctxt`  $\iota$ ; `wp`; ...

The tactic `deadcode` removes the random assignment to  $\mathbf{j}$  in  $G_1$ , while `eqobs_ctxt` removes the common prefix and suffix in both games except for the instruction  $\mathbf{L} \leftarrow \text{nil}$  because it affects the invariant  $\Phi_1 \langle 2 \rangle$ . The intermediate goal after applying these first two tactics is

$$\vdash \mathbf{L} \leftarrow \text{nil} \sim \mathbf{M}, \mathbf{L} \leftarrow \text{nil} : =_{\{pk, sk\}} \Rightarrow =_{\{\mathbf{L}\}} \wedge \Phi_1 \langle 2 \rangle$$

The tactic `wp` is then used to compute the weakest relational pre-condition of  $=_{\{\mathbf{L}\}} \wedge \Phi_1 \langle 2 \rangle$  with respect to the two resulting program fragments; the ellipsis stands for a straightforward script to prove that this weakest pre-condition holds. Games  $\mathbf{G}_{\text{EF}}$  and  $\mathbf{G}_1$  are thus equivalent on  $h, pk, \sigma, m$ , and  $\mathbf{S}$ , which implies

$$\Pr[\mathbf{G}_{\text{EF}} : h = f(pk, \sigma) \wedge m \notin \mathbf{S}] = \Pr[\mathbf{G}_1 : h = f(pk, \sigma) \wedge m \notin \mathbf{S}] \quad (5.4)$$

Games  $\mathbf{G}_{\text{EF}}$  and  $\mathbf{G}_1$  are also equivalent on all the variables appearing free in  $\Psi$ , so that  $\Psi$  is a post-condition of  $\mathbf{G}_1$  as well. Furthermore, since the game makes a last hash call for  $m$ ,  $m \in \text{dom}(\mathbf{L})$  is a post-condition of  $\mathbf{G}_1$ . We have that

$$\Psi \wedge \Phi_1 \wedge m \in \text{dom}(\mathbf{L}) \implies \exists i \leq q. m = \mathbf{M}[i]$$

This means that there exists at least one index  $i$  in  $\{0, \dots, q\}$  such that  $m = \mathbf{M}[i]$ . (In fact, there exists exactly one, but we do not need to prove this.) The probability of  $j$  being one of such indices is at least  $(q+1)^{-1}$  and is obviously independent of the success of the forgery, thus

$$\frac{\Pr[\mathbf{G}_1 : h = f(pk, \sigma) \wedge m \notin \mathbf{S}]}{q+1} \leq \Pr[\mathbf{G}_1 : h = f(pk, \sigma) \wedge m \notin \mathbf{S} \wedge m = \mathbf{M}[j]] \quad (5.5)$$

We apply now a semantics preserving transformation. Game  $\mathbf{G}_2$  eagerly samples the value  $\hat{y}$  that is given as answer to the  $j$ -th hash query, and that will later become the challenge to the inverter.

<p><b>Game <math>\mathbf{G}_2</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f</math>;  <math>\hat{y} \leftarrow G</math>;  <math>j \leftarrow \{0, \dots, q\}</math>;  <math>\mathbf{M}, \mathbf{L}, \mathbf{S} \leftarrow \text{nil}</math>;  <math>(m, \sigma) \leftarrow \mathcal{A}(pk)</math>;  <math>h \leftarrow H(m)</math></p>	<p><b>Oracle <math>H(m)</math> :</b>  if <math>m \notin \text{dom}(\mathbf{L})</math> then    if <math> \mathbf{L}  = j</math> then <math>h \leftarrow \hat{y}</math>    else <math>h \leftarrow G</math>;  <math>\mathbf{L} \leftarrow (m, h) :: \mathbf{L}</math>;  <math>\mathbf{M} \leftarrow m :: \mathbf{M}</math>  return <math>\mathbf{L}[m]</math></p> <p><b>Oracle <math>\text{Sign}(m)</math> :</b>  <math>\mathbf{S} \leftarrow m :: \mathbf{S}</math>; <math>h \leftarrow H(m)</math>;  return <math>f^{-1}(sk, h)</math></p>
--	--

We obtain

$$\vdash \mathbf{G}_1 \simeq_{\{j, \mathbf{M}, \mathbf{L}, \mathbf{S}, pk, m, \sigma, h\}}^{\emptyset} \mathbf{G}_2$$

Therefore,

$$\Pr[\mathbf{G}_1 : h = f(pk, \sigma) \wedge m \notin \mathbf{S} \wedge m = \mathbf{M}[j]] = \Pr[\mathbf{G}_2 : h = f(pk, \sigma) \wedge m \notin \mathbf{S} \wedge m = \mathbf{M}[j]] \quad (5.6)$$

In the next game we modify the way hash queries are computed. For all but the  $j$ -th query we return the image under  $f$  of a uniformly sampled element in its domain, and we store this element in a list  $\mathbf{P}$ . This is a local change that does not modify the distribution of the answers. Indeed, since  $f$  is a permutation we have

$$\vdash h \leftarrow G \simeq_{\{h\}}^{\emptyset} r \leftarrow G; h \leftarrow f(pk, r)$$

In preparation for the next transformation, we also introduce a flag **bad** to signal whether the simulation failed, i.e. whether the adversary asked for the signature of  $\mathbf{M}[j]$ .

<b>Game <math>G_3</math> :</b> $(pk, sk) \leftarrow \mathcal{KG}_f();$ $\hat{y} \leftarrow G;$ $j \leftarrow \{0, \dots, q\};$ $P \leftarrow \text{nil};$ $M, L, S \leftarrow \text{nil};$ $(m, \sigma) \leftarrow \mathcal{A}(pk);$ $h \leftarrow H(m)$	<b>Oracle <math>H(m)</math> :</b> if $m \notin \text{dom}(L)$ then if $ L  = j$ then $h \leftarrow \hat{y}$ else $r \leftarrow G; h \leftarrow f(pk, r)$ $P \leftarrow (m, r) :: P;$ $L \leftarrow (m, h) :: L;$ $M \leftarrow m :: M$ return $L[m]$ <b>Oracle <math>\text{Sign}(m)</math> :</b> $S \leftarrow m :: S; h \leftarrow H(m);$ if $m = M[j]$ then $\text{bad} \leftarrow \text{true};$ return $f^{-1}(sk, h)$ else return $f^{-1}(sk, h)$
---	--

We prove the equivalence

$$\vdash G_2 \sim G_3 : \text{true} \Rightarrow =_{\{j, M, L, S, pk, m, \sigma, h\}} \wedge (M[j] \notin S \Rightarrow \neg \text{bad}) \langle 2 \rangle$$

Hence

$$\begin{aligned} \Pr[G_2 : h = f(pk, \sigma) \wedge m \notin S \wedge m = M[j]] \\ = \Pr[G_3 : h = f(pk, \sigma) \wedge m \notin S \wedge m = M[j]] \end{aligned} \quad (5.7)$$

In game  $G_4$  we modify the signing oracle so that if the signature of the message  $M[j]$  is ever asked, instead of actually computing the preimage of its digest using  $sk$ , the signing oracle simply returns the corresponding  $P$ -entry. This entry will be undefined, but this poses no problem because as long as the guess  $j$  is correct this piece of code is unreachable.

<b>Game <math>G_4</math> :</b> $(pk, sk) \leftarrow \mathcal{KG}_f();$ $\hat{y} \leftarrow G;$ $j \leftarrow \{0, \dots, q\};$ $P \leftarrow \text{nil};$ $M, L, S \leftarrow \text{nil};$ $(m, \sigma) \leftarrow \mathcal{A}(pk);$ $h \leftarrow H(m)$	<b>Oracle <math>H(m)</math> :</b> if $m \notin \text{dom}(L)$ then if $ L  = j$ then $h \leftarrow \hat{y}$ else $r \leftarrow G; h \leftarrow f(pk, r)$ $P \leftarrow (m, r) :: P;$ $L \leftarrow (m, h) :: L;$ $M \leftarrow m :: M$ return $L[m]$ <b>Oracle <math>\text{Sign}(m)</math> :</b> $S \leftarrow m :: S; h \leftarrow H(m);$ if $m = M[j]$ then $\text{bad} \leftarrow \text{true};$ return $P[m]$ else return $f^{-1}(sk, h)$
---	---

Games  $G_3$  and  $G_4$  differ only in a portion of code that appears after **bad** is set, therefore they are syntactically equal up to the failure event **bad**. The Fundamental Lemma gives us

$$\begin{aligned} \Pr[G_3 : h = f(\mathbf{pk}, \sigma) \wedge m \notin \mathbf{S} \wedge m = \mathbf{M}[j] \wedge \neg \mathbf{bad}] \\ = \Pr[G_4 : h = f(\mathbf{pk}, \sigma) \wedge m \notin \mathbf{S} \wedge m = \mathbf{M}[j] \wedge \neg \mathbf{bad}] \end{aligned}$$

Since  $\mathbf{M}[j] \notin \mathbf{S} \implies \neg \mathbf{bad}$  is a post-condition of  $G_3$ , it does not change anything if we remove the last term  $\neg \mathbf{bad}$  from the event under consideration,

$$\begin{aligned} \Pr[G_3 : h = f(\mathbf{pk}, \sigma) \wedge m \notin \mathbf{S} \wedge m = \mathbf{M}[j]] \\ = \Pr[G_4 : h = f(\mathbf{pk}, \sigma) \wedge m \notin \mathbf{S} \wedge m = \mathbf{M}[j] \wedge \neg \mathbf{bad}] \quad (5.8) \\ \leq \Pr[G_4 : h = f(\mathbf{pk}, \sigma) \wedge m \notin \mathbf{S} \wedge m = \mathbf{M}[j]] \end{aligned}$$

When answering a signing query for a message  $m \neq \mathbf{M}[j]$ , we may obtain the preimage of its hash value from  $\mathbf{P}$  rather than using  $f^{-1}$ , as in game  $G_5$ .

<p><b>Game <math>G_5</math> :</b>  <math>(\mathbf{pk}, sk) \leftarrow \mathcal{KG}_f();</math>  <math>\hat{y} \xleftarrow{\\$} G;</math>  <math>j \xleftarrow{\\$} \{0, \dots, q\};</math>  <math>\mathbf{P}, \mathbf{L} \leftarrow \text{nil};</math>  <math>(m, \sigma) \leftarrow \mathcal{A}(\mathbf{pk});</math>  <math>y \leftarrow \hat{y};</math>  <math>x \leftarrow \sigma</math></p>	<p><b>Oracle <math>H(m)</math> :</b>  if <math>m \notin \text{dom}(\mathbf{L})</math> then    if <math> \mathbf{L}  = j</math> then <math>h \leftarrow \hat{y}</math>    else <math>r \xleftarrow{\\$} G; h \leftarrow f(\mathbf{pk}, r)</math>  <math>\mathbf{P} \leftarrow (m, r) :: \mathbf{P};</math>  <math>\mathbf{L} \leftarrow (m, h) :: \mathbf{L}</math>  return <math>\mathbf{L}[m]</math>  <b>Oracle <math>\text{Sign}(m)</math> :</b>  <math>h \leftarrow H(m);</math>  return <math>\mathbf{P}[m]</math></p>
---	--

Define  $\Phi_4$  as

$$\begin{aligned} \forall m \in \text{dom}(\mathbf{L}). m \neq \mathbf{M}[j] \implies \mathbf{P}[m] = f^{-1}(sk, \mathbf{L}[m]) \wedge \\ (j \leq |\mathbf{M}| \implies \mathbf{L}[\mathbf{M}[j]] = \hat{y}) \end{aligned}$$

We use  $\Phi_4$  to prove that the signing oracles in games  $G_4$  and  $G_5$  are equivalent. To this end we show that

$$\vdash G_4 \sim G_5 : \text{true} \Rightarrow =_{\{\hat{y}, j, \mathbf{L}, m, \sigma, h\}} \wedge \Phi_4 \langle 1 \rangle$$

If  $m = \mathbf{M}[j]$ , then certainly  $j \leq |\mathbf{M}|$  when the game finishes. In this case, post-condition  $\Phi_4$  implies that  $h = \mathbf{L}[\mathbf{M}[j]] = \hat{y}$ , which in turn gives

$$\Pr[G_4 : h = f(\mathbf{pk}, \sigma) \wedge m \notin \mathbf{S} \wedge m = \mathbf{M}[j]] \leq \Pr[G_5 : f^{-1}(sk, y) = x] \quad (5.9)$$

We finally prove that  $\vdash G_5 \simeq_{\{sk, x, y\}}^{\emptyset} G_{\text{OW}}$ . We need to inline the call to the inverter  $\mathcal{I}$  in  $G_{\text{OW}}$  and use  $y$  instead of  $\hat{y}$  in  $G_5$ . The proof script is straightforward,

alloc\_l  $\hat{y}$   $y$ ; sinline\_r  $\iota$   $\mathcal{I}$ .  
 eqobs\_tl  $\iota$ ; deadcode; eqobs\_in.

Using the above equivalence we derive

$$\Pr [\mathbf{G}_5 : f^{-1}(sk, y) = x] = \Pr [\mathbf{G}_{\text{OW}} : f^{-1}(sk, y) = x] \quad (5.10)$$

Putting all the above results together, we conclude

$$\frac{\Pr [\mathbf{G}_{\text{EF}} : h = f(pk, \sigma) \wedge m \notin \mathcal{S}]}{q + 1} \leq \Pr [\mathbf{G}_{\text{OW}} : f^{-1}(sk, y) = x] \quad (5.11)$$

□

A complete proof of asymptotic security follows trivially. Since  $f$  is a one-way permutation family,  $\Pr [\mathbf{G}_{\text{OW}} : f^{-1}(sk, y) = x]$  is negligible in the security parameter provided  $\mathcal{I}$  runs in PPT. This is indeed the case, and is proved automatically in CertiCrypt. The number of queries  $q$  made by the forger to the hash and signing oracles must necessarily be polynomial on the security parameter  $\eta$ . Since the product of a negligible function and a polynomial is still a negligible function, it follows from (5.11) that the probability of a successful existential forgery is negligible.

## 5.2 Improved Bound

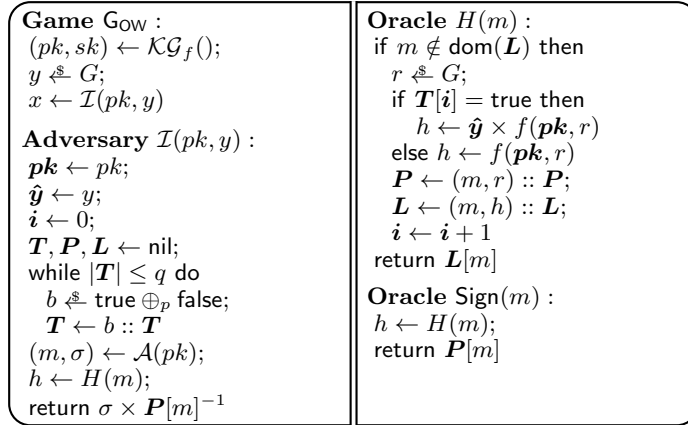
A tighter security bound for FDH appears in [Coron 2000]; this bound is independent of the number of hash queries. This is of much practical significance since the number of hash values a real-world forger can compute is only limited by the time and computational resources it invests, whereas the number of signatures it gets could be limited by the owner of the private key. Once the owner of a key has used it to sign a certain quantity of messages, he could simply discard that key and generate a new one.

**Theorem 5.4 (Improved bound).** *Assume the underlying trapdoor permutation  $(\mathcal{KG}_f, f, f^{-1})$  is homomorphic with respect to the group operation in its domain, i.e. for every  $(pk, sk)$  that might be output by  $\mathcal{KG}_f$ , and every  $x, y$ ,  $f(pk, x \times y) = f(pk, x) \times f(pk, y)$ . Let  $\mathcal{A}$  be an adversary against the existential unforgeability of FDH that makes at most  $q_{\text{H}}$  and  $q_{\text{S}}$  queries to the hash and signing oracles respectively. Suppose  $\mathcal{A}$  succeeds in forging a signature for a fresh message within time  $t$  with probability  $\epsilon$  during experiment  $\mathbf{G}_{\text{EF}}$ . Then, there exists an inverter  $\mathcal{I}$  that finds the preimage of an element uniformly drawn from the range of  $f$  with probability  $\epsilon'$  within time  $t'$  during experiment  $\mathbf{G}_{\text{OW}}$ , where*

$$\epsilon' \geq \frac{1}{q_{\text{S}} + 1} \left( 1 - \frac{1}{q_{\text{S}} + 1} \right)^{q_{\text{S}}} \epsilon \quad (5.12)$$

$$t' \leq t + (q_{\text{H}} + q_{\text{S}}) O(t_f) \quad (5.13)$$

These bounds hold for the inverter shown in Figure 5.2. The inverter first samples  $q + 1$  bits at random, choosing **true** with probability  $p$  and **false** with probability  $(1 - p)$ , and stores them in a list  $\mathbf{T}$ . It answers to the  $i$ -th hash query as follows: it picks uniformly a value  $r$  from the domain of  $f$  and stores it in a list  $\mathbf{P}$ , then replies according to the  $i$ -th entry in  $\mathbf{T}$ : if it is **true**, answers with  $y \times f(\mathbf{pk}, r)$  where  $y$  is its challenge, if it is **false** answers with simply  $f(\mathbf{pk}, r)$ . In both cases the answers are indistinguishable from those of a random function. When the adversary asks for the signature of a message  $m$ , the inverter makes the corresponding hash query itself and then answers with the  $m$  entry in the list  $\mathbf{P}$ . The simulation is correct provided the entries in  $\mathbf{T}$  corresponding to messages appearing in signing queries are **false**, because in this case the corresponding entries in  $\mathbf{P}$  coincide with the preimage of their hash value. The aim of the inverter is to inject its challenge in as many hash queries as possible, while at the same time maximizing the probability of the simulation being correct. The parameter  $p$  is left unspecified through the proof and will be chosen later to find the best compromise between these two competing goals.



**Fig. 5.2.** The inverter  $\mathcal{I}$  in the context of the one-wayness game for the family of trapdoor permutations  $(\mathcal{KG}_f, f, f^{-1})$ . We use  $(\text{true} \oplus_p \text{false})$  to denote a Bernoulli distribution with success probability  $p$ , i.e. the discrete distribution that takes value **true** with probability  $p$  and **false** with probability  $(1 - p)$ .

The success of the inverter is not guaranteed by the sole success of the forger. It depends on two additional conditions: that the simulation is consistent, so that the forger behaves as expected, and that the forgery can be used to compute the preimage of the challenge  $y$ .

Let us analyze first the probability of the simulation being consistent. The forger  $\mathcal{A}$  must not be able to distinguish the simulation from the experiment in  $G_{EF}$ . The forger's view, and in particular the distribution of the answers it gets from the hash and signing oracles must be the same as in the original experiment. The responses





$$\begin{aligned}
& \Pr [G_1, \mu : h = f(pk, \sigma) \wedge \text{success}] \\
&= \llbracket G_1 \rrbracket \mu \mathbb{1}_{(h=f(pk, \sigma) \wedge \text{success})} \\
&= \llbracket G'_1 \rrbracket \mu (\lambda \mu'. \llbracket \text{Init}_{\mathbf{T}} \rrbracket \mu' (\lambda \mu''. \mathbb{1}_{h=f(pk, \sigma)} \mu'' \times \mathbb{1}_{\text{success}} \mu'')) \\
&= \llbracket G'_1 \rrbracket \mu (\lambda \mu'. \mathbb{1}_{(h=f(pk, \sigma) \wedge m \notin \mathbf{S})} \mu' \times \Pr [\text{Init}_{\mathbf{T}}, \mu' : \text{success}])
\end{aligned} \tag{5.15}$$

Thus, computing the original probability is equivalent to measuring a function that can be expressed as a product of two factors, one of which is the probability of event `success` after initializing  $\mathbf{T}$  in the intermediate memory  $\mu'$ . The advantage of doing this is that the second factor in the product is upper bounded by a constant under some conditions on  $\mu'$  that we can prove to hold. Namely, we can show that the following is a post-condition of the piece of code  $G'_1$

$$|\mathbf{L}| \leq q + 1 \wedge |\mathbf{S}| \leq q_S \wedge \text{ran}(\mathbf{I}) = [|\mathbf{L}| - 1 .. 0] \tag{5.16}$$

Furthermore, we can assume that  $m$  is not in  $\mathbf{S}$  in memory  $\mu'$  because otherwise the measured function would be null. Under this conditions, we will show that

$$p(1-p)^{q_S} \leq \Pr [\text{Init}_{\mathbf{T}}, \mu' : \text{success}]$$

Consider the loop that initializes  $\mathbf{T}$ ,

$$c \stackrel{\text{def}}{=} \text{while } |\mathbf{T}| \leq q \text{ do } (b \stackrel{\$}{\leftarrow} \text{true} \oplus_p \text{false}; \mathbf{T} \leftarrow b :: \mathbf{T})$$

and define

$$\begin{aligned}
F_\mu(i, n) &\stackrel{\text{def}}{=} \text{if } n < i \text{ then } \llbracket \mathbf{T}[i-n] \rrbracket \mu \text{ else } p \\
G_\mu(l, n) &\stackrel{\text{def}}{=} \prod_{i \in l} \text{if } n < i \text{ then } \llbracket \neg \mathbf{T}[i-n] \rrbracket \mu \text{ else } 1-p
\end{aligned}$$

This may seem abstruse at first sight, but intuitively,  $F_\mu(i, q-k)$  equals the probability of  $\mathbf{T}[i]$  being `true` after executing `Init $\mathbf{T}$`  in a memory  $\mu$  where  $|\mathbf{T}| = k$ , while  $G_\mu(i, q-k)$  is a lower bound on the probability of  $\forall i \in l. \neg \mathbf{T}[i]$ . We prove the following invariant about the while loop  $c$ : for every index  $i$  and list of indices  $l$  such that  $i \notin l$ ,

$$F_\mu(i, q - \llbracket \mathbf{T} \rrbracket \mu) G_\mu(l, q - \llbracket \mathbf{T} \rrbracket \mu) \leq \Pr [c, \mu : \mathbf{T}[i] \wedge \forall i \in l. \neg \mathbf{T}[i]] \tag{5.17}$$

For any memory  $\mu$  satisfying (5.16) and where  $m \notin \mathbf{S}$ , if we take  $i = \mathbf{I}[m]$  and  $l = \mathbf{I}[\mathbf{S}]$ , we obtain

$$\begin{aligned}
p(1-p)^{q_S} &\leq F_\mu(i, q) G_\mu(l, q) \\
&\leq \Pr [\text{Init}_{\mathbf{T}}, \mu : \mathbf{T}[i] \wedge \forall i \in l. \neg \mathbf{T}[i]] \\
&= \Pr [\text{Init}_{\mathbf{T}}, \mu : \text{success}]
\end{aligned}$$

Since any intermediate memory  $\mu'$  in (5.15) satisfies (5.16), it follows from the above inequality that

$$\begin{aligned}
& p (1-p)^{qs} \Pr[\mathbf{G}_{\text{EF}} : h = f(pk, \sigma) \wedge m \notin \mathbf{S}] \\
&= p (1-p)^{qs} \Pr[\mathbf{G}_1 : h = f(pk, \sigma) \wedge m \notin \mathbf{S}] \\
&\leq \Pr[\mathbf{G}_1 : h = f(pk, \sigma) \wedge \text{success}]
\end{aligned} \tag{5.18}$$

In the next game,  $\mathbf{G}_2$ , we modify the hash oracle to answer to the  $i$ -th query according to the corresponding entry in list  $\mathbf{T}$ : the oracle samples a random element  $r$  in the group, stores it in list  $\mathbf{P}$ , and responds  $\hat{\mathbf{y}} \times f(\mathbf{pk}, r)$  if  $\mathbf{T}[i]$  is true, and simply  $f(\mathbf{pk}, r)$  otherwise. In this latter case, the inverse image of the hash value can be recovered from  $\mathbf{P}$ .

<p><b>Game <math>\mathbf{G}_2</math> :</b></p> <p><math>\hat{\mathbf{y}} \xleftarrow{\\$} G;</math>  <math>\mathbf{T} \leftarrow \text{nil};</math>  while <math> \mathbf{T}  \leq q</math> do  <math>b \xleftarrow{\\$} \text{true} \oplus_p \text{false};</math>  <math>\mathbf{T} \leftarrow b :: \mathbf{T}</math>  <math>(\mathbf{pk}, \mathbf{sk}) \leftarrow \mathcal{KG}_f();</math>  <math>i \leftarrow 0; \mathbf{I} \leftarrow \text{nil};</math>  <math>\mathbf{P}, \mathbf{L}, \mathbf{S} \leftarrow \text{nil};</math>  <math>(m, \sigma) \leftarrow \mathcal{A}(\mathbf{pk});</math>  <math>h \leftarrow H(m)</math></p>	<p><b>Oracle <math>H(m)</math> :</b></p> if $m \notin \text{dom}(\mathbf{L})$ then $r \xleftarrow{\$} G;$ if $\mathbf{T}[i]$ then $h \leftarrow \hat{\mathbf{y}} \times f(\mathbf{pk}, r)$ else $h \leftarrow f(\mathbf{pk}, r)$ $\mathbf{P} \leftarrow (m, r) :: \mathbf{P};$ $\mathbf{L} \leftarrow (m, h) :: \mathbf{L};$ $\mathbf{I} \leftarrow (m, i) :: \mathbf{I};$ $i \leftarrow i + 1$ return $\mathbf{L}[m]$ <b>Oracle <math>\text{Sign}(m)</math> : ...</b>
---	---

Observe that the transformation of  $\mathbf{G}_1$  into  $\mathbf{G}_2$  can be justified by locally reasoning on the code of the hash oracle, without needing to apply the lazy sampling technique (as we had to do in the previous proof), thanks to the fact that  $f$  is a permutation and  $f(\mathbf{pk}, r)$  acts as a one-time pad,

$$\vdash h \xleftarrow{\$} G \simeq_{\{h\}}^{\emptyset} r \xleftarrow{\$} G; h \leftarrow \hat{\mathbf{y}} \times f(\mathbf{pk}, r)$$

We now modify the way the signing oracle answers to a query  $m$  when the corresponding entry  $\mathbf{T}[\mathbf{I}[m]]$  is true: instead of answering with a proper signature, it answers with just  $\mathbf{P}[m]$ . By using the Fundamental Lemma we will see that this change does not modify the probability of the event that interests us.

<p><b>Game <math>\overline{[\mathbf{G}_3]} \overline{[\mathbf{G}_4]}</math> :</b></p> <p><math>\hat{\mathbf{y}} \xleftarrow{\\$} G;</math>  <math>\mathbf{T} \leftarrow \text{nil};</math>  while <math> \mathbf{T}  \leq q</math> do  <math>b \xleftarrow{\\$} \text{true} \oplus_p \text{false};</math>  <math>\mathbf{T} \leftarrow b :: \mathbf{T}</math>  <math>(\mathbf{pk}, \mathbf{sk}) \leftarrow \mathcal{KG}_f();</math>  <math>i \leftarrow 0; \mathbf{I} \leftarrow \text{nil};</math>  <math>\mathbf{P}, \mathbf{L}, \mathbf{S} \leftarrow \text{nil};</math>  <math>(m, \sigma) \leftarrow \mathcal{A}(\mathbf{pk});</math>  <math>h \leftarrow H(m)</math></p>	<p><b>Oracle <math>H(m)</math> : ...</b></p> <p><b>Oracle <math>\text{Sign}(m)</math> :</b>  <math>\mathbf{S} \leftarrow m :: \mathbf{S}; h \leftarrow H(m);</math>  if <math>\mathbf{T}[\mathbf{I}[m]]</math> then  <math>\text{bad} \leftarrow \text{true};</math>  return <math>\overline{[f^{-1}(\mathbf{sk}, h)]} \overline{[\mathbf{P}[m]]}</math>  else return <math>f^{-1}(\mathbf{sk}, h)</math></p>
--	---

Game  $G_3$  is semantically identical to  $G_2$  (ignoring **bad**) and differs from  $G_4$  only in a portion of code appearing after flag **bad** is set. Thus,

$$\Pr[G_3 : h = f(pk, \sigma) \wedge \text{success} \wedge \neg \text{bad}] = \Pr[G_4 : h = f(pk, \sigma) \wedge \text{success} \wedge \neg \text{bad}]$$

But since  $\text{success} \implies \neg \text{bad}$  is a post-condition of  $G_3$ , it does not change anything if we remove the last term  $\neg \text{bad}$  from the event under consideration,

$$\begin{aligned} \Pr[G_3 : h = f(pk, \sigma) \wedge \text{success}] &= \Pr[G_4 : h = f(pk, \sigma) \wedge \text{success} \wedge \neg \text{bad}] \\ &\leq \Pr[G_4 : h = f(pk, \sigma) \wedge \text{success}] \end{aligned} \quad (5.19)$$

The signing oracle might as well respond with  $P[m]$  to every query, regardless of the value of  $T[I[m]]$ ,

<p><b>Game <math>G_5</math> :</b>  <math>\hat{y} \xleftarrow{\\$} G</math>;  <math>T \leftarrow \text{nil}</math>;  while <math> T  \leq q</math> do  <math>b \xleftarrow{\\$} \text{true} \oplus_p \text{false}</math>;  <math>T \leftarrow b :: T</math>  <math>(pk, sk) \leftarrow \mathcal{KG}_f()</math>;  <math>i \leftarrow 0</math>; <math>I \leftarrow \text{nil}</math>;  <math>P, L, S \leftarrow \text{nil}</math>;  <math>(m, \sigma) \leftarrow \mathcal{A}(pk)</math>;  <math>h \leftarrow H(m)</math>;  <math>y \leftarrow \hat{y}</math>;  <math>x \leftarrow \sigma \times P[m]^{-1}</math></p>	<p><b>Oracle <math>H(m) : \dots</math></b>  if <math>m \notin \text{dom}(L)</math> then  <math>r \xleftarrow{\\$} G</math>;  if <math>T[i]</math> then <math>h \leftarrow \hat{y} \times f(pk, r)</math>  else <math>h \leftarrow f(pk, r)</math>  <math>P \leftarrow (m, r) :: P</math>;  <math>L \leftarrow (m, h) :: L</math>;  <math>I \leftarrow (m, i) :: I</math>;  <math>i \leftarrow i + 1</math>  return <math>L[m]</math>  <b>Oracle <math>\text{Sign}(m) :</math></b>  <math>S \leftarrow m :: S</math>; <math>h \leftarrow H(m)</math>;  return <math>P[m]</math></p>
---	--

We can guarantee that the response given is a proper signature when  $T[I[m]]$  is false by proving the following post-condition of game  $G_5$ :

$$\begin{aligned} \forall (m, h) \in L. \quad T[I[m]] &\implies h = \hat{y} \times f(pk, P[m]) \wedge \\ \neg T[I[m]] &\implies h = f(pk, P[m]) \end{aligned}$$

This allows us to show that the signing oracles in  $G_4$  and  $G_5$  are equivalent and, using the homomorphic property of  $f$ , to show

$$\begin{aligned} \Pr[G_4 : h = f(pk, \sigma) \wedge \text{success}] &= \Pr[G_5 : h = f(pk, \sigma) \wedge \text{success}] \\ &\leq \Pr[G_5 : y \times f(pk, P[m]) = f(pk, \sigma)] \\ &= \Pr[G_{OW} : f^{-1}(sk, y) = x] \end{aligned}$$

Therefore, we can conclude

$$p(1-p)^{q_S} \Pr[G_{EF} : h = f(pk, \sigma) \wedge m \notin S] \leq \Pr[G_{OW} : f^{-1}(sk, y) = x]$$

We get the bound in the statement of the theorem by choosing  $p = (q_S + 1)^{-1}$ , which maximizes the factor  $p(1-p)^{q_S}$ . For this value of  $p$ , the factor approximates  $\exp(-1) q_S^{-1}$  for large values of  $q_S$ .

### 5.3 Practical Interpretation

If one accepts that it is reasonable to draw practical conclusions from a security proof in the random oracle model, then the results above may be used to choose the scheme parameters based on an estimate of the time needed to invert the underlying trapdoor permutation.

The best known method to invert the RSA function is to factor its modulus. The General Number Field Sieve (GNFS) [Lenstra and Jr. 1993] is currently the most efficient way of factoring large composite integers like RSA moduli; it has been used to factor several RSA moduli from the RSA Factoring Challenge, including the 512-bit RSA-155 number and the largest (non-special) integer factored with a general-purpose algorithm, the 768-bit RSA-768 number [Kleinjung et al. 2010]. On heuristic grounds, an odd composite number  $N$  can be factored using GNFS in time sub-exponential with respect to its size in bits. Concretely, let

$$L_N[\alpha, \beta] \stackrel{\text{def}}{=} \exp((\beta + o(1)) \log(N)^\alpha \log(\log(N))^{1-\alpha}) \quad (5.20)$$

where the constant term that accompanies  $\beta$  tends towards zero as  $N$  increases. A good implementation of GNFS takes about  $L_N[1/3, 1.923]$  time to factor a number  $N$ . This heuristic cannot be used directly to estimate the number of operations required to factor a certain  $N$  [Lenstra and Verheul 2001]. However, experimental data suggests that it can be used for limited range extrapolation. If one knows, empirically, that factoring an RSA modulus  $N$  using GNFS takes time  $t$ , then factoring an RSA modulus  $M > N$  will take approximately time

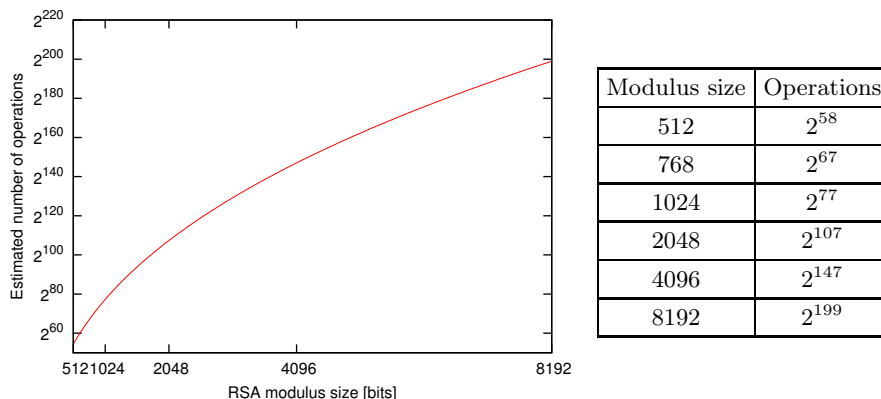
$$t \times \frac{L_M[1/3, 1.923]}{L_N[1/3, 1.923]} \quad (5.21)$$

(omitting the constant term  $o(1)$ ). If  $M \gg N$ , however, the effect of the constant term can no longer be ignored and the extrapolation will overestimate the time needed to factor  $M$ .

The computational effort involved in the factorization of the 512-bit number RSA-155 has been estimated at around 8400 MIPS-years<sup>1</sup>, or slightly less than  $2^{58}$  operations [Cavallar et al. 2000]. In comparison, the computational effort involved in the factorization of the 768-bit number RSA-768 has been estimated at around  $2^{67}$  operations [Kleinjung et al. 2010]. Extrapolating from this estimate using (5.21), we can make a rough prediction of the computational effort that would take to break larger RSA moduli (Figure 5.3).

GNFS could factor a 1024-bit number in around  $2^{77}$  operations, and a 2048-bit number in around  $2^{107}$  operations. Assume some safe bounds for  $q_H$  and  $q_S$ ,  $q_H \leq 2^{60}$ ,  $q_S \leq 2^{20}$ . To ensure that no forger within these bounds could forge a RSA-FDH signature within  $t = 2^{80}$  operations, one should pick an RSA modulus such that factoring it takes at least  $q_S(t + (q_H + q_S)O(t_{\text{RSA}}))$  operations, otherwise

<sup>1</sup> One MIPS-year is the equivalent of a computation running during a full year at a sustained rate of one million instructions per second. Consumer desktop PCs at the time of this writing attain speeds of up to 80,000 MIPS.



**Fig. 5.3.** Estimates for the computational effort required to factor large RSA-moduli.

one can iterate the construction in Theorem 5.4 about  $q_S$  times to invert the RSA function in less time than using the GNFS algorithm. For a modulus of size  $k$ ,  $t_{\text{RSA}} = O(k^2)$  when the public exponent is small. A 1024-bit modulus would not be enough, but a 2048-bit modulus would do. In contrast, if one were to choose the modulus according to the original security bound, even a 2048-bit modulus would not suffice.

The above guidelines for selecting key sizes by extrapolation should be taken with care. One should not forget that besides the fact that the analysis is based on heuristic ground, we are ignoring the  $o(1)$  factor from (5.20). Shamir and Tromer [2003] said to this respect:

“To determine what key sizes are appropriate for a given application, one needs concrete estimates for the cost of factoring integers of various sizes. Predicting these costs has proved notoriously difficult, for two reasons. First, the performance of modern factoring algorithms is not understood very well: their complexity analysis is often asymptotic and heuristic, and leaves large uncertainty factors. Second, even when the exact algorithmic complexity is known, it is hard to estimate the concrete cost of a suitable hypothetical large-scale computational effort using current technology; it’s even harder to predict what this cost would be at the end of the key’s planned lifetime, perhaps a decade or two into the future.”

Even leaving aside the possibility of a breakthrough in number theory or the discovery of a new factorization method that drastically improves on GNFS, there are a myriad of tweaks and hardware optimizations that can be readily—or at least hypothetically—applied to cut down the cost of factoring an RSA modulus using GNFS. It is conjectured that it could be possible to factor 1024-bit integers, and hence to break 1024-bit RSA keys, in 1 year using a special hardware device that could be built at a cost of US\$10M [Lenstra et al. 2003; Shamir and Tromer 2003].

## 5.4 Discussion and Related Work

An earlier example of a signature scheme that follows the same hash-then-decrypt pattern as FDH is the PKCS #1 scheme proposed by RSA Labs [RSA Data Security, Inc. 2002]. The latest version of PKCS #1 uses the following message encoding as a hash function

$$\text{ENCODE}(m, len) \stackrel{\text{def}}{=} 00 \parallel 01 \parallel \text{FF} \cdots \text{FF} \parallel 00 \parallel h(m)$$

Enough FF bytes are inserted to reach the intended length  $len$  of the encoded message. The value  $h(m)$  is (an encoding of) the digest of the message computed using one of MD2, MD5, SHA-1, SHA-256, SHA-384 or SHA-512. The problem with the above padding is that it does not map messages to the entire domain of the RSA function, but instead to a much smaller set of encoded messages. As a consequence, the scheme is not known to admit a security reduction to the standard RSA problem.

The generic Full-Domain Hash signature scheme based on a one-way trapdoor permutation was first described in the seminal work of Bellare and Rogaway [1993] to illustrate the applicability of the Random Oracle Model. In that work, the authors give in an appendix a sketch of a proof of its security against adaptive chosen-message attacks. A more detailed proof of the exact security of the RSA-based variant of FDH is given in [Bellare and Rogaway 1996, Theorem 3.1]; the bound corresponds to the one we showed in Section 5.1. The bound we proved in Section 5.2 is due to Coron [2000]. Unfortunately, it is not possible to further improve the security bound of FDH and prove that computing a forgery is as hard as inverting RSA. There is no tighter reduction than the one in Theorem 5.4 as showed by Coron [2002].

Bellare and Rogaway [1996] proposed the Probabilistic Signature Scheme (PSS) as a replacement for FDH; it has since then been incorporated into the PKCS #1 standard as an alternative signing method. PSS is roughly as efficient as FDH but admits a tight security reduction. The Probabilistic FDH scheme (PFDH) [Coron 2002] is a simple probabilistic variant of FDH that follows the same design principles as PSS and also admits a tight reduction, but at the cost of slightly longer signatures. As FDH, PFDH uses a hash function  $H : \{0, 1\}^* \rightarrow G$  and a trapdoor permutation  $f$ . Formally, PFDH is parametrized by the length  $k_0$  of the random “salt” it uses, and is composed of the following triple of algorithms:

$$\begin{aligned} \mathcal{KG} & \stackrel{\text{def}}{=} (pk, sk) \leftarrow \mathcal{KG}_f; \text{ return } (pk, sk) \\ \text{Sign}(sk, m) & \stackrel{\text{def}}{=} r \xleftarrow{\$} \{0, 1\}^{k_0}; \text{ return } (f^{-1}(sk, H(m \parallel r)), r) \\ \text{Verify}(pk, m, (\sigma, r)) & \stackrel{\text{def}}{=} \text{if } f(pk, \sigma) = H(m \parallel r) \text{ then return true else return false} \end{aligned}$$

Observe that FDH is obtained as a special case of this scheme by setting  $k_0 = 0$ . We have the following result about the security of PFDH.

**Theorem 5.5.** *Let  $\mathcal{A}$  be an adversary mounting a chosen-message existential forgery attack against PFDH that makes at most  $q_H$  queries to the hash oracle*

and at most  $q_S \leq 2^{k_0}$  queries to the signing oracle. Suppose the adversary succeeds in forging a signature for a fresh message with probability  $\epsilon$  within time  $t$ . Then, there exists an inverter  $\mathcal{I}$  that finds the preimage of an element uniformly drawn from the range of  $f$  with probability  $\epsilon'$  within time  $t'$ , where

$$\begin{aligned}\epsilon' &\geq \left(1 - \frac{1}{q_S}\right)^{q_S} \epsilon \approx \exp(-1) \epsilon \\ t' &\leq t + (q_H + q_S) O(t_f) + q_H q_S O(k_0)\end{aligned}$$

This means that when  $k_0 \geq \log_2(q_S)$ , the security of PFDH is tightly related to the problem of inverting the underlying trapdoor permutation. If a shorter salt is used, PFDH remains provably secure, but a tight reduction is not possible.

Katz and Wang [2003] describe yet another signature scheme that achieves a tight reduction using a single bit as random salt. The random salt can be removed altogether by computing this bit in a deterministic (but secret) way. This scheme uses the same key generation procedure as FDH and PSS; the signature and verification algorithms are as follows,

$$\begin{aligned}\text{Sign}(sk, m) &\stackrel{\text{def}}{=} \text{if } m \in \mathcal{S} \text{ then return } \mathcal{S}[m] \\ &\quad \text{else } b \xleftarrow{\$} \{0, 1\}; \sigma \leftarrow f^{-1}(sk, H(b \parallel m)); \mathcal{S}[m] \leftarrow \sigma; \\ &\quad \text{return } \sigma \\ \text{Verify}(pk, m, \sigma) &\stackrel{\text{def}}{=} \text{if } f(pk, \sigma) = H(0 \parallel m) \vee f(pk, \sigma) = H(1 \parallel m) \\ &\quad \text{then return true else return false}\end{aligned}$$

The signature algorithm is stateful: it will compute a signature for a message only once and return the same signature subsequently. To avoid maintaining state and remove the randomness, the bit  $b$  could be computed deterministically but in a secret way from message  $m$ . To avoid computing two hash values during verification, the bit  $b$  could be appended to the signature. We have the following result about the security of the above scheme,

**Theorem 5.6.** *Let  $\mathcal{A}$  be an adversary mounting a chosen-message existential forgery attack against the Katz-Wang scheme that makes at most  $q_H$  queries to the hash oracle and at most  $q_S$  queries to the signing oracle. Suppose the adversary succeeds in forging a signature for a fresh message with probability  $\epsilon$  within time  $t$ . Then, there exists an inverter  $\mathcal{I}$  that finds the preimage of an element uniformly drawn from the range of  $f$  with probability  $\epsilon'$  within time  $t'$ , where*

$$\begin{aligned}\epsilon' &\geq \frac{1}{2} \epsilon \\ t' &\leq t + (q_H + q_S + 1) O(t_f)\end{aligned}$$

Using just one *random* bit we have cut down the gap between the security of the scheme and the problem of inverting  $f$  by a factor of  $q_S$  !

The impossibility of finding a tight reduction from the security of a scheme to a hard problem does not necessarily mean that we should deem the scheme as insecure. It could be the case that the hard problem is not general enough, or

maybe a slightly modified version of the scheme would admit a tight reduction. Or perhaps the protocol is indeed secure for practical matters, even though a tight reduction does not exist. The absence of a tight reduction for RSA-FDH does not mean that its security is only loosely related to the RSA problem. The result about the Katz-Wang scheme described above suggests exactly the contrary. In fact, it is easy to see that the security of RSA-FDH is equivalent to the following problem:

You are given an RSA modulus  $N$  and its public exponent  $e$ , just as in the standard RSA problem. You are also given a set  $\{y_1, \dots, y_q\}$  of  $q$  values uniformly sampled from  $\mathbb{Z}_N$ . At any time you may choose a value  $y_i$  and get its  $e$ -th root modulo  $N$ , i.e. a solution  $x_i$  to  $x_i^e \equiv y_i \pmod{N}$ ; you may choose up to  $q_5$  of such values. Your goal is to compute the  $e$ -th root modulo  $N$  of one of the remaining  $y_i$  values.

It is hard to imagine how this problem could be any easier than solving the standard RSA problem.

Using `CryptoVerif` [Blanchet 2006], Blanchet and Pointcheval [2006] gave an alternative formal proof of the security of FDH against existential forgery under adaptive chosen-message attacks. This work has stirred considerable interest and shown the benefits of machine-checked verification. It also exposed one major weakness of `CryptoVerif`: it deviates from the style that is natural to cryptographers since it is difficult to recover a reductionist argument from the proof trace that the prover outputs, and even if one manages to do so, most likely the reduction will not be optimal. Indeed, only the original, suboptimal bound of Bellare and Rogaway has been proved in `CryptoVerif`.

Our machine-checked proofs follow quite closely the pen-and-paper game-based proofs of FDH (cf. [Pointcheval 2005]). There is however one important difference: in order to justify local transformations, machine-checked proofs must make invariants explicit and establish formally their validity. Proving that invariants hold constitutes a fair amount of work. More generally, machine-checked proofs must justify all reasoning, including reasoning about side conditions and about elementary mathematics (groups, probabilities) in terms of basic definitions. In contrast to game transformations, for which suitable tactics have been designed, this form of reasoning is not always amenable to automation, and thus accounts for a substantial amount of the effort and of the size of the proofs. Indeed, we estimate that about a third of the proof scripts are devoted to facts about probabilities. In spite of this, the size of machine-checked proofs remains reasonable: the formalizations of Bellare-Rogaway and Coron proofs are about 3,000 lines each. While the length of our proofs might look prohibitive in comparison to published proofs, we expect that machine-checked proofs will shrink substantially as `CertiCrypt` (and its underlying libraries) mature.

It must be noted that much of the proof lies outside of the trusted base: in order to trust the proofs of FDH, it is sufficient to trust our formalization of the scheme and of the security statement, the formalization of probabilistic programs provided by `CertiCrypt`, and the proof checker of `Coq`. In particular, trusting the proofs of FDH does not require trusting the sequence of games, nor the proofs of



transitions, nor the proofs of invariants. In this respect, CertiCrypt provides the highest possible level of assurance for the security of a cryptographic scheme, and breaks the symmetry between the effort of writing and checking a cryptographic proof. Both usually require a lot of expertise in cryptography, a lot of time, and a good understanding of the proof; in contrast, it is rather immediate and simple for a third party to check a proof in CertiCrypt.

# 6

## Ciphertext Indistinguishability of OAEP

WHEN we first discussed in Chapter 1 the definition of semantic security for public-key encryption schemes we observed that in order to achieve this notion of security an encryption scheme must be necessarily probabilistic. A deterministic asymmetric encryption scheme cannot be semantically secure because an adversary could trivially decide whether a given ciphertext is the encryption of a plaintext by encrypting the plaintext and comparing the resulting ciphertext to the one it was given. It is possible, however, to use a deterministic encryption scheme as a building block to construct a semantically secure probabilistic encryption scheme. The following definition gives a general way to construct a probabilistic encryption scheme from any family of trapdoor permutations.

**Definition 6.1 (Padding-based encryption scheme).** *Let  $\mu, \rho, k : \mathbb{N} \rightarrow \mathbb{N}$  be three functions such that  $\forall \eta. \mu(\eta) + \rho(\eta) \leq k(\eta)$ . Consider a family of mappings*

$$\begin{aligned} \pi_\eta &: \{0, 1\}^{\mu(\eta) + \rho(\eta)} \rightarrow \{0, 1\}^{k(\eta)} \\ \hat{\pi}_\eta &: \{0, 1\}^{k(\eta)} \rightarrow \{0, 1\}^{\mu(\eta)} \cup \{\perp\} \end{aligned}$$

such that  $\pi_\eta$  is injective and the following consistency condition is satisfied

$$\forall \eta. m \in \{0, 1\}^{\mu(\eta)}, r \in \{0, 1\}^{\rho(\eta)}. \hat{\pi}_\eta(\pi_\eta(m \parallel r)) = m$$

Then, given a family of trapdoor permutations  $(\mathcal{KG}_f, f, f^{-1})$  on  $\{0, 1\}^{k(\eta)}$ , one can construct a probabilistic padding-based encryption scheme  $(\mathcal{KG}, \mathcal{E}, \mathcal{D})$  as follows:

- Given  $\eta : \mathbb{N}$ , the key generation algorithm  $\mathcal{KG}(\eta)$  runs the key generation algorithm  $\mathcal{KG}_f(\eta)$  of the family of trapdoor permutations and returns the pair of keys  $(pk, sk)$  that it obtains as result. We assume that the description of  $(\pi_\eta, \hat{\pi}_\eta)$  is public, so there is no need to return it as part of the key generation process.
- Given a public key  $pk$  and a message  $m \in \{0, 1\}^{\mu(\eta)}$  the encryption algorithm  $\mathcal{E}(pk, m)$  chooses a uniformly random bitstring  $r \in \{0, 1\}^{\rho(\eta)}$  and returns a ciphertext computed as  $f(pk, \pi_\eta(m \parallel r)) \in \{0, 1\}^{k(\eta)}$ .

- Given a secret key  $sk$  and a ciphertext  $c \in \{0, 1\}^{k(n)}$ , the decryption algorithm  $\mathcal{D}(sk, c)$  returns  $m = \hat{\pi}_\eta(f^{-1}(sk, c))$ , which is either a message in  $\{0, 1\}^{\mu(n)}$  or the special symbol  $\perp$ , meaning that the ciphertext is invalid.


Many public-key encryption schemes can be viewed as particular instances of the above construction, including OAEP [Bellare and Rogaway 1994] and its variants OAEP<sup>+</sup> [Shoup 2001], SAEP, SAEP<sup>+</sup> [Boneh 2001], and REACT [Okamoto and Pointcheval 2001b]. In this chapter we will discuss the security of OAEP.

Optimal Asymmetric Encryption Padding (OAEP) [Bellare and Rogaway 1994] is a prominent public-key encryption scheme based on trapdoor permutations, most commonly used in combination with the RSA and Rabin functions. OAEP is widely deployed; many variants of OAEP are recommended by several standards, including IEEE P1363, PKCS, ISO 18033-2, ANSI X9, CRYPTREC and SET. Yet, the history of OAEP security is fraught with difficulties. The original paper of Bellare and Rogaway [1994] proves that, under the hypothesis that the underlying trapdoor permutation family is one-way, OAEP is semantically secure under chosen-ciphertext attacks. Shoup [2001] discovered later that this proof only established the security of OAEP against non-adaptive chosen-ciphertext attacks (IND-CCA), and not, as was believed at that time, against the stronger version of ciphertext indistinguishability that allows the adversary to adaptively obtain the decryption of ciphertexts of its choice (IND-CCA2). In response, Shoup suggested a modified scheme, secure against adaptive attacks under the one-wayness of the underlying permutation, and gave a proof of the adaptive security of the original scheme when it is used in combination with RSA with public exponent  $e = 3$ . Simultaneously, Fujisaki et al. [2004] proved that OAEP in its original formulation is indeed secure against adaptive attacks, but under the assumption that the underlying permutation family is *partial-domain* one-way. Since for the particular case of RSA this latter assumption is no stronger than (full-domain) one-wayness, this finally established the adaptive IND-CCA2 security of RSA-OAEP. Unfortunately, when one takes into account the additional cost of reducing the problem of inverting RSA to the problem of partially-inverting it, the security bound becomes less attractive. We note that there exist variants of OAEP that admit more efficient reductions when used in combination with the RSA and Rabin functions, notably SAEP, SAEP+ [Boneh 2001], and alternative schemes with tighter generic reductions, e.g. REACT [Okamoto and Pointcheval 2001b].

Here we focus on a machine-checked proof of the IND-CPA security of OAEP in the random oracle model, with a security bound that improves on the bound of Bellare and Rogaway [2006] game-based proof. We report as well on a significantly more challenging machine-checked proof of the IND-CCA2 security of OAEP in the random oracle model.

OAEP uses as a padding scheme a two-round Feistel network based on two hash functions  $G, H$ , and is commonly used together with the RSA and Rabin permutations. Figure 6.1 shows the Feistel network representation of the padding mappings  $(\pi, \hat{\pi})$  used in OAEP for encryption and decryption.

During encryption, a message  $m \in \{0, 1\}^\mu$  is first padded with enough zeros to obtain a bitstring of length  $k - \rho$ , which is then fed to the encryption Feistel



**Fig. 6.1.** Feistel network representation of the padding scheme used in OAEP for encryption (left) and decryption (right).

network together with a random bitstring  $r \in \{0, 1\}^\rho$ ; the ciphertext is the image under  $f$  of the resulting bitstring. To decrypt a ciphertext  $c$ , first compute its preimage  $f^{-1}(sk, c)$  to obtain a bitstring of length  $k$ , and then run it through the reverse Feistel network to obtain a bitstring  $(m \parallel z) \in \{0, 1\}^{k-\rho}$ . If  $z$  is not the all-zero bitstring, the ciphertext is rejected as invalid, otherwise,  $m$  is returned as its decryption. For concreteness, let us write down the definition of the OAEP scheme for a generic trapdoor permutation.

**Definition 6.2 (OAEP encryption scheme).** Let  $(\mathcal{KG}_f, f, f^{-1})$  be a family of trapdoor permutations on  $\{0, 1\}^k$ , and let

$$G : \{0, 1\}^\rho \rightarrow \{0, 1\}^{k-\rho} \quad H : \{0, 1\}^{k-\rho} \rightarrow \{0, 1\}^\rho$$

be two hash functions, with  $\rho + \mu < k$ . Let  $k_1 = k - \rho - \mu$ . The OAEP scheme is composed of the following triple of algorithms:

$$\begin{aligned} \mathcal{KG}(\eta) &\stackrel{\text{def}}{=} (pk, sk) \leftarrow \mathcal{KG}_f(\eta); \text{ return } (pk, sk) \\ \mathcal{E}(pk, m) &\stackrel{\text{def}}{=} r \xleftarrow{\$} \{0, 1\}^\rho; s \leftarrow G(r) \oplus (m \parallel 0^{k_1}); t \leftarrow H(s) \oplus r; \text{ return } f(pk, s \parallel t) \\ \mathcal{D}(sk, c) &\stackrel{\text{def}}{=} (s \parallel t) \leftarrow f^{-1}(sk, c); r \leftarrow t \oplus H(s); m \leftarrow s \oplus G(r); \\ &\quad \text{if } [m]_{k_1} = 0^{k_1} \text{ then return } [m]^\mu \text{ else return } \perp \end{aligned}$$

where  $[x]_n$  (resp.  $[x]^n$ ) denotes the  $n$  least (resp. most) significant bits of  $x$ .

The way padding is handled plays a crucial role in the proof of IND-CCA2 security of the scheme in the random oracle model that we describe in Section 6.2. To prove that OAEP is IND-CCA2 secure, it is necessary to devise an efficient way to simulate the decryption oracle without having the trapdoor to the underlying permutation. It turns out that there is an efficient way to simulate the decryption oracle from the history of queries that an adversary made to the hash oracles  $G$  and  $H$ . Suppose the adversary queries the decryption oracle with a ciphertext  $c = f(s \parallel t)$ , then there are two possibilities:

1. Either  $s$  has been queried to  $H$  and  $r = H(s) \oplus t$  has been queried to  $G$ . In that case, the corresponding plaintext can be found by inspecting the list of queries  $(r', G(r'))$ ,  $(s', H(s'))$  made to  $G$  and  $H$  respectively. For each pair of queries, check whether  $c$  coincides with  $f(pk, s' \parallel (H(s') \oplus r'))$  and the  $(k - \rho - \mu)$  least significant bits of  $G(r') \oplus s'$  are all zeros. If one such pair of queries is found, the  $\mu$  most significant bits of  $G(r') \oplus s'$  form the corresponding plaintext;
2. Or else there is only a minute chance that the ciphertext be valid because it depends on a uniformly random value (either one of  $G(r)$  and  $H(s)$ , or both). If the permutation is partial-domain one-way, the ciphertext may be safely rejected as invalid.

### 6.1 Indistinguishability under Chosen-Plaintext Attacks

Unsurprisingly, padding does not play any role in the proof of semantic security, because there is no need to simulate the decryption oracle. Therefore, in the rest of this section we assume without loss of generality that  $\mu = k - \rho$ .

**Theorem 6.3 (IND-CPA security of OAEP).** *Let  $\mathcal{A}$  be an adversary against the semantic security of OAEP under chosen-message attacks that makes at most  $q_G$  queries to the hash oracle  $G$  and at most  $q_H$  queries to  $H$ . Suppose this adversary succeeds in guessing the hidden bit  $b$  with probability  $\epsilon$  within time  $t$ . Then, there exists an inverter  $\mathcal{I}$  that finds the preimage of an element uniformly drawn from the domain of the permutation  $f$  with probability  $\epsilon'$  within time  $t'$ , where*

$$\epsilon' \geq \epsilon - \left( \frac{1}{2} + \frac{q_G}{2^\rho} \right) \quad (6.1)$$

$$t' \leq t + q_G q_H O(t_f) \quad (6.2)$$

and  $t_f$  is an upper bound for the time needed to compute the image of a bitstring under  $f$ .

*Proof.* We claim that the inverter  $\mathcal{I}$  shown in the context of game  $G_{OW}$  in Figure 6.2 achieves the probability and time bounds in (6.1) and (6.2). The inverter runs the adversary  $\mathcal{A}$  in a simulated environment. It intercepts queries to the hash oracles  $G$  and  $H$  and answers exactly as a random oracle would do, but keeps record of the queries and their responses. Instead of computing the challenge ciphertext to the IND-CPA adversary as an LR-oracle would do in the real IND-CPA game, it replaces the challenge ciphertext with its own challenge  $y$ . When adversary  $\mathcal{A}$  halts, the inverter  $\mathcal{I}$  inspects the history of queries that were made to the hash oracles  $G$  and  $H$  and tries to reconstruct a preimage of its challenge  $y$  under  $f$ . It turns out that the probability that the inverter succeeds in reconstructing such a preimage is closely related to the probability with which adversary  $\mathcal{A}$  wins the IND-CPA game.

The upper bound on the execution time of the inverter can be justified by examining just the game  $G_{OW}$ . The only non-constant overhead incurred is in the

<p><b>Game <math>G_{OW}</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>y \xleftarrow{\\$} \{0, 1\}^k;</math>  <math>x \leftarrow \mathcal{I}(pk, y)</math></p> <p><b>Adversary <math>\mathcal{I}(pk, y)</math> :</b>  <math>L, M \leftarrow \text{nil};</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y);</math>          if <math>\exists r \in \text{dom}(L), (s, h) \in M.</math>              <math>f(pk, s \parallel h \oplus r) = y</math>          then return <math>s \parallel (h \oplus r)</math>          else return <math>0^k</math></p>	<p><b>Oracle <math>G(r)</math> :</b>          if <math>r \notin \text{dom}(L)</math> then              <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>              <math>L \leftarrow (r, g) :: L</math>          return <math>L[r]</math></p> <p><b>Oracle <math>H(s)</math> :</b>          if <math>s \notin \text{dom}(M)</math> then              <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>              <math>M \leftarrow (s, h) :: M</math>          return <math>M[s]</math></p>
--	---

**Fig. 6.2.** The inverter  $\mathcal{I}$  in the context of the one-wayness game for the family of trapdoor permutations  $(\mathcal{KG}_f, f, f^{-1})$ .

reconstruction of the preimage of the challenge by inspecting the history of queries. In the worst case, it amounts to  $q_H$  computations of the permutation  $f$  for each of the  $q_G$  queries, from which (6.2) follows directly.

To prove that the inverter achieves the probability bound (6.1), we will exhibit a sequence of games relating the probability

$$\epsilon' = \Pr [G_{OW} : x = f^{-1}(sk, y)]$$

of  $\mathcal{I}$  successfully inverting  $f$  on its challenge  $y$ , to the probability

$$\epsilon = \Pr [G_{INDCPA} : b = \tilde{b}]$$

of the adversary  $\mathcal{A}$  correctly guessing the value of the hidden bit  $b$  in the IND-CPA game. We start from the IND-CPA game where oracles  $G$  and  $H$  are implemented as random oracles—this is justified by the random oracle model.

<p><b>Game <math>G_{INDCPA}</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}();</math>  <math>L, M \leftarrow \text{nil};</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>b \xleftarrow{\\$} \{0, 1\};</math>  <math>y \leftarrow \mathcal{E}(pk, m_b);</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle <math>G(r)</math> :</b>          if <math>r \notin \text{dom}(L)</math> then              <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>              <math>L \leftarrow (r, g) :: L</math>          return <math>L[r]</math></p> <p><b>Oracle <math>H(s)</math> :</b>          if <math>s \notin \text{dom}(M)</math> then              <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>              <math>M \leftarrow (s, h) :: M</math>          return <math>M[s]</math></p>
---	---

The hypothesis on the bound on the number of queries the adversary makes to  $G$  can be readily encoded as:

$$\Pr[\mathbf{G}_{\text{INDCPA}} : |L| \leq q_G] = 1$$

In the following game, we inline the key generation and the encryption of  $m_b$ , which uses a random seed  $\hat{r}$ , and we eagerly sample the response that the random oracle  $G$  gives back if  $\hat{r}$  is ever queried:

<p><b>Game <math>G_1</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>L, M \leftarrow \text{nil};</math>  <math>\hat{r} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{g} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>b \xleftarrow{\\$} \{0, 1\};</math>  <math>g \leftarrow G(\hat{r}); s \leftarrow g \oplus m_b;</math>  <math>h \leftarrow H(s); t \leftarrow h \oplus r;</math>  <math>y \leftarrow f(pk, s \parallel t);</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle <math>G(r)</math> :</b>  if <math>r \notin \text{dom}(L)</math> then    if <math>r = \hat{r}</math> then <math>g \leftarrow \hat{g}</math>    else <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho}</math>  <math>L \leftarrow (r, g) :: L</math>  return <math>L[r]</math></p> <p><b>Oracle <math>H(s)</math> :</b>  if <math>s \notin \text{dom}(M)</math> then    <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>M \leftarrow (s, h) :: M</math>  return <math>M[s]</math></p>
--	---

The resulting game  $G_1$  is semantically equivalent to the initial game, and thus we have

$$\Pr[\mathbf{G}_{\text{INDCPA}} : b = \tilde{b}] = \Pr[G_1 : b = \tilde{b}] \quad (6.3)$$

Our objective is now to eliminate  $\hat{g}$  from the code of the oracle  $G$ , because if we manage to do so, we will be able to make  $s$  completely random and remove the dependency of  $y$  on  $b$ .

In the next game we modify the oracle  $G$  so that if  $\hat{r}$  is ever queried, a flag **bad** is set to true and the answer is not recorded in the oracle memory.

<p><b>Game <math>G_2</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>L, M \leftarrow \text{nil};</math>  <math>\hat{r} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{g} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>b \xleftarrow{\\$} \{0, 1\};</math>  <math>s \leftarrow \hat{g} \oplus m_b;</math>  <math>h \leftarrow H(s); t \leftarrow h \oplus \hat{r};</math>  <math>y \leftarrow f(pk, s \parallel t);</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle <math>G(r)</math> :</b>  if <math>r = \hat{r}</math> then    <b>bad</b> <math>\leftarrow</math> true; return <math>\hat{g}</math>  else    if <math>r \notin \text{dom}(L)</math> then      <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>      <math>L \leftarrow (r, g) :: L</math>    return <math>L[r]</math></p> <p><b>Oracle <math>H(s)</math> :</b>  if <math>s \notin \text{dom}(M)</math> then    <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>    <math>M \leftarrow (s, h) :: M</math>  return <math>M[s]</math></p>
---	---

In order to justify this transformation, we define the following relational invariant:

$$\phi_{12} \stackrel{\text{def}}{=} (\hat{r} \in \text{dom}(\mathbf{L}) \implies \mathbf{L}[\hat{r}] = \hat{g}\langle 1 \rangle \wedge (\forall r. r \neq \hat{r}\langle 1 \rangle \implies \mathbf{L}[r]\langle 1 \rangle = \mathbf{L}[r]\langle 2 \rangle)$$

The above invariant allows us to prove that oracle  $G$  answers queries in the same way in games  $G_1$  and  $G_2$ . We prove that this invariant is established after  $\mathbf{L}$  is initialized in  $G_1$  and  $G_2$ , and that the implementations of  $G$  and  $H$  preserve it. We know from the fact that  $\mathcal{A}$  is well-formed that it cannot directly modify the global variables  $\mathbf{L}$ ,  $\hat{r}$ , and  $\hat{g}$ . Therefore, the invariant  $\phi_{12}$  is preserved through calls to the adversary. We then inline the call to  $G$  in the body of game  $G_1$ . At this point the invariant holds and either the adversary  $\mathcal{A}$  has already queried  $\hat{r}$ , in which case it follows from the invariant that the answer to the call is  $\hat{g}$ , or it has never queried  $\hat{r}$ , in which case the same follows directly by applying dead code elimination and constant propagation. Hence,

$$\Pr [G_1 : b = \tilde{b}] = \Pr [G_2 : b = \tilde{b}] \quad (6.4)$$

We will now use the Fundamental Lemma to remove  $\hat{g}$  altogether from the code of oracle  $G$ . We define a game  $G_3$  syntactically identical up to **bad** to the previous game.

<p><b>Game <math>G_3</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>\mathbf{L}, \mathbf{M} \leftarrow \text{nil};</math>  <math>\hat{r} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{g} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>b \xleftarrow{\\$} \{0, 1\};</math>  <math>s \leftarrow \hat{g} \oplus m_b;</math>  <math>h \leftarrow H(s); t \leftarrow h \oplus \hat{r};</math>  <math>y \leftarrow f(pk, s \parallel t);</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle <math>G(r)</math> :</b>          if <math>r = \hat{r}</math> then            <b>bad</b> <math>\leftarrow</math> true;            if <math>r \notin \text{dom}(\mathbf{L})</math> then              <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>              <math>\mathbf{L} \leftarrow (r, g) :: \mathbf{L}</math>              return <math>\mathbf{L}[r]</math>            else              if <math>r \notin \text{dom}(\mathbf{L})</math> then                <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>                <math>\mathbf{L} \leftarrow (r, g) :: \mathbf{L}</math>                return <math>\mathbf{L}[r]</math>            return <math>\mathbf{L}[r]</math>  <b>Oracle <math>H(s)</math> :</b> ...</p>
---	--

It follows from the Fundamental Lemma that

$$\left| \Pr [G_2 : b = \tilde{b}] - \Pr [G_3 : b = \tilde{b}] \right| \leq \Pr [G_3 : \text{bad}] \quad (6.5)$$

Since  $\hat{g}$  no longer appears in  $G$ , we can now sample it later in the game. We also simplify the implementation of oracle  $G$  for the sake of readability by coalescing the portion of code appearing in both branches of the conditional.



<b>Game <math>G_4</math> :</b> $(pk, sk) \leftarrow \mathcal{KG}_f();$ $L, M \leftarrow \text{nil};$ $\hat{r} \xleftarrow{\$} \{0, 1\}^\rho;$ $(m_0, m_1) \leftarrow \mathcal{A}_1(pk);$ $b \xleftarrow{\$} \{0, 1\};$ $\hat{g} \xleftarrow{\$} \{0, 1\}^{k-\rho};$ $s \leftarrow \hat{g} \oplus m_b;$ $h \leftarrow H(s); t \leftarrow h \oplus \hat{r};$ $y \leftarrow f(pk, s \parallel t);$ $\tilde{b} \leftarrow \mathcal{A}_2(y)$	<b>Oracle <math>G(r)</math> :</b> if $r = \hat{r}$ then $\text{bad} \leftarrow \text{true}$ if $r \notin \text{dom}(L)$ then $g \xleftarrow{\$} \{0, 1\}^{k-\rho};$ $L \leftarrow (r, g) :: L$ return $L[r]$  <b>Oracle <math>H(s)</math> :</b> if $s \notin \text{dom}(M)$ then $h \xleftarrow{\$} \{0, 1\}^\rho;$ $M \leftarrow (s, h) :: M$ return $M[s]$
---	---

We have

$$\Pr [G_3 : b = \tilde{b}] = \Pr [G_4 : b = \tilde{b}] \quad (6.6)$$

$$\Pr [G_3 : \text{bad}] = \Pr [G_4 : \text{bad}] \quad (6.7)$$

We make now a transformation that relies on algebraic properties of the exclusive or operator: instead of sampling  $\hat{g}$  and defining  $s$  in terms of it, we can sample  $s$  and define  $\hat{g}$  in terms of  $s$ . This is justified by the following program equivalence:

$$\vdash \hat{g} \xleftarrow{\$} \{0, 1\}^{k-\rho}; s \leftarrow \hat{g} \oplus m_b \simeq_{\{\hat{g}, s, m_b\}}^{\{m_b\}} s \xleftarrow{\$} \{0, 1\}^{k-\rho}; \hat{g} \leftarrow s \oplus m_b$$

This transformation is sometimes called *optimistic sampling* and is a pattern that appears recurrently in game-based proofs; we gave a proof using the relational Hoare logic in Section 3.2.2.

We can now eliminate the assignment to  $\hat{g}$  as dead code, and sample  $s$  at the beginning of the game. Since  $y$  no longer depends on the hidden bit  $b$ , we can sample  $b$  at the end of the game. The resulting game  $G_5$  is:

<b>Game <math>G_5</math> :</b> $(pk, sk) \leftarrow \mathcal{KG}_f();$ $L, M \leftarrow \text{nil};$ $\hat{r} \xleftarrow{\$} \{0, 1\}^\rho;$ $s \xleftarrow{\$} \{0, 1\}^{k-\rho};$ $(m_0, m_1) \leftarrow \mathcal{A}_1(pk);$ $h \leftarrow H(s); t \leftarrow h \oplus \hat{r};$ $y \leftarrow f(pk, s \parallel t);$ $\tilde{b} \leftarrow \mathcal{A}_2(y);$ $b \xleftarrow{\$} \{0, 1\}$	<b>Oracle <math>G(r)</math> :</b> if $r = \hat{r}$ then $\text{bad} \leftarrow \text{true}$ if $r \notin \text{dom}(L)$ then $g \xleftarrow{\$} \{0, 1\}^{k-\rho};$ $L \leftarrow (r, g) :: L$ return $L[r]$  <b>Oracle <math>H(s)</math> :</b> if $s \notin \text{dom}(M)$ then $h \xleftarrow{\$} \{0, 1\}^\rho;$ $M \leftarrow (s, h) :: M$ return $M[s]$
---	---

We have  $\Pr [G_4 : \text{bad}] = \Pr [G_5 : \text{bad}]$ . It is obvious that, in the above game,  $\tilde{b}$  and  $b$  are independent and thus

$$\Pr [\mathbf{G}_4 : b = \tilde{b}] = \Pr [\mathbf{G}_5 : b = \tilde{b}] = \frac{1}{2} \quad (6.8)$$

The objective now is to bound the probability of **bad** being set. We no longer care about  $b$ , so we remove the instruction sampling it. We reallocate  $s$  to a global variable  $\hat{s}$  and we eagerly sample the value of  $H(\hat{s})$ .

<p><b>Game <math>\mathbf{G}_6</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>\mathbf{L}, \mathbf{M} \leftarrow \text{nil};</math>  <math>\hat{r} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{s} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>\hat{h} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>h \leftarrow H(\hat{s}); t \leftarrow h \oplus \hat{r};</math>  <math>y \leftarrow f(pk, \hat{s} \parallel t);</math>  <math>b \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle <math>G(r)</math> :</b>  if <math>r = \hat{r}</math> then <b>bad</b> <math>\leftarrow</math> true  if <math>r \notin \text{dom}(\mathbf{L})</math> then  <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>\mathbf{L} \leftarrow (r, g) :: \mathbf{L}</math>  return <math>\mathbf{L}[r]</math></p> <p><b>Oracle <math>H(s)</math> :</b>  if <math>s \notin \text{dom}(\mathbf{M})</math> then  if <math>s = \hat{s}</math> then <math>h \leftarrow \hat{h}</math>  else <math>h \xleftarrow{\\$} \{0, 1\}^\rho</math>  <math>\mathbf{M} \leftarrow (s, h) :: \mathbf{M}</math>  return <math>\mathbf{M}[s]</math></p>
--	---

We then modify oracle  $H$  so that it does not record the answer in its memory if  $\hat{s}$  is ever queried, and we inline the call to  $H$  in the body of the game. To prove that this modification to  $H$  is transparent to the adversary we prove the following relational invariant between  $\mathbf{G}_6$  and  $\mathbf{G}_7$ :

$$(\hat{s} \in \text{dom}(\mathbf{M}) \implies \mathbf{M}[\hat{s}] = \hat{h}) \langle 1 \rangle \wedge (\forall s. s \neq \hat{s} \langle 1 \rangle \implies \mathbf{M}[s] \langle 1 \rangle = \mathbf{M}[s] \langle 2 \rangle)$$

<p><b>Game <math>\mathbf{G}_7</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>\mathbf{L}, \mathbf{M} \leftarrow \text{nil};</math>  <math>\hat{r} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{s} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>\hat{h} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>t \leftarrow \hat{h} \oplus \hat{r};</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>y \leftarrow f(pk, \hat{s} \parallel t);</math>  <math>b \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle <math>G(r)</math> :</b>  if <math>r = \hat{r}</math> then <b>bad</b> <math>\leftarrow</math> true  if <math>r \notin \text{dom}(\mathbf{L})</math> then  <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>\mathbf{L} \leftarrow (r, g) :: \mathbf{L}</math>  return <math>\mathbf{L}[r]</math></p> <p><b>Oracle <math>H(s)</math> :</b>  if <math>s = \hat{s}</math> then return <math>\hat{h}</math>  else if <math>s \notin \text{dom}(\mathbf{M})</math> then  <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\mathbf{M} \leftarrow (s, h) :: \mathbf{M}</math>  return <math>\mathbf{M}[s]</math></p>
---	--

Thus

$$\Pr [\mathbf{G}_5 : \mathbf{bad}] = \Pr [\mathbf{G}_6 : \mathbf{bad}] = \Pr [\mathbf{G}_7 : \mathbf{bad}] \quad (6.9)$$

We then revert to the previous implementation of oracle  $H$ , which stores the answer to a query  $\hat{s}$  in its memory. This allows us to bound the probability of **bad** being set by analyzing two different cases:

1. either the adversary queried  $\hat{s}$  to  $H$  before **bad** is set, in which case we set a flag **bad**<sub>1</sub>;
2. or it did not, in which case we set a flag **bad**<sub>2</sub>.

<p><b>Game G<sub>8</sub> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>L, M \leftarrow \text{nil};</math>  <math>\hat{r} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{s} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>\hat{h} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>t \leftarrow \hat{h} \oplus \hat{r};</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>y \leftarrow f(pk, \hat{s} \parallel t);</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle G(r) :</b>  if <math>r = \hat{r}</math> then    if <math>\hat{s} \in \text{dom}(M)</math>    then <b>bad</b><sub>1</sub> <math>\leftarrow</math> true    else <b>bad</b><sub>2</sub> <math>\leftarrow</math> true  ...  <b>Oracle H(s) :</b>  if <math>s \notin \text{dom}(M)</math> then    if <math>s = \hat{s}</math> then <math>h \leftarrow \hat{h}</math>    else <math>h \xleftarrow{\\$} \{0, 1\}^\rho</math>    <math>M \leftarrow (s, h) :: M</math>  return <math>M[s]</math></p>
---	---

We prove the following relational invariant between G<sub>7</sub> and G<sub>8</sub>:

$$\mathbf{bad}\langle 1 \rangle \implies (\mathbf{bad}_1 \vee \mathbf{bad}_2)\langle 2 \rangle$$

Hence we have by the union bound,

$$\Pr[G_7 : \mathbf{bad}] \leq \Pr[G_8 : \mathbf{bad}_1 \vee \mathbf{bad}_2] \leq \Pr[G_8 : \mathbf{bad}_1] + \Pr[G_8 : \mathbf{bad}_2] \quad (6.10)$$

We split the sequence of games and bound **bad**<sub>1</sub> and **bad**<sub>2</sub> separately; we deal with **bad**<sub>1</sub> first.

We slice the assignment to **bad**<sub>2</sub> off the code of  $G$ , we apply again the optimistic sampling transformation to sample  $t$  instead of  $\hat{r}$ , and we reallocate  $t$ ,  $y$  and  $pk$  to global variables, obtaining:

<p><b>Game P<sub>1</sub> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>L, M \leftarrow \text{nil};</math>  <math>\hat{s} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>\hat{t} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{h} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{r} \leftarrow \hat{h} \oplus \hat{t};</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>y \leftarrow f(pk, \hat{s} \parallel \hat{t});</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle G(r) :</b>  if <math>r = \hat{r} \wedge \hat{s} \in \text{dom}(M)</math> then    <b>bad</b><sub>1</sub> <math>\leftarrow</math> true  if <math>r \notin \text{dom}(L)</math> then    <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>    <math>L \leftarrow (r, g) :: L</math>  return <math>L[r]</math>  <b>Oracle H(s) :</b>  if <math>s \notin \text{dom}(M)</math> then    if <math>s = \hat{s}</math> then <math>h \leftarrow \hat{h}</math>    else <math>h \xleftarrow{\\$} \{0, 1\}^\rho</math>    <math>M \leftarrow (s, h) :: M</math>  return <math>M[s]</math></p>
---	---

We now replace the condition under which  $\mathbf{bad}_1$  is set to true by an equivalent one. We prove that  $P_1$  satisfies the invariant

$$(\forall (s, h) \in M. s = \hat{s} \implies h = \hat{h}) \wedge (\hat{r} = \hat{h} \oplus \hat{t} \wedge \mathbf{y} = f(\mathbf{pk}, \hat{s} \parallel \hat{t}))$$

From this invariant we have that

$$r = \hat{r} \wedge \hat{s} \in \text{dom}(M) \iff \exists (s, h) \in M. f(\mathbf{pk}, s \parallel (h \oplus r)) = \mathbf{y}$$

Therefore we can reformulate the condition under which  $\mathbf{bad}_1$  is set in  $G$ , and remove  $\hat{r}$  since it is no longer used.

<p><b>Game <math>P_2</math> :</b>  <math>(\mathbf{pk}, sk) \leftarrow \mathcal{KG}_f();</math>  <math>L, M \leftarrow \text{nil};</math>  <math>\hat{s} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>\hat{t} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{h} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(\mathbf{pk});</math>  <math>\mathbf{y} \leftarrow f(\mathbf{pk}, \hat{s} \parallel \hat{t});</math>  <math>\hat{b} \leftarrow \mathcal{A}_2(\mathbf{y})</math></p>	<p><b>Oracle <math>G(r)</math> :</b>          if <math>\exists (s, h) \in M.</math>            <math>f(\mathbf{pk}, s \parallel h \oplus r) = \mathbf{y}</math> then            <math>\mathbf{bad}_1 \leftarrow \text{true}</math>          if <math>r \notin \text{dom}(L)</math> then            <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>            <math>L \leftarrow (r, g) :: L</math>            return <math>L[r]</math>  <b>Oracle <math>H(s)</math> :</b>          if <math>s \notin \text{dom}(M)</math> then            if <math>s = \hat{s}</math> then <math>h \leftarrow \hat{h}</math>            else <math>h \xleftarrow{\\$} \{0, 1\}^\rho</math>            <math>M \leftarrow (s, h) :: M</math>            return <math>M[s]</math></p>
--	---

We have

$$\Pr[G_8 : \mathbf{bad}_1] = \Pr[P_1 : \mathbf{bad}_1] = \Pr[P_2 : \mathbf{bad}_1] \quad (6.11)$$

Next, we revert  $H$  to the original random oracle implementation and we slice away  $\hat{h}$  as dead code. In the resulting game  $\hat{s}$  and  $\hat{t}$  are used only to compute  $\mathbf{y} = f(\mathbf{pk}, \hat{s} \parallel \hat{t})$ . Since  $\hat{s}$  and  $\hat{t}$  are uniformly sampled bitstrings, their concatenation is uniformly distributed, and since  $f$  is a permutation, the value of  $\mathbf{y}$  is uniformly distributed. This reasoning is summarized in the following program equivalence

$$\vdash \hat{s} \xleftarrow{\$} \{0, 1\}^{k-\rho}; \hat{t} \xleftarrow{\$} \{0, 1\}^\rho; \mathbf{y} \leftarrow f(\mathbf{pk}, \hat{s} \parallel \hat{t}) \simeq_{\{\mathbf{pk}\}}^{\{\mathbf{pk}, \mathbf{y}\}} \mathbf{y} \xleftarrow{\$} \{0, 1\}^k$$

<p><b>Game P<sub>3</sub> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>L, M \leftarrow \text{nil};</math>  <math>y \xleftarrow{\\$} \{0, 1\}^k;</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle <math>G(r)</math> :</b>  if <math>\exists(s, h) \in M.</math>  <math>f(pk, s \parallel h \oplus r) = y</math> then  <math>\mathbf{bad}_1 \leftarrow \text{true}</math>  if <math>r \notin \text{dom}(L)</math> then  <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>L \leftarrow (r, g) :: L</math>  return <math>L[r]</math></p> <p><b>Oracle <math>H(s)</math> :</b>  if <math>s \notin \text{dom}(M)</math> then  <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>M \leftarrow (s, h) :: M</math>  return <math>M[s]</math></p>
--	--

We now revert  $G$  to its original implementation and prove the following relational invariant between  $P_3$  and  $P_4$  which gives a necessary condition for  $\mathbf{bad}_1$  to be set:

$$\mathbf{bad}_1 \langle 1 \rangle \implies (\exists r \in \text{dom}(L), (s, h) \in M. f(pk, s \parallel (h \oplus r)) = y) \langle 2 \rangle$$

Thus,

$$\Pr[P_3 : \mathbf{bad}_1] \leq \Pr[P_4 : \exists r \in \text{dom}(L), (s, h) \in M. f(pk, s \parallel (h \oplus r)) = y] \quad (6.12)$$

<p><b>Game P<sub>4</sub> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>L, M \leftarrow \text{nil};</math>  <math>y \xleftarrow{\\$} \{0, 1\}^k;</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y)</math></p>	<p><b>Oracle <math>G(r)</math> :</b>  if <math>r \notin \text{dom}(L)</math> then  <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>L \leftarrow (r, g) :: L</math>  return <math>L[r]</math></p> <p><b>Oracle <math>H(s)</math> :</b>  if <math>s \notin \text{dom}(M)</math> then  <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>M \leftarrow (s, h) :: M</math>  return <math>M[s]</math></p>
--	---

We reallocate variables  $pk$  and  $y$  to local variables, and at the end of the game compute  $x$  as the inverter in Figure 6.2 would do.

<b>Game <math>P_5</math> :</b> $(pk, sk) \leftarrow \mathcal{KG}_f();$ $y \xleftarrow{\$} \{0, 1\}^k;$ $\mathbf{L}, \mathbf{M} \leftarrow \text{nil};$ $(m_0, m_1) \leftarrow \mathcal{A}_1(pk);$ $\tilde{b} \leftarrow \mathcal{A}_2(y);$ if $\exists r \in \text{dom}(\mathbf{L}), (s, h) \in \mathbf{M}.$ $f(pk, s \parallel h \oplus r) = y$ then $x \leftarrow s \parallel (h \oplus r)$ else $x \leftarrow 0^k$	<b>Oracle <math>G(r)</math> :</b> if $r \notin \text{dom}(\mathbf{L})$ then $g \xleftarrow{\$} \{0, 1\}^{k-\rho};$ $\mathbf{L} \leftarrow (r, g) :: \mathbf{L}$ return $\mathbf{L}[r]$  <b>Oracle <math>H(s)</math> :</b> if $s \notin \text{dom}(\mathbf{M})$ then $h \xleftarrow{\$} \{0, 1\}^\rho;$ $\mathbf{M} \leftarrow (s, h) :: \mathbf{M}$ return $\mathbf{M}[s]$
--	--

Compare this last game to game  $G_{OW}$  in Figure 6.2: both games semantically equivalent once the call to  $\mathcal{I}$  in  $G_{OW}$  is inlined. Furthermore,

$$\begin{aligned} \Pr [P_4 : \exists r \in \text{dom}(\mathbf{L}), (s, h) \in \mathbf{M}. f(pk, s \parallel (h \oplus r)) = y] \\ \leq \Pr [P_5 : x = f^{-1}(sk, y)] \\ = \Pr [G_{OW} : x = f^{-1}(sk, y)] \end{aligned} \quad (6.13)$$

From (6.11)–(6.13) we get

$$\Pr [G_8 : \mathbf{bad}_1] \leq \Pr [G_{OW} : x = f^{-1}(sk, y)] \quad (6.14)$$

We are now left with the problem of bounding  $\mathbf{bad}_2$  in game  $G_8$ . Define the game  $Q$  parametrized by two instructions  $c_1$  and  $c_2$  as follows

<b>Game <math>Q(c_1, c_2)</math> :</b> $(pk, sk) \leftarrow \mathcal{KG}_f();$ $\mathbf{L}, \mathbf{M} \leftarrow \text{nil};$ $\hat{r} \xleftarrow{\$} \{0, 1\}^\rho;$ $\hat{s} \xleftarrow{\$} \{0, 1\}^{k-\rho};$ $\hat{h} \xleftarrow{\$} \{0, 1\}^\rho;$ $t \leftarrow \hat{h} \oplus \hat{r};$ $(m_0, m_1) \leftarrow \mathcal{A}_1(pk);$ $y \leftarrow f(pk, \hat{s} \parallel t);$ $\tilde{b} \leftarrow \mathcal{A}_2(y)$	<b>Oracle <math>G(r)</math> :</b> if $r = \hat{r} \wedge \hat{s} \notin \text{dom}(\mathbf{M})$ then $\mathbf{bad}_2 \leftarrow \text{true}$ ...  <b>Oracle <math>H(s)</math> :</b> if $s \notin \text{dom}(\mathbf{M})$ then $h \xleftarrow{\$} \{0, 1\}^\rho;$ if $s = \hat{s} \wedge \mathbf{bad}_2$ then $\mathbf{bad}_2, \mathbf{bad}_3 \leftarrow \text{true}; c_1$ if $s = \hat{s} \wedge \neg \mathbf{bad}_2$ then $\mathbf{bad}_4 \leftarrow \text{true}; c_2$ $\mathbf{M} \leftarrow (s, h) :: \mathbf{M}$ return $\mathbf{M}[s]$
---	--

Flag  $\mathbf{bad}_3$  is set when  $\hat{s}$  is first queried to  $H$  and  $G(\hat{r})$  has already been queried. Flag  $\mathbf{bad}_4$  is set when  $\hat{s}$  is first queried to  $H$  and  $G(\hat{r})$  has not yet been queried. Consider the games

$$\begin{aligned} Q_1 &\stackrel{\text{def}}{=} Q(h \leftarrow \hat{h}, h \leftarrow \hat{h}) \\ Q_2 &\stackrel{\text{def}}{=} Q(h \xleftarrow{\$} \{0, 1\}^\rho, h \leftarrow \hat{h}) \\ Q_3 &\stackrel{\text{def}}{=} Q(h \xleftarrow{\$} \{0, 1\}^\rho, h \xleftarrow{\$} \{0, 1\}^\rho) \end{aligned}$$

Using the fundamental lemma together with appropriate program invariants, we obtain

$$\begin{aligned}
\Pr[G_8 : \mathbf{bad}_2] &= \Pr[Q_1 : \mathbf{bad}_2] \\
&= \Pr[Q_2 : \mathbf{bad}_2] && (\text{FL, } \mathbf{bad}_2) \\
&= \Pr[Q_2 : \mathbf{bad}_2 \wedge \mathbf{bad}_4] + \Pr[Q_2 : \mathbf{bad}_2 \wedge \neg \mathbf{bad}_4] \\
&= \Pr[Q_2 : \mathbf{bad}_2 \wedge \mathbf{bad}_4] + \Pr[Q_3 : \mathbf{bad}_2 \wedge \neg \mathbf{bad}_4] && (\text{FL, } \mathbf{bad}_4) \\
&= 0 && + \Pr[Q_3 : \mathbf{bad}_2 \wedge \neg \mathbf{bad}_4] \quad (\mathbf{bad}_4 \implies \neg \mathbf{bad}_2) \\
&= \Pr[Q_3 : \mathbf{bad}_2 \wedge \mathbf{bad}_4] + \Pr[Q_3 : \mathbf{bad}_2 \wedge \neg \mathbf{bad}_4] \quad (\mathbf{bad}_4 \implies \neg \mathbf{bad}_2) \\
&= \Pr[Q_3 : \mathbf{bad}_2]
\end{aligned}$$

We use optimistic sampling and the fact that  $f$  is a permutation to prove the following equivalence:

$$\begin{array}{lll}
\hat{r} \xleftarrow{\$} \{0, 1\}^\rho; & \hat{r} \xleftarrow{\$} \{0, 1\}^\rho; & \\
\hat{s} \xleftarrow{\$} \{0, 1\}^{k-\rho}; & \hat{s} \xleftarrow{\$} \{0, 1\}^{k-\rho}; & \\
\hat{h} \xleftarrow{\$} \{0, 1\}^\rho; & \simeq_{\{\hat{r}, y\}}^\emptyset t \xleftarrow{\$} \{0, 1\}^\rho; & \simeq_{\{\hat{r}, y\}}^\emptyset \hat{r} \xleftarrow{\$} \{0, 1\}^\rho; \\
t \leftarrow \hat{h} \oplus \hat{r}; & \hat{h} \leftarrow t \oplus \hat{r}; & y \xleftarrow{\$} \{0, 1\}^k; \\
y \leftarrow f(pk, \hat{s} \parallel t) & y \leftarrow f(pk, \hat{s} \parallel t) &
\end{array}$$

In the next game, we apply the above equivalence and we revert oracles  $G$  and  $H$  to the original random oracles,

<p><b>Game <math>Q_4</math> :</b>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>L, M \leftarrow \text{nil};</math>  <math>y \xleftarrow{\\$} \{0, 1\}^k;</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>\tilde{b} \leftarrow \mathcal{A}_2(y);</math>  <math>\hat{r} \xleftarrow{\\$} \{0, 1\}^\rho;</math></p>	<p><b>Oracle <math>G(r)</math> :</b>  if <math>r \notin \text{dom}(L)</math> then  <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>L \leftarrow (r, g) :: L</math>  return <math>L[r]</math></p> <p><b>Oracle <math>H(s)</math> :</b>  if <math>s \notin \text{dom}(M)</math> then  <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>M \leftarrow (s, h) :: M</math>  return <math>M[s]</math></p>
--	---

If  $\mathbf{bad}_2$  is set to true in  $Q_3$  then  $\hat{r}$  must be in  $L$ , i.e. the following is a relational invariant of  $Q_3$  and  $Q_4$ ,

$$\mathbf{bad}_2\langle 1 \rangle \implies (\hat{r} \in \text{dom}(L))\langle 2 \rangle$$

and we can readily bound the probability of  $\hat{r}$  being in the domain of  $L$  in  $Q_4$  because we know that the adversary makes at most  $q_G$  calls to  $G$ ,

$$\Pr[G_8 : \mathbf{bad}_2] = \Pr[Q_3 : \mathbf{bad}_2] \leq \Pr[Q_4 : \hat{r} \in \text{dom}(L)] \leq \frac{q_G}{2^\rho} \quad (6.15)$$

Putting (6.10), (6.14) and (6.15) together,

$$\Pr[G_7 : \text{bad}] \leq \Pr[G_{\text{OW}} : x = f^{-1}(sk, y)] + \frac{q_G}{2^\rho} \quad (6.16)$$

Finally from (6.3)–(6.9) and (6.16),

$$\left| \Pr[G_{\text{INDCPA}} : b = \tilde{b}] - \frac{1}{2} \right| \leq \Pr[G_{\text{OW}} : x = f^{-1}(sk, y)] + \frac{q_G}{2^\rho} \quad (6.17)$$

From which follows the bound (6.1).  $\square$

## 6.2 Indistinguishability under Chosen-Ciphertext Attacks

We use as a starting point the proof of Pointcheval [2005], which fills several gaps in the reduction of Fujisaki et al. [2004], resulting in a weaker security bound. Our proof unveils minor glitches in the proof of Pointcheval [2005] and marginally improves on its exact security bound (reducing the coefficients) by performing an aggressive analysis of oracle queries earlier in the sequence of games. The initial and final games of the reduction appear in Figure 6.3; the proof in CertiCrypt is about 10,000 lines long.

**Theorem 6.4 (IND-CCA2 security of OAEP).** *Let  $\mathcal{A}$  be an adversary against the ciphertext indistinguishability of OAEP under an adaptive chosen-ciphertext attack that makes at most  $q_G$  and  $q_H$  queries to the hash oracles  $G$  and  $H$ , respectively, and at most  $q_D$  queries to the decryption oracle  $\mathcal{D}$ .<sup>1</sup> Suppose  $\mathcal{A}$  achieves an advantage  $\epsilon$  within time  $t$  during game  $G_{\text{INDCCA}}$ . Then, there exists an inverter  $\mathcal{I}$  that finds a partial preimage (the  $k - \rho$  most significant bits of a preimage) of an element uniformly drawn from the domain of  $f$  with probability  $\epsilon'$  within time  $t'$  during experiment  $G_{\text{PDOW}}$ , where*

$$t' \leq t + q_G q_H q_D O(t_f) \quad \epsilon' \geq \frac{1}{q_H} \left( \epsilon - \frac{3q_D q_G + q_D^2 + 4q_D + q_G}{2^\rho} - \frac{2q_D}{2^{k_1}} \right)$$

### 6.2.1 Proof Outline

Figure 6.4 outlines the structure of the proof; the first step from  $G_{\text{INDCCA}}$  to  $G_1$  and the final step from  $G_5$  to  $G_{\text{PDOW}}$  are not displayed. The reduction successively eliminates all situations in which the plaintext extractor used by the inverter to simulate decryption may fail.

Starting from game  $G_{\text{INDCCA}}$ , we use the logic of swapping statements to fix the hash  $\hat{g}$  that  $G$  gives in response to the random seed of the challenge ciphertext; the computation of the challenge ciphertext unfolds to:

<sup>1</sup> The machine-checked proof slightly relaxes this condition; it requires that the length of  $L_G$  be at most  $q_D + q_G$ , so that the adversary could trade queries to  $\mathcal{D}$  for queries to  $G$ .




<p><b>Game</b> <math>G_{\text{INDCCA}}</math> :</p> $L_G, L_H, L_D \leftarrow \text{nil};$ $(pk, sk) \leftarrow \mathcal{KG}();$ $(m_0, m_1) \leftarrow \mathcal{A}_1(pk);$ $b \xleftarrow{\$} \{0, 1\};$ $\hat{c} \leftarrow \mathcal{E}(m_b);$ $\hat{c}_{\text{def}} \leftarrow \text{true};$ $\tilde{b} \leftarrow \mathcal{A}_2(\hat{c})$	<p><b>Oracle</b> <math>G(r)</math> :</p> if $r \notin \text{dom}(L_G)$ then $g \xleftarrow{\$} \{0, 1\}^{k-\rho};$ $L_G \leftarrow (r, g) :: L_G$ return $L_G[r]$ <p><b>Oracle</b> <math>H(s)</math> :</p> if $s \notin \text{dom}(L_H)$ then $h \xleftarrow{\$} \{0, 1\}^\rho;$ $L_H \leftarrow (s, h) :: L_H$ return $L_H[s]$	<p><b>Oracle</b> <math>\mathcal{D}(c)</math> :</p> $L_D \leftarrow (\hat{c}_{\text{def}}, c) :: L_D;$ $(s, t) \leftarrow f^{-1}(sk, c);$ $h \leftarrow H(s);$ $r \leftarrow t \oplus h;$ $g \leftarrow G(r);$ if $[s \oplus g]_{k_1} = 0^{k_1}$ then return $[s \oplus g]^\mu$ else return $\perp$
<p><b>Game</b> <math>G_{\text{PDOW}}</math> :</p> $(pk, sk) \leftarrow \mathcal{KG}_f();$ $s \xleftarrow{\$} \{0, 1\}^{k-\rho};$ $t \xleftarrow{\$} \{0, 1\}^\rho;$ $\tilde{s} \leftarrow \mathcal{I}(pk, f(pk, s \parallel t))$ <p><b>Adversary</b> <math>\mathcal{I}(pk, y)</math> :</p> $L_G, L_H \leftarrow \text{nil};$ $\hat{pk} \leftarrow pk;$ $(m_0, m_1) \leftarrow \mathcal{A}_1(pk);$ $\tilde{b} \leftarrow \mathcal{A}_2(y);$ $s \xleftarrow{\$} \text{dom}(L_H);$ return $s$	<p><b>Oracle</b> <math>G(r)</math> :</p> if $r \notin \text{dom}(L_G)$ then $g \xleftarrow{\$} \{0, 1\}^{k-\rho};$ $L_G \leftarrow (r, g) :: L_G$ return $L_G[r]$ <p><b>Oracle</b> <math>H(s)</math> :</p> if $s \notin \text{dom}(L_H)$ then $h \xleftarrow{\$} \{0, 1\}^\rho;$ $L_H \leftarrow (s, h) :: L_H$ return $L_H[s]$	<p><b>Oracle</b> <math>\mathcal{D}(c)</math> :</p> if $\exists(s, h) \in L_H, (r, g) \in L_G.$ $c = f(\hat{pk}, s \parallel (r \oplus h)) \wedge$ $[s \oplus g]_{k_1} = 0^{k_1}$ then return $[s \oplus g]^\mu$ else return $\perp$

**Fig. 6.3.** Initial and final games in the reduction of the IND-CCA2 security of OAEP to the problem of partially inverting the underlying permutation. We exclude cheating adversaries who query the decryption oracle with the challenge ciphertext during the second phase of the experiment by requiring  $(\text{true}, \hat{c}) \notin L_D$  to be a post-condition of the initial game.


$$\begin{aligned}
\hat{r} &\xleftarrow{\$} \{0, 1\}^\rho; \\
\hat{s} &\leftarrow \hat{g} \oplus (m_b \parallel 0^{k_1}); \\
\hat{h} &\leftarrow H(\hat{s}); \\
\hat{t} &\leftarrow \hat{h} \oplus \hat{r}; \\
\hat{c} &\leftarrow f(\hat{pk}, \hat{s} \parallel \hat{t})
\end{aligned}$$

where  $\hat{g}$  is sampled from  $\{0, 1\}^{k-\rho}$  before the first call to  $\mathcal{A}$ . We then make  $G$  respond to an adversary query  $\hat{r}$  with a freshly sampled value instead of  $\hat{g}$ ; this only makes a difference if flag **bad** is set in game  $G_1$ . Since at this point  $\hat{g}$  is uniformly distributed and independent from the adversary's view, the value  $\hat{s}$  computed as  $\hat{g} \oplus (m_b \parallel 0^{k_1})$  is as well uniformly distributed and independent from the adversary's view. This removes the dependence of the adversary output on the hidden bit  $b$ , and thus the probability of a correct guess is exactly  $1/2$ . Using the Fundamental Lemma we obtain the bound:

$$\begin{aligned}
\Pr [G_{\text{IND-CCA2}} : \tilde{b} = b] - \Pr [G_1 : \tilde{b} = b] &= \Pr [G_{\text{IND-CCA2}} : \tilde{b} = b] - \frac{1}{2} \\
&\leq \Pr [G_1 : \text{bad}] \tag{6.18}
\end{aligned}$$



**Fig. 6.4.** Outline of the reduction showing the *lossy* transitions. Fragments of code that change between games are highlighted on a gray background.



**Fig. 6.4.** Outline of the reduction showing the *lossy* transitions. Fragments of code that change between games are highlighted on a gray background.

The transition from  $G_1$  to  $G_2$  modifies the decryption oracle successively by inlining the call to  $G$ , and by applying the Fundamental and Failure Event lemmas to reject the ciphertext when there is a negligible chance it matches the padding. Overall, we prove:

$$\Pr[G_1 : \text{bad}] \leq \Pr[G_2 : \text{bad}] + \frac{q_D^2 + q_D q_G + q_D}{2^\rho} + \frac{q_D}{2^{k_1}} \quad (6.19)$$

Next, we eliminate fresh calls to  $G$  in the decryption oracle. These calls correspond to the two assignments  $L_G[r] \leftarrow g$ , since calls to  $G$  have been inlined previously. We perform an aggressive elimination and remove both calls. As a result, in game  $G_3$  the length of list  $L_G$  (i.e. the number of calls to  $G$ ) is bounded by  $q_G$  rather than  $q_D + q_G$ . This is the key to improve on the security bound of Pointcheval

[2005], who only removes the second call. The proof relies on the logic of swapping statements to show that values of discarded calls are “uniformly distributed and independent from the adversary’s view”. Details appear below. Overall, we prove:

$$\Pr[\mathbf{G}_2 : \mathbf{bad}] \leq \Pr[\mathbf{G}_3 : \mathbf{bad}] + \frac{q_{\mathcal{D}}}{2^{k_1}} \quad (6.20)$$

Likewise, we eliminate calls to  $H$  in  $\mathcal{D}$ , yielding a new game  $\mathbf{G}_4$  in which the decryption oracle does not add any new values to the memories of  $G$  and  $H$ . Using the Fundamental and Failure Event lemmas, we obtain:

$$\Pr[\mathbf{G}_3 : \mathbf{bad}] \leq \Pr[\mathbf{G}_4 : \mathbf{bad}] + \frac{q_{\mathcal{D}}q_G + q_{\mathcal{D}}}{2^\rho} \quad (6.21)$$

We next fix the value  $\hat{\mathbf{h}}$  that oracle  $H$  gives in response to  $\hat{\mathbf{s}}$ , and then make  $H$  return a freshly sampled value instead of  $\hat{\mathbf{h}}$ . This allows us to bound the probability of  $\mathbf{bad}$  in terms of the probability of a newly introduced event  $\mathbf{bad}_H$ . The proof uses the hypothesis that  $\mathcal{A}_2$  cannot query the decryption oracle with the challenge ciphertext, and yields:

$$\Pr[\mathbf{G}_4 : \mathbf{bad}] \leq \Pr[\mathbf{G}_5 : \mathbf{bad}_H] + \frac{q_{\mathcal{D}}q_G + 2q_{\mathcal{D}} + q_G}{2^\rho} \quad (6.22)$$

Finally, we prove that the probability of  $\mathbf{bad}_H$  in  $\mathbf{G}_5$  is upper bounded by the probability that the inverter  $\mathcal{I}$  succeeds in partially inverting the permutation  $f$ . The proof uses the (standard, non-relational) invariant on  $\mathbf{G}_5$ :

$$\mathbf{bad}_H \implies \hat{\mathbf{s}} \in \text{dom}(\mathbf{L}_H)$$

The inverter  $\mathcal{I}$  that we build (shown in Figure 6.3) gives its own challenge  $y$  as the challenge ciphertext to the IND-CCA2 adversary  $\mathcal{A}$  and returns a random element in the list of queries made to  $H$ . Thus,

$$\Pr[\mathbf{G}_5 : \mathbf{bad}_H] \leq \Pr[\mathbf{G}_5 : \hat{\mathbf{s}} \in \text{dom}(\mathbf{L}_H)] \leq \frac{1}{q_H} \Pr[\mathbf{G}_{\text{PDOW}} : s = \tilde{s}] \quad (6.23)$$

Where the last inequality follows from the bound on the number of queries to oracle  $H$  and an instance of the optimistic sampling equivalence:

$$\vdash \hat{\mathbf{h}} \stackrel{\$}{\leftarrow} \{0, 1\}^\rho; \hat{\mathbf{t}} \leftarrow \hat{\mathbf{h}} \oplus \hat{\mathbf{r}} \stackrel{\{\hat{\mathbf{r}}\}}{\simeq_{\{\hat{\mathbf{h}}, \hat{\mathbf{t}}, \hat{\mathbf{r}}\}}} \hat{\mathbf{t}} \stackrel{\$}{\leftarrow} \{0, 1\}^\rho; \hat{\mathbf{h}} \leftarrow \hat{\mathbf{t}} \oplus \hat{\mathbf{r}}$$

Putting together (6.18)–(6.23) concludes the proof of the statement in Theorem 6.4.

### Detailed proof of the transition from $\mathbf{G}_2$ to $\mathbf{G}_3$

We use the five intermediate games shown in Figure 6.5. The first transition from  $\mathbf{G}_2$  to  $\mathbf{G}_2^1$  consists in adding a tag to queries in the memory of  $G$  indicating whether the query has been made directly by the adversary or indirectly, through the decryption oracle. The decryption oracle tests this tag when accessing the memory of  $G$ : if the

<p><b>Game <math>\overline{G_2^1} \overline{G_2^2}</math> :</b>  <math>L_G, L_H, L_D \leftarrow \text{nil};</math>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>b \xleftarrow{\\$} \{0, 1\};</math>  <math>\hat{r} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{s} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>\hat{h} \leftarrow H(\hat{s});</math>  <math>\hat{t} \leftarrow \hat{h} \oplus \hat{r};</math>  <math>\hat{c} \leftarrow f(pk, \hat{s} \parallel \hat{t});</math>  <math>\hat{c}_{\text{def}} \leftarrow \text{true};</math>  <math>\hat{b} \leftarrow \mathcal{A}_2(\hat{c})</math></p>	<p><b>Oracle <math>G(r)</math> :</b>  if <math>r \notin \text{dom}(L_G)</math> then  if <math>r = \hat{r}</math> then    <b>bad</b> <math>\leftarrow</math> true    <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>    <math>L_G[r] \leftarrow (\text{false}, g)</math>  else    <math>(d, g) \leftarrow L_G[r];</math>    <math>L_G[r] \leftarrow (\text{false}, g)</math>    return <math>g</math></p> <p><b>Oracle <math>H(s)</math> :</b>  if <math>s \notin \text{dom}(L_H)</math> then    <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>    <math>L_H[s] \leftarrow h</math>  else <math>h \leftarrow L_H[s]</math>  return <math>h</math></p>	<p><b>Oracle <math>\mathcal{D}(c)</math> :</b>  if <math>(\hat{c}_{\text{def}} \wedge \hat{c} = c) \vee q_D &lt;  L_D  \vee q_D + q_G &lt;  L_G </math> then  return <math>\perp</math>  else    <math>L_D \leftarrow (\hat{c}_{\text{def}}, c) :: L_D; (s, t) \leftarrow f^{-1}(sk, c);</math>    if <math>s \in \text{dom}(L_H)</math> then      <math>r \leftarrow t \oplus H(s);</math>      if <math>r \in \text{dom}(L_G)</math> then        <math>(d, g) \leftarrow L_G[r];</math>        if <math>d = \text{true}</math> then          if <math>[s \oplus g]_{k_1} = 0^{k_1}</math> then            <b>bad</b><math>_1 \leftarrow</math> true; <math>\boxed{\text{return } [s \oplus g]^\mu}</math> <math>\boxed{\text{return } \perp}</math>          else return <math>\perp</math>        else          if <math>[s \oplus g]_{k_1} = 0^{k_1}</math> then return <math>[s \oplus g]^\mu</math>          else return <math>\perp</math>        else          if <math>r = \hat{r}</math> then <b>bad</b> <math>\leftarrow</math> true;          <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho}; L_G[r] \leftarrow (\text{true}, g);</math> return <math>\perp</math>        else          <math>r \leftarrow t \oplus H(s);</math>          if <math>r \notin \text{dom}(L_G)</math> then            <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho}; L_G[r] \leftarrow (\text{true}, g);</math>          return <math>\perp</math></p>
<p><b>Game <math>\overline{G_2^3} \overline{G_2^4} \overline{G_2^5}</math> :</b>  <math>L_G, L_H, L_D \leftarrow \text{nil};</math>  <math>(pk, sk) \leftarrow \mathcal{KG}_f();</math>  <math>(m_0, m_1) \leftarrow \mathcal{A}_1(pk);</math>  <math>b \xleftarrow{\\$} \{0, 1\};</math>  <math>\hat{r} \xleftarrow{\\$} \{0, 1\}^\rho;</math>  <math>\hat{s} \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>  <math>\hat{h} \leftarrow H(\hat{s});</math>  <math>\hat{t} \leftarrow \hat{h} \oplus \hat{r};</math>  <math>\hat{c} \leftarrow f(pk, \hat{s} \parallel \hat{t});</math>  <math>\hat{c}_{\text{def}} \leftarrow \text{true};</math>  <math>\hat{b} \leftarrow \mathcal{A}_2(\hat{c})</math></p> <div style="border: 1px dashed black; padding: 5px;"> <p><math>L \leftarrow L_G;</math>  while <math>L \neq \text{nil}</math> do    <math>(r, (b, g)) \leftarrow \text{hd}(L);</math>    if <math>b = \text{true}</math> then      <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>      <math>L_G[r] \leftarrow (\text{true}, g)</math>    <math>L \leftarrow \text{tl}(L)</math></p> </div>	<p><b>Oracle <math>G(r)</math> :</b>  if <math>r \notin \text{dom}(L_G)</math> then  if <math>r = \hat{r}</math> then    <b>bad</b> <math>\leftarrow</math> true    <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho};</math>    <math>L_G[r] \leftarrow (\text{false}, g)</math>  else    <math>(d, g) \leftarrow L_G[r];</math>    if <math>d = \text{true}</math> then      <math>\boxed{g \xleftarrow{\\$} \{0, 1\}^{k-\rho};}</math>      <math>\boxed{g \xleftarrow{\\$} \{0, 1\}^{k-\rho};}</math>      <math>L_G[r] \leftarrow (\text{false}, g);</math>      <b>bad</b><math>_2 \leftarrow P(g, r)</math>    return <math>g</math></p> <p><b>Oracle <math>H(s)</math> :</b>  if <math>s \notin \text{dom}(L_H)</math> then    <math>h \xleftarrow{\\$} \{0, 1\}^\rho;</math>    <math>L_H[s] \leftarrow h</math>  else <math>h \leftarrow L_H[s]</math>  return <math>h</math></p>	<p><b>Oracle <math>\mathcal{D}(c)</math> :</b>  if <math>(\hat{c}_{\text{def}} \wedge \hat{c} = c) \vee q_D &lt;  L_D  \vee q_D + q_G &lt;  L_G </math> then  return <math>\perp</math>  else    <math>L_D \leftarrow (\hat{c}_{\text{def}}, c) :: L_D; (s, t) \leftarrow f^{-1}(sk, c);</math>    if <math>s \in \text{dom}(L_H)</math> then      <math>r \leftarrow t \oplus H(s);</math>      if <math>r \in \text{dom}(L_G)</math> then        <math>(d, g) \leftarrow L_G[r];</math>        if <math>d = \text{true}</math> then return <math>\perp</math>      else        if <math>[s \oplus g]_{k_1} = 0^{k_1}</math> then return <math>[s \oplus g]^\mu</math>        else return <math>\perp</math>      else        if <math>r = \hat{r}</math> then <b>bad</b> <math>\leftarrow</math> true;        <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho}; L_G[r] \leftarrow (\text{true}, g);</math> return <math>\perp</math>      else        <math>r \leftarrow t \oplus H(s);</math>        if <math>r \notin \text{dom}(L_G)</math> then          <math>g \xleftarrow{\\$} \{0, 1\}^{k-\rho}; L_G[r] \leftarrow (\text{true}, g);</math>        return <math>\perp</math></p>

$P(g, r) \stackrel{\text{def}}{=} \exists (d, c) \in L_D. \text{ let } (s, t) = f^{-1}(sk, c) \text{ in } s \in \text{dom}(L_H) \wedge r = t \oplus L_H[s] \wedge [s \oplus g]_{k_1} = 0^{k_1}$

**Fig. 6.5.** Games in the transition from  $G_2$  to  $G_3$ . Fragments of code inside a box appear only in the game whose name is surrounded by the matching box.

ciphertext queried is valid and its random seed appeared in a previous decryption query, but not yet in a direct query to  $G$ , the decryption oracle raises a flag  $\mathbf{bad}_1$ . We show that this can happen with probability  $2^{-k_1}$  for any single query, since the random seed is uniformly distributed and independent from the adversary's view. In this case, the decryption oracle can safely reject the ciphertext, as done in game  $G_2^2$ . The proof proceeds in two steps: We show that game  $G_2$  is observationally equivalent to game  $G_2^1$  using the relational invariant

$$\mathbf{L}_G\langle 1 \rangle = (\text{map } (\lambda(r, (b, g)).(r, g)) \mathbf{L}_G)\langle 2 \rangle$$

Therefore  $\Pr[G_2 : \mathbf{bad}] = \Pr[G_2^1 : \mathbf{bad}]$ . Game  $G_2^2$  is identical to  $G_2^1$ , except that it rejects ciphertexts that raise the  $\mathbf{bad}_1$  flag. Applying the Fundamental Lemma (i.e. Lemma 3.3), we show that

$$\Pr[G_2^1 : \mathbf{bad}] \leq \Pr[G_2^2 : \mathbf{bad}] + \Pr[G_2^2 : \mathbf{bad}_1]$$

Our next goal is to show that answers to queries tagged as true can be resampled. However, one cannot apply the logic of swapping statements at this stage to resample these answers in  $G$ , because flag  $\mathbf{bad}_1$  is set on  $\mathcal{D}$  and depends on them. The solution is to introduce a new game  $G_2^3$  that sets another flag  $\mathbf{bad}_2$  in the code of  $G$  instead of setting  $\mathbf{bad}_1$  in the decryption oracle<sup>2</sup>. Flag  $\mathbf{bad}_2$  is raised whenever the adversary queries  $G$  with the random seed of a valid ciphertext previously submitted to the decryption oracle. We prove that games  $G_2^2$  and  $G_2^3$  satisfy the relational invariant:

$$\mathbf{bad}_1\langle 1 \rangle \implies (\mathbf{bad}_2 \vee \phi)\langle 2 \rangle$$

where

$$\begin{aligned} \phi \stackrel{\text{def}}{=} & \exists(d, c) \in \mathbf{L}_{\mathcal{D}}. \text{let } (s, t) = f^{-1}(\mathbf{sk}, c), r = t \oplus \mathbf{L}_H[s] \text{ in} \\ & r \in \text{dom}(\mathbf{L}_G) \wedge s \in \text{dom}(\mathbf{L}_H) \wedge \text{fst}(\mathbf{L}_G[r]) = \text{false} \wedge [s \oplus \text{snd}(\mathbf{L}_G[r])]_{k_1} = 0^{k_1} \end{aligned}$$

Therefore:

$$\Pr[G_2^2 : \mathbf{bad}] + \Pr[G_2^2 : \mathbf{bad}_1] \leq \Pr[G_2^3 : \mathbf{bad}] + \Pr[G_2^3 : \mathbf{bad}_2 \vee \phi]$$

We now consider game  $G_2^4$  where oracle  $G$  resamples the answers to queries previously sampled in the decryption oracle. As such answers are uniformly distributed and independent from the adversary's view, the logic for swapping statements can be used to establish that this transformation preserves semantics. Hence:

$$\Pr[G_2^3 : \mathbf{bad}] + \Pr[G_2^3 : \mathbf{bad}_2 \vee \phi] = \Pr[G_2^4 : \mathbf{bad}] + \Pr[G_2^4 : \mathbf{bad}_2 \vee \phi]$$

Note that in order to prove semantic equivalence we need to resample the values in  $\mathbf{L}_G$  associated to queries tagged as true—made by the  $\mathcal{D}$ —at the end of the game. Using Lemma 3.4, we upper bound the probability of  $\mathbf{bad}_2 \vee \phi$  in  $G_2^4$ :

<sup>2</sup> As  $\mathbf{bad}_1$  is not set anymore, we simplify the code of  $\mathcal{D}$  by coalescing branches in the innermost conditional.

$$\Pr [\mathbf{G}_2^4 : \mathbf{bad}_2 \vee \phi] \leq \frac{q_{\mathcal{D}}}{2^{k_1}}$$

We are now only interested in bounding  $\mathbf{bad}$ , so we can remove as dead code the fragment of code at the end of  $\mathbf{G}_2^4$  that resamples values in  $\mathbf{L}_G$ , obtaining  $\mathbf{G}_2^5$ , and prove that

$$\Pr [\mathbf{G}_2^4 : \mathbf{bad}] = \Pr [\mathbf{G}_2^5 : \mathbf{bad}]$$

We finally prove that game  $\mathbf{G}_2^5$  is observationally equivalent to  $\mathbf{G}_3$ , in which the code for the oracle  $G$  is reverted to its original form and the decryption oracle no longer tampers with the memory of  $G$ . Thus,

$$\Pr [\mathbf{G}_2 : \mathbf{bad}] \leq \Pr [\mathbf{G}_2^5 : \mathbf{bad}] + \frac{q_{\mathcal{D}}}{2^{k_1}} = \Pr [\mathbf{G}_3 : \mathbf{bad}] + \frac{q_{\mathcal{D}}}{2^{k_1}}$$

□

*Comparison with the security bound of Pointcheval [2005]*

Pointcheval obtains a slightly different bound:

$$\epsilon' \geq \frac{1}{q_{\mathbf{H}}} \left( \epsilon - \frac{4q_{\mathcal{D}}q_{\mathbf{G}} + 2q_{\mathcal{D}}^2 + 4q_{\mathcal{D}} + 8q_{\mathbf{G}}}{2^{\rho}} - \frac{3q_{\mathcal{D}}}{2^{k_1}} \right)$$

We marginally improve on this bound by reducing the coefficients. As previously mentioned, the improvement stems from the transition from  $\mathbf{G}_2$  to  $\mathbf{G}_3$ , where we eliminate both calls to  $G$  in the decryption oracle, whereas only one of them is eliminated in [Pointcheval 2005]. In fact, eliminating both calls is not only needed to give a better bound, but is essential for the correctness of the proof. Indeed, the transition from  $\mathbf{G}_3$  to  $\mathbf{G}_4$  would not be possible if  $\mathcal{D}$  modified the memory of  $G$ . Concretely, the justification of Equation (27) in [Pointcheval 2005] contains two minor glitches: firstly, the remark “which just cancels  $r'$  from  $\mathbf{L}_G$ ” oversees the possibility of this removal having an impact on future queries. Secondly, “the probability for  $r'$  to be in  $\mathbf{L}_G$  is less than  $q_{\mathbf{G}}/2^{\rho}$ ” oversees that the length of  $\mathbf{L}_G$  is upper bounded by  $q_{\mathbf{G}} + q_{\mathcal{D}}$  rather than just  $q_{\mathbf{G}}$ , as the decryption oracle still adds values to  $\mathbf{L}_G$ ; a correct bound for this probability in [Pointcheval 2005] is  $(q_{\mathbf{G}} + q_{\mathcal{D}})/2^{\rho}$ .

## 6.2.2 Notes about the Proved Security Bound

We note that although we exhibit an explicit inverter that achieves the advantage in the statement of Theorem 6.4, we do not prove formally that it executes within the given time bound. We would be loath to say that the proof is incomplete, because the time complexity of the inverter is evident from its formulation. Nonetheless, we can prove that the inverter executes in probabilistic polynomial-time (under a reasonable cost model for constructions in the language) and thus the asymptotic security of OAEP under the hypothesis that the underlying permutation family is partial-domain one-way. Formally, if the winning probability  $\Pr [\mathbf{G}_{\text{PDOW}} : s = \tilde{s}]$  of

any PPT inverter is negligible, we can prove that the advantage of any IND-CCA2 adversary is also negligible, provided  $k_1$  and  $\eta$  increase at least linearly with the security parameter, i.e.,

$$\exists \alpha. \exists \eta_0. \forall \eta \geq \eta_0. \rho \geq \alpha \eta \wedge k_1 \geq \alpha \eta$$

Moreover, we observe that by using a standard time-space trade-off argument (cf. [Pointcheval 2005]) one can reduce the factor  $q_G q_H q_D$  in the time bound of the inverter to  $q_G q_H$ .

### 6.3 About the Security of OAEP in the Standard Model

From a practical point of view, in reductionist arguments like the ones we studied in this chapter, the tightness of the security bound makes a whole world of difference. The tighter the bound is, the closer the problem of breaking the security of the scheme is to the problem of solving the computational hard problem under consideration. Another aspect of practical importance is the model where the proof is carried out. A proof in the random oracle model, like the proofs we presented, only rules out *generic attacks* that do not exploit the implementation details of the hash functions  $G$  and  $H$ . In contrast, a proof in the standard model of cryptography considers the possibility that an adversary might be able to exploit weaknesses in the actual implementation of the hash functions to attack the security of the scheme.

Several authors have studied the possibility of proving OAEP IND-CCA2 secure in the standard model, either when used in conjunction with the RSA function or with a trapdoor permutation satisfying particular classes of properties. Shoup [2001] showed the impossibility of finding a blackbox reduction from the IND-CCA2 security of OAEP to the one-wayness of the underlying trapdoor permutation, either in the standard or random oracle model, but he exhibited a blackbox reduction in the random oracle model to the problem of partially inverting the permutation. Brown [2006] showed a pathological instantiation of the hash functions that would render OAEP insecure, which means that a security proof should at least assume some property about the hash functions. He ruled out as well the possibility of proving the IND-CCA2 security of OAEP in the standard model using certain types of reductions. Finally, Kiltz and Pietrzak [2009] prove a blackbox separation result: no padding-based encryption scheme can be proved IND-CCA2 secure in the standard model, even assuming the underlying trapdoor permutation is *ideal*. This means in particular, that is impossible to prove the IND-CCA2 security of OAEP in the standard model under most standard assumptions about the trapdoor permutation, including one-wayness, partial-domain one-wayness, or claw-freeness.

In contrast, Kiltz et al. [2010] recently proved the IND-CPA security of generic OAEP in the standard model under non-interactive and non-interdependent assumptions on the underlying trapdoor permutation and the hash functions. In particular, they prove that RSA-OAEP is IND-CPA secure when instantiated using a  $t$ -wise independent hash function (for an appropriate  $t$ ) provided RSA is a



*lossy* trapdoor permutation. This latter condition can be proved under the phi-hiding assumption, a number-theoretic assumption slightly stronger than the RSA assumption.

Despite the above negative results, most people agree that a security proof of a scheme in the random oracle model gives strong evidence about its practical security, and that it is very unlikely that someone comes out with an attack exploiting details of the concrete instantiations of random oracles.

# 7

## Machine-Checked Formalization of Zero-Knowledge Protocols

PROOFS of knowledge [Goldreich 2002; Goldwasser et al. 1989] are two-party interactive protocols where one party, called the *prover*, convinces the other one, called the *verifier*, that it knows *something*. Typically, both parties share a common input  $x$  and *something* refers to a witness  $w$  of membership of the input  $x$  to an  $\mathcal{NP}$  language. Proofs of knowledge are useful to enforce honest behavior of potentially malicious parties [Backes et al. 2009]: the knowledge witness acts as an authentication token used to establish that the prover is a legitimate user of a service provided by the verifier, or as evidence that a message sent by the prover has been generated in accordance to the rules of a protocol. Proofs of knowledge must be complete, so that a prover that has indeed knowledge of a witness can convince a honest verifier, and sound, so that a dishonest prover has little chance of being convincing. In addition, practical applications often require to preserve secrecy or anonymity; in such cases the proof should not leak anything about the witness. Zero-knowledge proofs are computationally convincing proofs of knowledge that achieve this goal, i.e. they are convincing and yet the verifier does not learn anything from interacting with the prover beyond the fact that the prover knows a witness for their common input. This property has an elegant formulation: a protocol is said to be zero-knowledge when transcripts of protocol runs between a prover  $P$  and a (possibly dishonest) verifier  $V$  can be efficiently simulated without ever interacting with the prover—but with access to the strategy of  $V$ . In particular, this implies that proofs are not transferable; a conversation is convincing only for the verifier interacting with the prover and cannot be replayed to convince a third party.

In his PhD dissertation, Cramer [1996] introduced the concept of  $\Sigma$ -protocols, a class of three-move interactive protocols that are suitable as a basis for the design of many efficient and secure cryptographic services. Cramer described  $\Sigma$ -protocols as abstract modules and showed that they are realizable when instantiated for most

computational assumptions commonly considered in cryptography, including the difficulty of computing discrete logarithms or factoring integers, or the existence of some abstract function families (e.g. one-way group homomorphisms). In addition, he gave an effective method to combine  $\Sigma$ -protocols to obtain zero-knowledge proofs of any Boolean formula constructed using the AND and OR operators from formulae for which  $\Sigma$ -protocols exist. This means that  $\Sigma$ -protocols can be used in a practical setting as building blocks to achieve various cryptographic goals. Applications of  $\Sigma$ -protocols notably include secure multi-party computation, identification schemes, secret-ballot electronic voting, and anonymous attestation credentials.

This chapter reports on a fully machine-checked formalization of a comprehensive theory of  $\Sigma$ -protocols using `CertiCrypt`. The formalization consists of more than 10,000 lines of `Coq` code, and covers the basics of  $\Sigma$ -protocols: definitions, relations between different notions of security, general constructions and composability theorems. We show its applicability by formalizing several well-known protocols, including the Schnorr, Guillou-Quisquater, Okamoto, and Feige-Fiat-Shamir protocols. The highlight of the formalization is a generic account of  $\Sigma^\phi$ -protocols, that prove knowledge of a preimage under a group homomorphism  $\phi$ . We use the module system of `Coq` to define and relate the classes of  $\Sigma^\phi$ - and  $\Sigma$ -protocols. Our formalization of  $\Sigma^\phi$ -protocols provides sufficient conditions (the so-called specialness conditions) on the group homomorphism  $\phi$  so that every  $\Sigma^\phi$ -protocol can be construed as a  $\Sigma$ -protocol. Moreover, we show that special homomorphisms are closed under direct product, which yields a cheap way of AND-combining  $\Sigma^\phi$  proofs. We exploit the generality of our result to achieve short proofs of completeness, special soundness, and (honest verifier) zero-knowledge for many protocols in the literature.


## 7.1 Sigma-Protocols

A  $\Sigma$ -protocol is a 3-step interactive protocol where a prover  $P$  interacts with a verifier  $V$ . Both parties have access to a common input  $x$ , and the goal of the prover is to convince the verifier that it knows some value  $w$  suitably related to  $x$ , without revealing anything beyond this assertion. The protocol begins with the prover sending a commitment  $r$  to the verifier, who responds by sending back a random challenge  $c$ ; the prover then computes a response  $s$  to the challenge and sends it to the verifier, who either accepts or rejects the conversation. Figure 7.1 shows a diagram of a run of a  $\Sigma$ -protocol.

Formally, a  $\Sigma$ -protocol is defined with respect to a knowledge relation  $R$ . This terminology comes from interpreting the proof system as proving membership of the common input to an  $\mathcal{NP}$  language  $L$ . Each  $\mathcal{NP}$  language admits an efficient membership verification procedure via a polynomial-time recognizable relation  $R_L$  such that

$$L = \{x \mid \exists w. R_L(x, w)\}$$

Proving that  $x$  belongs to the language amounts to proving knowledge of a witness  $w$  related to  $x$  via  $R_L$ . In `CertiCrypt`, the class of  $\Sigma$ -protocols is formalized as a



**Fig. 7.1.** Characteristic 3-step interaction in a run of a  $\Sigma$ -protocol.

module type parametrized over a knowledge relation  $R$ , and a number of procedures specifying the different phases of the prover and the verifier; the module type specifies as well the properties that any given protocol instance must satisfy. In the remainder of this section we describe in detail our formal definition of  $\Sigma$ -protocols and comment on an alternative—but in some sense equivalent—specification of the zero-knowledge property.

**Definition 7.1 ( $\Sigma$ -protocol).** *A  $\Sigma$ -protocol for a knowledge relation  $R$  is a 3-step protocol between a prover  $P$  and a verifier  $V$ , whose interaction is described by the following parametrized program:*

**Protocol** $(x, w)$  :

$(r, state) \leftarrow P_1(x, w);$   
 $c \leftarrow V_1(x, r);$   
 $s \leftarrow P_2(x, w, state, c);$   
 $b \leftarrow V_2(x, r, c, s)$

In the above program specifying a  $\Sigma$ -protocol, the two phases of the prover are described by the procedures  $P_1$  and  $P_2$ , while the phases of the verifier are described by  $V_1$  and  $V_2$ . Note that the protocol explicitly passes state between the phases of the participants; we could have used instead global variables shared between  $P_1$  and  $P_2$  on one hand, and  $V_1$  and  $V_2$  on the other, but that would unnecessarily complicate the proofs because we would need to specify that the procedures representing one party do not have access to the shared state of the other party. All the protocols that we consider in the following are public-coin, meaning that a honest verifier chooses the challenge uniformly from some predefined set  $C$ . A  $\Sigma$ -protocol must satisfy the following three properties,

1. *Completeness:* Given a public input  $x$  together with a witness  $w$  such that  $R(x, w)$ , the prover is always able to convince the verifier:

$$\forall m. R(m(x), m(w)) \implies \Pr[\text{Protocol}, m : b = \text{true}] = 1$$

2. *Special Honest Verifier Zero-Knowledge (sHVZK):* There exists a probabilistic polynomial-time simulator  $S$  that given  $x \in \text{dom}(R)$  and a challenge  $c$ , computes triples  $(r, c, s)$  with the same distribution as a valid conversation. The property is formalized in terms of a version of the protocol where the challenge  $c$  is fixed as a parameter,

$$\boxed{\begin{array}{l} \mathbf{Protocol}(x, w, c) : \\ (r, \text{state}) \leftarrow \mathbf{P}_1(x, w); \\ s \leftarrow \mathbf{P}_2(x, w, \text{state}, c); \\ b \leftarrow \mathbf{V}_2(x, r, c, s) \end{array}} \underset{\{r, c, s\}}{\overset{\{x, c\} \wedge R(x, w)}} \approx (r, s) \leftarrow \mathbf{S}(x, c)$$

3. *Special Soundness*: Given two accepting conversations  $(r, c_1, s_1)$ ,  $(r, c_2, s_2)$  for an input  $x$ , with the same commitment  $r$ , but with different challenges, there exists a PPT knowledge extractor  $\mathbf{KE}$  that computes a witness  $w$  such that  $R(x, w)$ . Formally, for any memory  $m$ ,

$$\left. \begin{array}{l} m(c_1) \neq m(c_2) \\ \Pr[b \leftarrow \mathbf{V}_2(x, r, c_1, s_1), m : b = \text{true}] = 1 \\ \Pr[b \leftarrow \mathbf{V}_2(x, r, c_2, s_2), m : b = \text{true}] = 1 \end{array} \right\} \implies \Pr[w \leftarrow \mathbf{KE}(x, r, c_1, c_2, s_1, s_2), m : R(x, w)] = 1$$

Special soundness might seem a relatively weak property at first sight. It can be shown using a rewinding argument (although we did not formalize this result in Coq) that thanks to special soundness, any public-coin  $\Sigma$ -protocol with challenge set  $C$  can be seen as a *proof of knowledge* with *soundness error*  $|C|^{-1}$  [Damgård and Pfitzmann 1998]. Informally, this means that any efficient prover (possibly dishonest) that manages to convince a honest verifier for a public input  $x$  with a probability greater than  $|C|^{-1}$  can be turned into an efficient procedure that computes a witness for  $x$ .

### 7.1.1 Relation between sHVZK and HVZK

Some authors require that  $\Sigma$ -protocols satisfy a somewhat weaker property known as honest verifier zero-knowledge rather than the *special* version of this property mentioned above. The difference is that in the former the simulator is allowed to choose the challenge while in the latter the challenge is fixed. In other words, plain HVZK requires that there exists a PPT simulator  $\mathbf{S}$  that given just  $x \in \text{dom}(R)$  computes a triple  $(r, c, s)$  with the same distribution as the verifier's view of a conversation. The relation between the two notions has been studied by Cramer [1996]. As an illustration of the use of CertiCrypt and the  $\Sigma$ -protocol framework, the formalization of this relation is discussed below. Without loss of generality we assume that the challenge set of the protocols we consider is  $\{0, 1\}^k$ .

**Theorem 7.2 (sHVZK implies HVZK).** *If a  $\Sigma$ -protocol satisfies sHVZK, it also satisfies HVZK.*

*Proof.* A HVZK simulator  $\mathbf{S}'$  can be built from the sHVZK simulator  $\mathbf{S}$ :

$$\mathbf{Simulator } \mathbf{S}'(x) : c \leftarrow \{0, 1\}^k; (r, s) \leftarrow \mathbf{S}(x, c); \text{ return } (r, c, s)$$

The fact that  $\mathbf{S}'$  perfectly simulates conversations of the protocol can be proved by means of the following sequence of games:

$$\begin{aligned}
\text{Protocol}(x, w) &\simeq_{\{r, c, s\}}^{\{x\} \wedge R(x, w)} c \stackrel{\$}{\leftarrow} \{0, 1\}^k; \text{Protocol}(x, w, c) \\
&\simeq_{\{r, c, s\}}^{\{x\} \wedge R(x, w)} c \stackrel{\$}{\leftarrow} \{0, 1\}^k; (r, s) \leftarrow S(x, c) \\
&\simeq_{\{r, c, s\}}^{\{x\} \wedge R(x, w)} (r, c, s) \leftarrow S'(x)
\end{aligned}$$

The first and last equivalences are easily proved by inlining procedure calls using the tactic `inline`, and reordering instructions in the resulting programs using `swap`. To prove the second equivalence, the tactic `eqobs_hd` is used to get rid of the instruction  $c \stackrel{\$}{\leftarrow} \{0, 1\}^k$  that is common to both games; the resulting goal matches exactly the definition of `sHVZK` for `S`.  $\square$

In a sense, `sHVZK` is a stronger property than `HVZK`, because a protocol satisfying `sHVZK` can be shown to satisfy `HVZK`, while the converse is not generally true. However, from every protocol  $(P, V)$  that satisfies `HVZK` it is possible to construct a protocol  $(P', V')$  that satisfies `sHVZK` and is nearly as efficient as the original protocol:

$$\begin{aligned}
P'_1(x, w) &\stackrel{\text{def}}{=} (r, \text{state}) \leftarrow P_1(x, w); c' \stackrel{\$}{\leftarrow} \{0, 1\}^k; \text{return } ((r, c'), (\text{state}, c')) \\
P'_2(x, w, (\text{state}, c'), c) &\stackrel{\text{def}}{=} s \leftarrow P_2(x, w, \text{state}, c \oplus c'); \text{return } s \\
V'_1(x, (r, c')) &\stackrel{\text{def}}{=} c \leftarrow V_1(x, r); \text{return } (c \oplus c') \\
V'_2(x, (r, c'), c, s) &\stackrel{\text{def}}{=} b \leftarrow V_2(x, r, c \oplus c', s); \text{return } b
\end{aligned}$$

Essentially, the construction creates a new protocol for which `HVZK` and `sHVZK` coincide. The difference is that in the new protocol the challenge that the verifier chooses is xor-ed with a random bitstring sampled by the prover at the beginning of the protocol.

**Theorem 7.3 (sHVZK from HVZK).** *If a protocol  $(P, V)$  is a  $\Sigma$ -protocol as in Definition 7.1 but satisfying `HVZK` instead of `sHVZK`, then the protocol  $(P', V')$  defined above is a  $\Sigma$ -protocol.*

*Proof.*

*Completeness*

Follows easily from the completeness of protocol  $(P, V)$  and the absorption property of the exclusive or operator, i.e.,  $(c \oplus c') \oplus c' = c$ .

*Special Honest Verifier Zero-Knowledge*

The following is a `sHVZK` simulator for the protocol

$$\text{Simulator } S'(x, c) : (\hat{r}, \hat{c}, \hat{s}) \leftarrow S(x); \text{return } ((\hat{r}, c \oplus \hat{c}), \hat{s})$$

(The variables of the original protocol are decorated with a hat.) We prove this by means of a sequence of program equivalences,

$$\begin{aligned}
\text{Protocol}'(x, w, c) &\simeq_{\{r, c, s\}}^{\{x, c\} \wedge R(x, w)} \text{Protocol}(x, w); r \leftarrow (\hat{r}, c \oplus \hat{c}); s \leftarrow \hat{s} \\
&\simeq_{\{r, c, s\}}^{\{x, c\} \wedge R(x, w)} (\hat{r}, \hat{c}, s) \leftarrow \mathbf{S}(x); r \leftarrow (\hat{r}, c \oplus \hat{c}) \\
&\simeq_{\{r, c, s\}}^{\{x, c\} \wedge R(x, w)} (r, s) \leftarrow \mathbf{S}'(x, c)
\end{aligned}$$

The first and last equivalences are proved without much difficulty using the program transformation tactics described in Chapter 3, while the second can be reduced to the HVZK of  $\mathbf{S}$  using the `alloc` and `eqobs_t1` tactics to simplify the goal.

### Soundness

From a conversation  $((r, c'), (c \oplus c'), s)$  of  $(\mathbf{P}', \mathbf{V}')$  a conversation  $(r, c, s)$  of the original protocol can be trivially recovered. Thus, the following knowledge extractor proves special soundness of  $(\mathbf{P}', \mathbf{V}')$ :

$$\text{KE}'(x, (r, c'), c_1, c_2, s_1, s_2) : w \leftarrow \text{KE}(x, r, c' \oplus c_1, c' \oplus c_2, s_1, s_2); \text{return } w$$

□

## 7.2 Sigma Protocols Based on Special Homomorphisms

An important class of  $\Sigma$ -protocols are the so-called  $\Sigma^\phi$ -protocols, that prove knowledge of a preimage under a homomorphism. The Schnorr protocol [Schnorr 1991], one of the most archetypal zero-knowledge proofs, is an instance of a  $\Sigma^\phi$ -protocol that proves knowledge of a discrete logarithm in a cyclic group, i.e. the homomorphism is in this case exponentiation,  $\phi(x) = g^x$ , where  $g$  is a generator of the group.

Our formalization of  $\Sigma^\phi$ -protocols is constructive. We provide a functor that, given a homomorphism  $\phi$  together with proofs that it satisfies certain properties, builds a concrete  $\Sigma$ -protocol for proving knowledge of a preimage under  $\phi$ . This protocol comes with proofs of completeness, soundness, and sHVZK. Thus, all that it takes to build an instance of a  $\Sigma^\phi$ -protocol is to specify a homomorphism and prove that it has the necessary properties. In this way, we give several examples of  $\Sigma^\phi$ -protocols, including the Schnorr, Guillou-Quisquater and Feige-Fiat-Shamir protocols. Although using the  $\Sigma^\phi$  construction spares us the hassle of proving each time the properties in Definition 7.1, these instantiations remain non-trivial because one needs to formalize the homomorphisms themselves, which in turn requires to give representations of the groups over which they are defined.

In the remaining of this subsection we let  $(G, \oplus)$  be a finite additive group and  $(H, \otimes)$  a multiplicative group.

**Definition 7.4 ( $\Sigma^\phi$ -protocol).** *Let  $\phi : G \rightarrow H$  be a homomorphism, and define  $R \stackrel{\text{def}}{=} \{(x, w) \mid x = \phi(w)\}$ . The  $\Sigma^\phi$ -protocol for relation  $R$  with challenge set  $C \subset \mathbb{N}$  is the  $\Sigma$ -protocol  $(\mathbf{P}, \mathbf{V})$  defined as follows:*

$$\begin{aligned}
P_1(x, w) &\stackrel{\text{def}}{=} y \xleftarrow{\$} G; \text{ return } (\phi(y), y) \\
P_2(x, w, y, c) &\stackrel{\text{def}}{=} \text{ return } (y \oplus cw) \\
V_1(x, r) &\stackrel{\text{def}}{=} c \xleftarrow{\$} C; \text{ return } c \\
V_2(x, r, c, s) &\stackrel{\text{def}}{=} \text{ return } (\phi(s) = r \otimes x^c)
\end{aligned}$$

It can be shown that the above protocol satisfies the properties of a  $\Sigma$ -protocol when  $C = \{0, 1\}$ . However, a cheating prover could convince the verifier with probability  $1/2$ ; this probability may be reduced to  $1/2^n$  (at the cost of efficiency) by repeating the protocol  $n$  rounds. We will see that a certain class of homomorphisms defined below admits a much larger challenge set, and thus achieves a lower soundness error in a single execution of the protocol.

**Definition 7.5 (Special homomorphism).** *An homomorphism  $\phi : G \rightarrow H$  is special if there exists a value  $v \in \mathbb{Z} \setminus \{0\}$  (called special exponent) and a PPT algorithm that given  $x \in H$  computes  $u \in G$  such that  $\phi(u) = x^v$ .*

To formalize  $\Sigma^\phi$ -protocols, we extended the language of CertiCrypt with types for the groups  $G, H$  and operators for computing the group operation, exponentiation/product, and inverse; we also added operators  $\phi(\cdot)$ ,  $u(\cdot)$ , and a constant expression  $v$  denoting the special exponent of the homomorphism as in Definition 7.5. In addition, we wrote an expression normalizer that simplifies arithmetic expressions by applying the homomorphic property of  $\phi$ ; normalization is done as part of the `ep` tactic.

A  $\Sigma^\phi$ -protocol built from a special homomorphism admits as a challenge set any natural interval of the form  $[0..c^+]$ , where  $c^+$  is smaller than the smallest prime divisor of the special exponent  $v$ .

**Theorem 7.6 ( $\Sigma^\phi$ -protocols for special homomorphisms).** *If  $\phi$  is special and  $c^+$  is smaller than any prime divisor of the special exponent  $v$ , then the protocol in Definition 7.4 is a  $\Sigma$ -protocol with challenge set  $C = [0..c^+]$ .*

*Proof.*

*Completeness*

We must prove that a honest prover always succeeds in convincing a verifier, i.e.

$$\forall m. R(m(x), m(w)) \implies \Pr[\text{Protocol}, m : b = \text{true}] = 1$$

Note that this can be reformulated in terms of a program equivalence as follows

$$\text{Protocol}(x, w) \simeq_{\{b\}}^{R(x, w)} b \leftarrow \text{true}$$

To prove this, we inline all procedure calls in the protocol and simplify the resulting program performing expression propagation, normalization, and dead code elimination. We use the following proof script:


```
inline P1; inline P2; inline V1; inline V2; ep; deadcode.
```



The resulting goal has the form

$$y \stackrel{\$}{\leftarrow} G; c \stackrel{\$}{\leftarrow} [0..c^+]; b \leftarrow (\phi(y) \otimes \phi(w)^c = \phi(y) \otimes x^c) \simeq_{\{b\}}^{\phi(w)=x} b \leftarrow \text{true}$$

We use the tactic `ep_eq x  $\phi(w)$`  to simplify the last instruction in the game on the left hand side to `b  $\leftarrow$  true`, tactic `deadcode` to remove the first two instructions that are no longer relevant, and `eqobs_in` to conclude.



**Fig. 7.2.** A game-based proof that  $S$  is a sHVZK simulator for the  $\Sigma^\phi$ -protocol in Theorem 7.6.

### Special Honest Verifier Zero-Knowledge

The following is a sHVZK simulator for the protocol:

$$\text{Simulator } S(x, c) : s \stackrel{\$}{\leftarrow} G; r \leftarrow \phi(s) \otimes x^{-c}; \text{return } (r, c, s)$$

A proof that  $S$  perfectly simulates conversations of the protocol is illustrated in Figure 7.2; we briefly explain the numbered steps in the figure.

1. Similarly to the proof above, we inline calls to  $P_1$  and  $P_2$ , and simplify the goal using tactics `ep` and `deadcode`.

2. We introduce an intermediate game using the transitivity of observational equivalence. To prove that the new game is observationally equivalent to the previous one, we first reorder the instructions using `swap` to obtain a common suffix which we then remove using the tactic `eqobs_tl`. The resulting goal is

$$\vdash y \not\Leftarrow G \simeq_{\substack{\{x,w,c\} \wedge R(x,w) \\ \{y,w,c\}}} s' \not\Leftarrow G; y \leftarrow s' \oplus -cw$$

Since variables  $w$  and  $c$  are not modified, we can remove them from the output set using tactic `clean_nm`. We next use tactic `alloc y s'` to sample  $s'$  instead of  $y$  in the game on the left, and we weaken the pre-condition to true, which results in the goal

$$\vdash s' \not\Leftarrow G; y \leftarrow s' \simeq_{\{y\}} s' \not\Leftarrow G; y \leftarrow s' \oplus -cw$$

This equivalence holds because  $-cw$  acts as a one-time pad; we have proved this as a lemma called `sum_otp` that we apply to conclude the proof.

3. Using `ep`, we propagate throughout the code the value assigned to  $y$  and then remove the assignment using `deadcode`. The expression normalizer automatically simplifies  $(s' \oplus -cw) \oplus cw$  to  $s'$ , and  $\phi(s' \oplus -cw)$  to  $\phi(s') \otimes \phi(w)^{-c}$  using the homomorphic property of  $\phi$ .
4. We introduce a new intermediate game; to prove that is equivalent to the previous one, we allocate variable  $s$  into  $s'$ ; the resulting game is identical to the one on the left hand side.
5. We substitute variable  $s$  for  $s'$  in the game on the right hand side of the equivalence, and use the pre-condition  $R(x, w)$ —which boils down to  $x = \phi(w)$ —to substitute  $x$  by  $\phi(w)$ . The resulting games are identical modulo reordering of instructions.
6. We conclude by inlining the call to `S` in the simulation.

### Soundness

Soundness requires the existence of an algorithm `KE` that given two accepting conversations  $(x, r, c_1, s_1)$ ,  $(x, r, c_2, s_2)$ , with  $c_1 \neq c_2$ , efficiently computes a  $w$  such that  $x = \phi(w)$ . We propose the following knowledge extractor:

```

KE(x, c1, c2, s1, s2) :
(a, b, d) ← extended_gcd(c1 - c2, v);
w ← a(s1 ⊕ -s2) ⊕ b u(x);
return w

```

where `extended_gcd` efficiently implements the extended Euclidean algorithm. For integers  $a, b$ , `extended_gcd(a, b)` computes a triple of integers  $(x, y, d)$  such that  $d$  is the greatest common divisor of  $a$  and  $b$ , and  $x, y$  satisfy the Bézout's identity

$$ax + by = \gcd(a, b) = d$$

Since all computations done by the knowledge extractor can be efficiently implemented, `KE` is a PPT algorithm. We have to prove as well that `KE` computes a

preimage of the public input  $x$ . For two accepting conversations  $(x, r, c_1, s_1)$  and  $(x, r, c_2, s_2)$ , we have

$$\phi(s_1) = r \otimes x^{c_1} \wedge \phi(s_2) = r \otimes x^{c_2}$$

and thus

$$x^{c_1 - c_2} = \phi(s_1 \oplus -s_2) \quad (7.1)$$

Furthermore, since  $\phi$  is special we can efficiently compute  $u$  such that  $x^v = \phi(u)$ . The triple  $(a, b, d)$  given by the extended Euclidean algorithm satisfies the Bézout's identity

$$a(c_1 - c_2) + bv = \gcd(c_1 - c_2, v) = d \quad (7.2)$$

Both  $c_1$  and  $c_2$  are bounded by  $c^+$ , which is in turn smaller than the smallest prime that divides  $v$ . Thus, no divisor of  $|c_1 - c_2|$  can divide  $v$  and

$$d = \gcd(c_1 - c_2, v) = 1$$

In addition, since  $\phi$  is a homomorphism, from (7.1) and (7.2) we conclude

$$\phi(w) = \phi(a(s_1 \oplus -s_2) \oplus bu) = x^{a(c_1 - c_2)} \otimes x^{bv} = x$$

□

### 7.2.1 Concrete Instances of Sigma-Phi Protocols

We have formalized several  $\Sigma^\phi$ -protocols using the functor described in the previous section. For each protocol, we specify the groups  $G, H$  and the underlying special homomorphism  $\phi : G \rightarrow H$ , and provide appropriate interpretations for the operator  $u(\cdot)$  and the constant special exponent  $v$ . Table 7.1 summarizes all the protocols that we have formalized.

**Table 7.1.** Special homomorphisms in selected  $\Sigma^\phi$ -protocols. In the table,  $\mathbb{Z}_q^+$  stands for the additive group of integers modulo  $q$ ,  $\mathbb{Z}_p^*$  for the multiplicative group of integers modulo  $p$ ;  $N$  is an RSA modulus and  $e$  a public RSA exponent coprime with  $\varphi(N)$ .

Protocol	$G$	$H$	$\phi$	$u$	$v$
Schnorr	$\mathbb{Z}_q^+$	$\mathbb{Z}_p^*$	$x \mapsto g^x$	$x \mapsto 0$	$q$
Okamoto	$(\mathbb{Z}_q^+, \mathbb{Z}_q^+)$	$\mathbb{Z}_p^*$	$(x_1, x_2) \mapsto g_1^{x_1} \otimes g_2^{x_2}$	$x \mapsto (0, 0)$	$q$
Fiat-Shamir	$\mathbb{Z}_N^*$	$\mathbb{Z}_N^*$	$x \mapsto x^2$	$x \mapsto x$	2
Guillou-Quisquater	$\mathbb{Z}_N^*$	$\mathbb{Z}_N^*$	$x \mapsto x^e$	$x \mapsto x$	$e$
Feige-Fiat-Shamir	$\{-1, 1\} \times \mathbb{Z}_N^*$	$\mathbb{Z}_N^*$	$(s, x) \mapsto s x^2$	$x \mapsto (1, x)$	2

The Schnorr [1991] and Okamoto [1993] protocols are based on the discrete logarithm problem. For prime numbers  $p$  and  $q$  such that  $q$  divides  $p - 1$ , a Schnorr group is a multiplicative subgroup of  $\mathbb{Z}_p^*$  of order  $q$  with generator  $g$ . A  $\Sigma$ -protocol

for proving knowledge of discrete logarithms in the Schnorr group is obtained by instantiating the construction of Definition 7.4 with the homomorphism

$$\begin{aligned}\phi & : \mathbb{Z}_q^+ \rightarrow \mathbb{Z}_p^* \\ \phi(x) & = g^x\end{aligned}$$

Since the order  $q$  of the Schnorr group is known, it suffices to take  $q$  as the special exponent of the homomorphism, and  $u(x) = 0$  for all  $x \in \mathbb{Z}_q^+$ . The Okamoto protocol is similar to Schnorr protocol but it works with two Schnorr subgroups of  $\mathbb{Z}_p^*$  with generators  $g_1$  and  $g_2$ , respectively (it can be naturally generalized to any number of generators). In this case  $\phi$  maps a pair  $(x_1, x_2)$  to  $g_1^{x_1} \otimes g_2^{x_2}$ .

Let  $N$  be an RSA modulus with prime factors  $p$  and  $q$ , and let  $e$  be a public exponent;  $e$  must be co-prime with the totient  $\varphi = (p-1)(q-1)$  (i.e.  $\gcd(e, \varphi(N)) = 1$ ). The Guillou-Quisquater [Guillou and Quisquater 1988], Fiat-Shamir [Fiat and Shamir 1987], and Feige-Fiat-Shamir [Feige et al. 1988] protocols are based on the difficulty of solving the RSA problem: given  $N$ ,  $e$  and  $y \equiv x^e \pmod N$ , compute  $x$ , the  $e^{\text{th}}$ -root of  $y$  modulo  $N$ .

The Guillou-Quisquater protocol is obtained by taking

$$\begin{aligned}\phi & : \mathbb{Z}_N^* \rightarrow \mathbb{Z}_N^* \\ \phi(x) & = x^e\end{aligned}$$

The Fiat-Shamir protocol is obtained as a special case when  $e = 2$ . The Feige-Fiat-Shamir is obtained by taking

$$\begin{aligned}\phi & : \{-1, 1\} \times \mathbb{Z}_N^* \rightarrow \mathbb{Z}_N^* \\ \phi(s, x) & = s x^2\end{aligned}$$

*Remark*

We note that our results hold independently of any computational assumption. Certainly, it is the difficulty of inverting the underlying homomorphism what makes a  $\Sigma^\phi$ -protocol interesting, but this is inessential for establishing the properties we prove about the protocol.

### 7.2.2 Composition of Sigma-Phi Protocols

Let  $\phi_1 : G_1 \rightarrow H_1$  and  $\phi_2 : G_2 \rightarrow H_2$  be two special homomorphisms with special exponents  $v_1, v_2$  and associated algorithms  $u_1, u_2$ , respectively. We give below two useful ways of combining the  $\Sigma^\phi$ -protocols induced by these homomorphisms.

**Theorem 7.7 (Direct product of special homomorphisms).** *The following homomorphism from the direct product of  $G_1$  and  $G_2$  to the direct product of  $H_1$  and  $H_2$  is a special homomorphism:*

$$\begin{aligned}\phi & : G_1 \times G_2 \rightarrow H_1 \times H_2 \\ \phi(x_1, x_2) & \stackrel{\text{def}}{=} (\phi_1(x_1), \phi_2(x_2))\end{aligned}$$

*Proof.* It suffices to take

$$\begin{aligned} v &\stackrel{\text{def}}{=} \text{lcm}(v_1, v_2) \\ u(x_1, x_2) &\stackrel{\text{def}}{=} (u_1(x_1)^{v/v_1}, u_2(x_2)^{v/v_2}) \end{aligned}$$

Indeed,

$$\begin{aligned} \phi(u(x_1, x_2)) &= (\phi_1(u_1(x_1)^{v/v_1}), \phi_2(u_2(x_2)^{v/v_2})) \\ &= (x_1^{v_1 v/v_1}, x_2^{v_2 v/v_2}) \\ &= (x_1, x_2)^v \end{aligned}$$

□

This yields an effective means of AND-combining assertions proved by  $\Sigma^\phi$ -protocols. The result generalizes the protocol of Maurer [2009, Theorem 6.2]; we do not require that the special exponent be the same.

**Theorem 7.8 (Equality of preimages).** *Suppose that the domain of both homomorphisms is the same,  $G_1 = G_2 = G$ ,  $v_1 = v_2$ , and  $u_1, u_2$  are such that*

$$\forall x_1 \in H_1, x_2 \in H_2. u_1(x_1) = u_2(x_2)$$

*Then, the following homomorphism from  $G$  to the direct product of  $H_1$  and  $H_2$  is a special homomorphism:*

$$\begin{aligned} \phi &: G \rightarrow H_1 \times H_2 \\ \phi(x) &\stackrel{\text{def}}{=} (\phi_1(x), \phi_2(x)) \end{aligned}$$

*Proof.* Take  $v \stackrel{\text{def}}{=} v_1$  and  $u(x_1, x_2) \stackrel{\text{def}}{=} u_1(x_1)$ ,

$$\begin{aligned} \phi(u(x_1, x_2)) &= (\phi_1(u_1(x_1)), \phi_2(u_2(x_2))) \\ &= (x_1^{v_1}, x_2^{v_2}) \\ &= (x_1, x_2)^v \end{aligned}$$

□

We can use this latter theorem to construct a  $\Sigma$ -protocol that proves correctness of Diffie-Hellman keys. Given a group with prime order  $q$  and a generator  $g$ , this amounts to prove that triples of group elements of the form  $(\alpha, \beta, \gamma)$  are Diffie-Hellman triples, i.e. that if  $\alpha = g^a$  and  $\beta = g^b$ , then  $\gamma = g^{ab}$ . We instantiate the above construction for homomorphisms  $\phi_1(x) = g^x$ , and  $\phi_2(x) = \beta^x$ . Knowledge of a preimage  $a$  of  $(\alpha, \gamma)$  implies that  $(\alpha, \beta, \gamma)$  is a Diffie-Hellman triple (and thus that  $\gamma$  is a valid Diffie-Hellman shared key).

### 7.3 Sigma Protocols Based on Claw-Free Permutations

This section describes a general construction in the same flavor as the  $\Sigma^\phi$  construction discussed in the previous section, but based on pairs of *claw-free* permutations [Cramer 1996] rather than on special homomorphisms.

**Definition 7.9 (Claw-free permutation pair).** A pair of trapdoor permutations  $f = (f_0, f_1)$  on the same domain  $D$  is claw-free if it is unfeasible to compute  $x, y \in D$  such that  $f_0(x, pk) = f_1(y, pk)$ .

Given a claw-free permutation pair  $f$ , and a bitstring  $a \in \{0, 1\}^k$ , we define

$$f_{[a]}(b) \stackrel{\text{def}}{=} f_{a_1}(f_{a_2}(\dots(f_{a_k}(b))\dots))$$

where  $a_i$  denotes the  $i^{\text{th}}$  bit of  $a$ .

**Theorem 7.10 ( $\Sigma$ -protocol based on claw-free permutations).** Let  $(f_0, f_1)$  be a pair of claw-free permutations on  $D$  and let  $R$  be such that

$$R(pk, sk) \iff \forall x. f_0^{-1}(sk, f_0(pk, x)) = x \wedge f_1^{-1}(sk, f_1(pk, x)) = x$$

The following protocol is a  $\Sigma$ -protocol for relation  $R$ :

$$\begin{aligned} P_1(pk, sk) &\stackrel{\text{def}}{=} y \xleftarrow{\$} D; \text{ return } (y, y) \\ P_2(pk, sk, y, c) &\stackrel{\text{def}}{=} \text{ return } f_{[c]}^{-1}(sk, y) \\ V_1(pk, r) &\stackrel{\text{def}}{=} c \xleftarrow{\$} \{0, 1\}^k; \text{ return } c \\ V_2(pk, r, c, s) &\stackrel{\text{def}}{=} \text{ return } (f_{[c]}(pk, s) = r) \end{aligned}$$

except that it might not satisfy the knowledge soundness property.

*Proof.*

*Completeness*

The proof follows almost the same structure as the completeness proof for  $\Sigma^\phi$ -protocols. After inlining procedure calls in the protocol, we are left with the goal

$$b \leftarrow f_{[c]}(pk, f_{[c]}^{-1}(sk, y)) = y \simeq_{\{b\}}^{R(pk, sk)} b \leftarrow \text{true}$$

We use the fact that the pair  $(pk, sk)$  is in  $R$  to prove that  $f_{[c]}(pk, f_{[c]}^{-1}(sk, y)) = y$  by induction on  $c$ .

*Special Honest Verifier Zero-Knowledge*

The following is a sHVZK simulator for the protocol,

$$\mathbf{Simulator } S(pk, c) : s \xleftarrow{\$} D; r \leftarrow f_{[c]}(pk, s); \text{ return } (r, s)$$

To prove that

$$\text{Protocol}(pk, sk, c) \simeq_{\{r, c, s\}}^{\{pk, c\} \wedge R(pk, sk)} (r, s) \leftarrow S(pk, c)$$

we inline every procedure call in both games and perform expression propagation and dead code elimination, we are left with the following goal:

$$\vdash r \stackrel{\$}{\leftarrow} D; s \leftarrow f_{[c]}^{-1}(sk, r) \simeq_{\{pk, sk, c\}}^{\{r, s\}} r \leftarrow f_{[c]}(pk, s)$$

which is provable from the fact that  $f$  is a permutation pair.  $\square$

We observed that the above protocol does not necessarily satisfy the special soundness property. Instead, it satisfies a property known as *collision intractability*: no efficient algorithm can find two accepting conversations with different challenges but same commitment (a collision) with non-negligible probability. Interactive proof protocols that are complete, sHVZK but only satisfy *collision intractability* have important applications as signature protocols.

**Theorem 7.11.** *It is unfeasible to find a collision for the protocol in Theorem 7.10.*

*Proof.* By contradiction. Assume two accepting conversations  $(r, c_1, s_1)$ ,  $(r, c_2, s_2)$  for a public input  $pk$  with  $c_1 \neq c_2$ . We show that it is possible to efficiently compute a claw  $(b, b')$  such that  $f_0(b) = f_1(b')$ . Since the two conversations are accepting,

$$f_{[c_1]}(pk, s_1) = f_{[c_2]}(pk, s_2) = r$$

The following algorithm computes a claw

```

find_claw( $s_1, c_1, s_2, c_2$ ) :
  if head( $c_1$ ) = head( $c_2$ )
  then find_claw(tail( $c_1$ ),  $s_1$ , tail( $c_2$ ),  $s_2$ )
  else if head( $c_1$ ) = 0
    then ( $f_{[\text{tail}(c_1)]}(pk, s_1), f_{[\text{tail}(c_2)]}(pk, s_2)$ )
    else ( $f_{[\text{tail}(c_2)]}(pk, s_2), f_{[\text{tail}(c_1)]}(pk, s_1)$ )

```

The algorithm executes in polynomial-time provided permutations  $f_0$  and  $f_1$  can be evaluated in polynomial time, and  $c_1, c_2$  are polynomially bounded. For a polynomially bounded challenge set, this contradicts the assumption that  $(f_0, f_1)$  is claw-free.  $\square$

## 7.4 Combination of Sigma-Protocols

There are two immediate, but essential, ways of combining two  $\Sigma$ -protocols  $(P^1, V^1)$  and  $(P^2, V^2)$  with knowledge relations  $R_1$  and  $R_2$  respectively: AND-combination, and OR-combination. The former allows a prover to prove knowledge of witnesses  $w_1, w_2$  such that  $R_1(x_1, w_1)$  and  $R_2(x_2, w_2)$ . The latter allows a prover to prove knowledge of a witness  $w$  such that either  $R_1(x_1, w)$  or  $R_2(x_2, w)$ , without revealing which is the case. This can be naturally extended to proofs of any monotone Boolean formula by nested combination (although there exist a direct, more efficient construction based on secret-sharing schemes, cf. [Cramer 1996]). Even though simple, such constructions are incredibly powerful and form the basis of many practical protocols, like secure electronic voting protocols.

## 7.4.1 AND-Combination

Two  $\Sigma$ -protocols can be combined into a  $\Sigma$ -protocol that proves simultaneous knowledge of witnesses for both underlying knowledge relations, i.e. a  $\Sigma$ -protocol with a knowledge relation:

$$R \stackrel{\text{def}}{=} \{((x_1, x_2), (w_1, w_2)) \mid R_1(x_1, w_1) \wedge R_2(x_2, w_2)\}$$

We have formalized a functor AND, that combines two public-coin  $\Sigma$ -protocols  $(P^1, V^1)$  and  $(P^2, V^2)$  in this form. Without loss of generality, we assume that both protocols mandate that honest verifiers choose challenges uniformly from a set of bitstrings of a certain length  $k$ . The construction is straightforward; the combination is essentially a parallel composition of the two sub-protocols using the same randomly chosen challenge:

```

P1((x1, x2), (w1, w2))  $\stackrel{\text{def}}{=}$ 
  (r1, state1)  $\leftarrow$  P11(x1, w1);
  (r2, state2)  $\leftarrow$  P12(x2, w2);
  return ((r1, r2), (state1, state2))

P2((x1, x2), (w1, w2), state1, state2, c)  $\stackrel{\text{def}}{=}$ 
  s1  $\leftarrow$  P21(x1, w1, state1, c);
  s2  $\leftarrow$  P22(x2, w2, state2, c);
  return (s1, s2)

V1((x1, x2), (r1, r2))  $\stackrel{\text{def}}{=} c \leftarrow_{\$} \{0, 1\}^k$ ; return c

V2((x1, x2), (r1, r2), c, (s1, s2))  $\stackrel{\text{def}}{=}$ 
  b1  $\leftarrow$  V21(x1, r1, c, s1)
  b2  $\leftarrow$  V22(x2, r2, c, s2)
  return (b1 = true  $\wedge$  b2 = true)

```

Observe that  $V_1$  is not built from  $V_1^1$  and  $V_1^2$ . The reason for this is that in order to prove soundness, two runs of the protocol for the same public input  $x$  with the same commitment  $r$ , but with different challenges  $c \neq c'$ , must yield two runs of each of the sub-protocols with distinct challenges. If the challenge for the main protocol were built from the challenges computed by  $V_1^1$  and  $V_1^2$ , e.g. by concatenating them, we would not be able to conclude that the challenges in each pair of conversations extracted for the sub-protocols are different—one could only conclude that this is the case for one of the sub-protocols. Instead, we make use of the public-coin property and simply draw in  $V_1$  a new random challenge that is used in both sub-protocols. This solves the above problem, but also requires that the sub-protocols satisfy the special honest verifier zero-knowledge property, since we need to be able to simulate the sub-protocols for any fixed challenge.

Since AND combination essentially amounts to pairing the two sub-protocols while respecting the structure of a  $\Sigma$ -protocol, all proofs have the same general form: procedure calls are first inlined, and then the goal is manipulated using program transformations to put it in a form where the properties of the sub-protocols can be applied to conclude. We give below a proof sketch of sHVZK and special soundness; a more detailed proof of these properties and a proof of completeness can be found in [Barthe et al. 2010b].



*Special Honest-Verifier Zero-Knowledge*

The sHVZK simulator for the protocol simply runs the simulators of the sub-protocols to obtain a conversation for each sub-protocol with the same challenge  $c$ , the conversations are then combined to obtain a conversation of the main protocol:

**Simulator**  $S((x_1, x_2), c)$  :

$(r_1, s_1) \leftarrow S_1(x_1, c);$   
 $(r_2, s_2) \leftarrow S_2(x_2, c);$   
**return** $((r_1, r_2), (s_1, s_2))$

*Soundness*

Soundness requires us to give a PPT knowledge extractor that computes a witness for the knowledge relation  $R$  from two accepting runs of the protocol with different challenges but the same commitment. This amounts to computing a witness for each of the sub-protocols and can be done using the corresponding knowledge extractors as follows:

**KE** $((x_1, x_2), (r_1, r_2), c, c', (s_1, s_2), (s'_1, s'_2))$  :

$w_1 \leftarrow KE^1(x_1, r_1, c, c', s_1, s'_1);$   
 $w_2 \leftarrow KE^2(x_2, r_2, c, c', s_2, s'_2);$   
**return**  $(w_1, w_2)$

Note that an accepting conversation of the main protocol yields an accepting conversation for each one of the sub-protocols. Moreover, since the challenge of the main protocol is used as the challenge in both sub-protocols and  $c \neq c'$ , the extracted conversations have different challenges. Concretely, from

$$\Pr [b \leftarrow V_2((x_1, x_2), (r_1, r_2), c, (s_1, s_2)), m : b = \text{true}] = 1$$

$$\Pr [b \leftarrow V_2((x_1, x_2), (r_1, r_2), c', (s'_1, s'_2)), m : b = \text{true}] = 1$$

we can prove that for  $i = 1, 2$ ,

$$\Pr [w_i \leftarrow KE^i(x_i, r_i, c, c', s_i, s'_i), m : R_i(x_i, w_i)] = 1$$

from the soundness of the sub-protocols and from the fact that

$$\Pr [b_i \leftarrow V_2^i(x_i, r_i, c, s_i), m : b_i = \text{true}] = 1$$

$$\Pr [b_i \leftarrow V_2^i(x_i, r_i, c', s'_i), m : b_i = \text{true}] = 1$$

□

**7.4.2 OR-Combination**

Two  $\Sigma$ -protocols can also be combined to obtain a protocol that proves knowledge of a witness for the knowledge relation of one of the sub-protocols, but without revealing which. The construction relies on the ability to simulate accepting runs; the basic idea is that the prover runs the real protocol for which it knows a witness,

and uses the simulator to generate a run of the other protocol. The knowledge relation suggested by, e.g. Damgård [2010],

$$\hat{R} \stackrel{\text{def}}{=} \{((x_1, x_2), w) \mid R_1(x_1, w) \vee R_2(x_2, w)\}$$

suffers from placing unrealistic demands on the simulator. As pointed out by Cramer [1996], it is important to allow the simulator to fail on an input  $x \notin \text{dom}(R)$ . However, in order to prove completeness for the above relation, the simulator must be able to perfectly simulate outside the domain of the respective knowledge relation. Instead, we can prove completeness (and sHVZK) of the combination with respect to a knowledge relation whose domain is restricted to the Cartesian product of the domains of the knowledge relations of the sub-protocols, i.e.

$$R \stackrel{\text{def}}{=} \left\{ ((x_1, x_2), w) \mid \begin{array}{l} (R_1(x_1, w) \wedge x_2 \in \text{dom}(R_2)) \vee \\ (R_2(x_2, w) \wedge x_1 \in \text{dom}(R_1)) \end{array} \right\}$$

Unfortunately, we cannot prove soundness with respect to  $R$ , we can only prove it with respect to  $\hat{R}$ . The reason for this is that an accepting run of the combined protocol only guarantees the existence of a witness for the public input of one of the protocols, the simulation of the other protocol may succeed even if the input is not in the domain of the respective relation. Otherwise said, from two accepting runs of the combined protocol with distinct challenges we might not be able to extract two accepting runs with distinct challenges for each of the sub-protocols; we can only guarantee we can do that for one of them. Observe that we do not really lose anything by proving completeness with respect to the smaller relation  $R$ . If we admitted pairs  $(x_1, x_2)$  as public input where one component does not belong to the domain of the corresponding knowledge relation, we would not be able to say anything about the success of the simulator. The simulator might as well fail, trivially revealing that the prover could not have known a witness for the corresponding input, and rendering the protocol pointless for such inputs.

Compared to the AND combination, the OR combination is harder to fit into the structure of a  $\Sigma$ -protocol. The reason for this is that the first phase of the prover needs to use the simulator of one of the sub-protocols, which results in a full (accepting) conversation that has to be passed over to the second phase of the prover. Given  $R_1(x_1, w)$ , the OR prover runs the prover of the first protocol and the simulator of the second, and returns as a commitment a pair with the commitments of each protocol; it passes over in the state the challenge and the reply of the simulated conversation,

$$\begin{aligned}
P_1((x_1, x_2), w) &\stackrel{\text{def}}{=} \\
&\text{if } (x_1, w) \in R_1 \text{ then} \\
&\quad (r_1, state_1) \leftarrow P_1^1(x_1, w); \\
&\quad c_2 \xleftarrow{\$} \{0, 1\}^k; \\
&\quad (r_2, s_2) \leftarrow S_2(x_2, c_2); \\
&\quad state \leftarrow (state_1, c_2, s_2) \\
&\text{else} \\
&\quad (r_2, state_2) \leftarrow P_1^2(x_2, w); \\
&\quad c_1 \xleftarrow{\$} \{0, 1\}^k; \\
&\quad (r_1, s_1) \leftarrow S_1(x_1, c_1); \\
&\quad state \leftarrow (state_2, c_1, s_1) \\
&\text{return } ((r_1, r_2), state)
\end{aligned}$$

Above, the test  $(x_1, w) \in R_1$  is an encoding of the fact that the prover knows to which knowledge relation corresponds the witness  $w$ , and thus which protocol it can run for real, while simulating the other one. The commitment  $(r_1, r_2)$  is passed along to the verifier that simply replies by returning a randomly chosen bitstring to the prover, the combination is a public-coin protocol,

$$V_1((x_1, x_2), (r_1, r_2)) \stackrel{\text{def}}{=} c \xleftarrow{\$} \{0, 1\}^k; \text{ return } c$$

Assume without loss of generality that  $R_1(x_1, w)$ . In the second phase the prover constructs the challenge for the first protocol by xor-ing the challenge  $c$  of the OR protocol with the challenge used in the simulation of the second protocol in the first phase. It then runs the second phase of the prover of the first protocol to compute a reply. The result of the second phase is constructed from the challenges for each protocol and the prover replies (the ones coming from the simulated protocol are recovered from the state):

$$\begin{aligned}
P_2((x_1, x_2), w, (state, c', s), c) &\stackrel{\text{def}}{=} \\
&\text{if } (x_1, w) \in R_1 \text{ then} \\
&\quad state_1 \leftarrow state; c_2 \leftarrow c'; s_2 \leftarrow s; \\
&\quad c_1 \leftarrow c_2 \oplus c; \\
&\quad s_1 \leftarrow P_2^1(x_1, w, state_1, c_1) \\
&\text{else} \\
&\quad state_2 \leftarrow state; c_1 \leftarrow c'; s_1 \leftarrow s; \\
&\quad c_2 \leftarrow c_1 \oplus c; \\
&\quad s_2 \leftarrow P_2^2(x_2, w, state_2, c_2) \\
&\text{return } ((c_1, s_1), (c_2, s_2))
\end{aligned}$$

The verifier accepts a conversation when the runs of both protocols are accepting and the challenge is the xor of the challenges used in each of the combined protocols,

$$\begin{aligned}
V_2((x_1, x_2), (r_1, r_2), c, ((c_1, s_1), (c_2, s_2))) &\stackrel{\text{def}}{=} \\
b_1 &\leftarrow V_2^1(x_1, r_1, c_1, s_1); \\
b_2 &\leftarrow V_2^2(x_2, r_2, c_2, s_2); \\
\text{return } (c = c_1 \oplus c_2 \wedge b_1 = \text{true} \wedge b_2 = \text{true})
\end{aligned}$$

### Completeness

The proof is slightly more involved than the proof for the AND combination, since only one of the protocols is run for real, while the other is just simulated, and this depends on the knowledge of the prover. Thus, the proof is split into two cases:

- **case**  $(x_1, w) \in R_1$ : the outline of the proof is as follows:

$$\begin{aligned} \text{Protocol}((x_1, x_2), w) &\simeq \text{Protocol}_1(x_1, w); c_2 \xleftarrow{\$} \{0, 1\}^k; (r_2, s_2) \leftarrow S_2(x_2, c_2) \\ &\simeq \text{Protocol}_1(x_1, w); \text{Protocol}_2(x_2, w') \end{aligned}$$

The first equivalence is immediate from inlining procedure calls and simplification. The second equivalence follows from the fact that  $\exists w'. R_2(x_2, w')$  and the sHVZK property of the second protocol. The proof concludes by application of the completeness property of each of the sub-protocols.

- **case**  $(x_2, w) \in R_2$ : Idem.

#### *Special Honest-Verifier Zero-Knowledge*

The simulator for the OR combination is easily built from the simulators of the sub-protocols:

**Simulator**  $S((x_1, x_2), c)$  :

$c_2 \xleftarrow{\$} \{0, 1\}^k$ ;  
 $c_1 \leftarrow c \oplus c_2$ ;  
 $(r_1, s_1) \leftarrow S_1(x_1, c_1)$ ;  
 $(r_2, s_2) \leftarrow S_2(x_2, c_2)$ ;  
**return**  $((r_1, r_2), ((c_1, s_1), (c_2, s_2)))$

As before, the proof is split into two cases:

- **case**  $(x_1, w) \in R_1$ :

$$\begin{aligned} \text{Protocol}((x_1, x_2), w) &\simeq \text{Protocol}_1(x_1, w); S_2(x_2) \\ &\simeq S_1(x_1); S_2(x_2) \\ &\simeq S((x_1, x_2), c) \end{aligned}$$

Where the first and last steps are immediate from inlining, and simplification, whereas the second step is a direct application of the HVZK property of  $S_1$  (which follows from sHVZK by Theorem 7.2).

- **case**  $(x_2, w) \in R_2$ : Idem.

#### *Soundness*

(With respect to  $\hat{R}$ ). Unlike the AND combination, the OR combination does not have the property that runs with distinct challenges guarantee that the challenges used in the sub-protocols are also distinct. This is not as problematic as in the case of the AND combination, since it suffices to compute a  $w$  such that either  $R_1(x_1, w)$  or  $R_2(x_2, w)$ . Furthermore, from

$$c = c_1 \oplus c_2 \neq c' = c'_1 \oplus c'_2$$

we have either  $c_1 \neq c'_1$  or else  $c_1 = c'_1$ , in which case necessarily  $c_2 \neq c'_2$ . Thus, the knowledge extractor simply needs to do a case analysis:

```

KE((x1, x2), (r1, r2), c, c',
  ((c1, s1), (c2, s2)), ((c'1, s'1), (c'2, s'2))) :
if c1 ≠ c'1 then
  w ← KE1(x1, r1, c1, c'1, s1, s'1)
else
  w ← KE2(x2, r2, c2, c'2, s2, s'2)
return w

```

Assume two accepting runs of the combined protocol with the same commitment and  $c \neq c'$ :

$$\begin{aligned} &((x_1, x_2), (r_1, r_2), c, ((c_1, s_1), (c_2, s_2))) \\ &((x_1, x_2), (r_1, r_2), c', ((c'_1, s'_1), (c'_2, s'_2))) \end{aligned}$$

We have to establish that for an  $i \in 1, 2$ ,

$$\Pr [w_i \leftarrow \text{KE}^i(x_i, r_i, c_i, c'_i, s_i, s'_i) : R_i(x_i, w_i)] = 1$$

depending on whether  $c_1 \neq c'_1$  or  $c_1 = c'_1 \wedge c_2 \neq c'_2$ ,

- **case**  $c_1 \neq c'_1$ : From the special soundness of Protocol<sub>1</sub>,

$$\Pr [w_1 \leftarrow \text{KE}^1(x_1, r_1, c_1, c'_1, s_1, s'_1) : R_1(x_1, w_1)] = 1$$

- **case**  $c_1 = c'_1$  (and thus  $c_2 \neq c'_2$ ): From the special soundness of Protocol<sub>2</sub>

$$\Pr [w_2 \leftarrow \text{KE}^2(x_2, r_2, c_2, c'_2, s_2, s'_2) : R_2(x_2, w_2)] = 1$$

□

## 7.5 Related Work

Our work participates to an upsurge of interest in  $\Sigma$ -protocols, and shares some motivations and commonalities with recently published papers. Specifically, our account of  $\Sigma^\phi$ -protocols coincides with Maurer's [2009] unifying treatment of proofs of knowledge for preimages of group homomorphisms. Concretely, Maurer exhibits a main protocol that uses a group homomorphism—which in our setting corresponds to the definition of the module of  $\Sigma^\phi$ -protocols in Section 7.2—and shows (in his Theorem 3) that under suitable hypotheses the main protocol is a  $\Sigma$ -protocol. He gives several instances of the main protocol by picking suitable group homomorphisms and showing that they satisfy these hypotheses.

Our work is also closely connected to the recent effort of Bangerter et al. [2008, 2010] to design and implement efficient zero-knowledge proofs of knowledge. They provide both a set of sufficient conditions on a homomorphism  $\phi$  under which the corresponding  $\Sigma^\phi$ -protocol can be viewed as a  $\Sigma$ -protocol [Bangerter et al. 2008, Theorem 1], and a generalization that allows to consider sets of linear relations among preimages of group homomorphisms [Bangerter et al. 2008, Theorem 2]. The latter result is used to justify the soundness of a compiler that generates efficient code from high-level descriptions of protocols. As future work, the authors

of [Bangerter et al. 2008] mention that they plan to make their compiler certifying, so that it would generate proofs accompanying the code. Doing so from scratch remains a daunting task. By building on CertiCrypt, our formalization could be readily used as a stepping stone for a modular certifying compiler, in which (high-level, unoptimized) code is certified, and then compiled to efficient code using a certified or certifying compiler; see e.g. [Barthe et al. 2009b; Leroy 2006] for an instance of applying ideas from certified/certifying compilation to cryptography.

Cryptographic primitives need not only be secure; they must also be used correctly. In a series of papers, Backes, Hritcu, and Maffei [2008b]; Backes, Maffei, and Unruh [2008c] develop sound analysis methods for protocols that use zero-knowledge proofs, and apply their analyses to verify the authentication and secrecy properties of the Direct Anonymous Attestation Protocol. One extremely ambitious objective would be to use their results, which complement ours, to fully certify the security of the protocol in the computational model. Intermediate results would involve formalizing computational soundness results [Abadi and Rogaway 2002; Cortier and Warinschi 2005], which represents a substantial amount of work on its own.

## 7.6 Perspectives

We have presented a formalization of  $\Sigma$ -protocols in CertiCrypt. The highlights of our formalization are its generic account of the class of  $\Sigma^\phi$ -protocols and the detailed treatment of the AND/OR composition. Our work complements recent advances in the field, and takes a first but important step towards formalizing a rich theory of zero-knowledge proofs. In our opinion, and judging by the myriad of small variations in definitions we have found in the literature, this effort would be worth pursuing for a field that strives for definitional clarity and consistency.

Compared to other applications of CertiCrypt, like the verification of security proofs of encryption and signature schemes discussed in previous chapters, the formalization presented here imposes challenges of a different nature to the user. In contrast to earlier case studies, for which we have developed a mature set of techniques that mechanize most of the reasoning patterns appearing in proofs, we found that the formalization of  $\Sigma$ -protocols does not require as much complex reasoning, but is more demanding with respect to the compositionality of proofs. This led us to revise some design choices of CertiCrypt and has given us ideas on how to improve the framework so that results can be reused and composed more easily. For instance, when composing proofs of observational equivalence statements the user often needs to manually rename variables to match the names of the context where the proof is being reused; currently the user has to appeal to the `alloc` tactic to do this, but a simply heuristic may suffice in most cases.

We can build on the existing formalization to verify other important results about zero-knowledge proofs. These include other means of composing protocols: sequential [Goldreich and Oren 1994] and concurrent [Damgård 2000; Garay et al.

2006] composition; transforming public-coin zero-knowledge proofs in general zero-knowledge proofs [Goldreich 2002], or different formulations like non-interactive zero-knowledge proofs [Blum et al. 1988] or properties, i.e. statistical zero-knowledge and computational zero-knowledge instead of perfect zero-knowledge. Moreover,  $\Sigma$ -protocols form the base for a number of important and intriguing protocols for electronic voting schemes [Cramer 1996], identity schemes [Cramer 1996], and commitment schemes [Cramer 1996; Damgård 1990]. All are prime targets for future formalizations.



## Related Work and Conclusion

### 8.1 Related Work

Cryptographic protocol verification is an established area of formal methods, and a wealth of automated and deductive methods have been developed to the purpose of verifying that protocols provide the expected level of security [Meadows 2003]. Traditionally, protocols have been verified in a symbolic model, for which effective decision procedures exist under suitable hypotheses [Abadi and Cortier 2006]. Although the symbolic model assumes perfect cryptography, soundness results such as [Abadi and Rogaway 2002]—see [Cortier et al. 2010] for a recent survey—relate the symbolic model with the computational model, provided the cryptographic primitives satisfy adequate notions of security. It is possible to combine symbolic methods and soundness proofs to achieve guarantees in the computational model, as done e.g. in [Backes and Laud 2006; Backes et al. 2010; Sprenger and Basin 2008]. One drawback of this approach is that the security proof relies on intricate soundness proofs and hypotheses that unduly restrict the usage of primitives. Besides, it is not clear whether computational soundness results will always exist to allow factoring verification through symbolic methods [Backes and Pfitzmann 2005]. Consequently, some authors attempt to provide guarantees directly at the computational level [Blanchet 2008; Laud 2001; Roy et al. 2008].

In contrast, the formal verification of cryptographic functionalities is an emerging trend. An early work of Barthe et al. [2004] proves the security of ElGamal in Coq, but the proof relies on the generic model, a very specialized and idealized model that elides many of the issues that are relevant to cryptography. Den Hartog 2008 also proves ElGamal semantic security using a probabilistic (non-relational) Hoare logic. However, their formalism is not sufficiently powerful to express precisely security goals: notions such as well-formed and effective adversary are not modeled.

Blanchet and Pointcheval [2006] were among the first to use verification tools to carry out game-based proofs of cryptographic schemes. They used `CryptoVerif` to



prove the existential unforgeability of the FDH signature scheme, with the original security bound given in Section 5.1, which is much looser than the one given in Section 5.2. `CryptoVerif` has also been used to verify the security of many protocols, including Kerberos [Blanchet et al. 2008]. It is difficult to assess `CryptoVerif` ability to handle automatically more complex cryptographic proofs (or tighter security bounds), e.g. for schemes such as OAEP; on the other hand, compiling `CryptoVerif` sequences of games in `CertiCrypt` is an interesting research direction that would increase automation in `CertiCrypt` and confidence in `CryptoVerif`—by generating independently verifiable proofs.

Impagliazzo and Kapron [2006] were the first to develop a logic to reason about indistinguishability. Their logic is built upon a more general logic whose soundness relies on non-standard arithmetic; they show the correctness of a pseudo-random generator and that next-bit unpredictability implies pseudo-randomness. Recently, Zhang [2009] developed a similar logic on top of Hofmann’s SLR system [Hofmann 1998] and reconstructed the examples of Impagliazzo and Kapron [2006]. These logics have limited applicability because they lack support for oracles or adaptive adversaries and so cannot capture many of the the standard patterns for reasoning about cryptographic schemes. More recently Barthe et al. [2010a] developed a general logic, called Computational Indistinguishability Logic (CIL), that captures reasoning patterns that are common in provable security, such as simulation and reduction, and deals with oracles and adaptive adversaries. They use CIL to prove the security of the Probabilistic Signature Scheme, a widely used signature scheme that forms part of the PKCS standard [Bellare and Rogaway 1996]. CIL subsumes an earlier logic by Courant et al. [2008], who developed a form of strongest post-condition calculus that can establish automatically asymptotic security (IND-CPA and IND-CCA2) of encryption schemes that use one-way functions and hash functions modeled as random oracles. They show soundness and provide a prototype implementation that covers many examples in the literature.

In parallel, several authors have initiated formalizations of game-based proofs in proof assistants and shown the security of basic examples. Nowak [2007] gives a game-based proof of ElGamal semantic security in `Coq`. Nowak uses a shallow embedding to model games; his framework ignores complexity issues and has limited support for proof automation: because there is no special syntax for writing games, mechanizing syntactic transformations becomes very difficult. Affeldt et al. [2007] formalize a game-based proof of the PRP/PRF switching lemma in `Coq`. However, their formalization is tailored towards the particular example they consider, which substantially simplifies their task and hinders generality. They deal with a weak (non-adaptive) adversary model and ignore complexity. In another attempt to build a system supporting provable security, Backes et al. [2008a] formalize a language for games in the `Isabelle` proof assistant and prove the Fundamental Lemma; however, no examples are reported. All in all, these works appear like preliminary experiments that are not likely to scale.

Leaving the realm of cryptography, `CertiCrypt` relies on diverse mathematical concepts and theories that have been modeled for their own sake. We limit ourselves to singling out Audebaud and Paulin-Mohring [2009] formalization of the measure

monad, which we use extensively, and the work of Hurd et al. [2005], who developed a mechanized theory in the HOL theorem prover for reasoning about pGCL programs, a probabilistic extension of Dijkstra’s guarded command language.

## 8.2 Conclusion

CertiCrypt is a fully formalized framework that supports machine-checked game-based proofs; we have validated its design through formalizing standard cryptographic proofs. Our work shows that machine-checked proofs of cryptographic schemes are not only plausible but indeed feasible. However, constructing machine-checked proofs requires a high-level of expertise in formal proofs and remains time consuming despite the high level of automation achieved. Thus, CertiCrypt only provides a first step towards the completion of Halevi’s program, in spite of the amount of work invested so far (the project was initiated in June 2006). A medium-term objective would be to develop a minimalist interface that eases the writing of games and provides a fixed set of mechanisms (tactics, proof-by-pointing) to prove some basic transitions, leaving the side conditions as hypotheses. We believe that such an interface would help cryptographers ensure that there are no obvious flaws in their definitions and proofs, and to build sketches of security proofs. In fact, it is our experience that the type system and the automated tactics provide valuable information in debugging proofs.

Numerous research directions remain to be explored. Our main priority is to improve proof automation. In particular, we expect that one can automate many proofs in pRHL, by relying on a combination of standard verification tools: weakest pre-condition generators, invariant inference tools, SMT solvers.

In addition, it would be useful to formalize cryptographic meta-results such as the equivalence between IND-CPA and IND-CCA2 under plaintext awareness, or the transformation of an IND-CPA-secure scheme into an IND-CCA2-secure scheme [Fujisaki and Okamoto 1999]. Another direction would be to formalize proofs of computational soundness of the symbolic model, see e.g. [Abadi and Rogaway 2002] and proofs of automated methods for proving security of primitives and protocols, see e.g. [Courant et al. 2008; Laud 2001]. Finally, it would also be worthwhile to explore applications of CertiCrypt outside the realm cryptography, in particular to randomized algorithms and complexity.

### *Complexity and termination analysis of probabilistic programs*

CertiCrypt provides the necessary ingredients to reason about termination and complexity of programs. Yet cryptographic applications only make a limited use of them; e.g. we only use simple closure properties of PPT programs. It would be instructive to extend our formalization to define standard complexity classes and to prove the complexity of well-known probabilistic algorithms. More generally, we are interested in developing automated methods to carry such analyses for programs with loops and recursive calls.

*Reasoning about probabilistic programs*

The pWHILE language is sufficiently powerful to program widely used randomized algorithms, and it would be attractive to endow the formal semantics in Coq with a mechanized program logic that allows proving formally properties of these algorithms, in the spirit of the work of Hurd et al. [2005].

## References

- M. Abadi and V. Cortier. Deciding knowledge in security protocols under equational theories. *Theor. Comput. Sci.*, 367(1-2):2–32, 2006. Cited in page 133.
- M. Abadi and P. Rogaway. Reconciling two views of cryptography (The computational soundness of formal encryption). *J. Cryptology*, 15(2):103–127, 2002. Cited in pages 131, 133, and 135.
- R. Affeldt, M. Tanaka, and N. Marti. Formal proof of provable security by game-playing in a proof assistant. In *1st International conference on Provable Security, ProVSec 2007*, volume 4784 of *Lecture Notes in Computer Science*, pages 151–168, Berlin, 2007. Springer. Cited in pages 66 and 134.
- T. Amtoft, S. Bandhakavi, and A. Banerjee. A logic for information flow in object-oriented programs. In *33rd ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages, POPL 2006*, pages 91–102, New York, 2006. ACM. Cited in page 46.
- P. Audebaud and C. Paulin-Mohring. Proofs of randomized algorithms in Coq. *Sci. Comput. Program.*, 74(8):568–589, 2009. Cited in pages 2, 15, 17, and 134.
- M. Backes and P. Laud. Computationally sound secrecy proofs by mechanized flow analysis. In *13th ACM conference on Computer and Communications Security, CCS 2006*, pages 370–379, New York, 2006. ACM. Cited in page 133.
- M. Backes and B. Pfitzmann. Limits of the cryptographic realization of Dolev-Yao-style XOR. In *Computer Security – ESORICS 2005, 10th European symposium on Research in Computer Security*, volume 3679 of *Lecture Notes in Computer Science*, pages 178–196, Berlin, 2005. Springer. Cited in page 133.
- M. Backes, M. Berg, and D. Unruh. A formal language for cryptographic pseudocode. In *15th International conference on Logic for Programming, Artificial Intelligence and Reasoning, LPAR 2008*, volume 5330 of *Lecture Notes in Computer Science*, pages 353–376. Springer, 2008a. Cited in page 134.
- M. Backes, C. Hritcu, and M. Maffei. Type-checking zero-knowledge. In *15th ACM conference on Computer and Communications Security, CCS 2008*, pages 357–370. ACM, 2008b. Cited in page 131.
- 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, S&P 2008*, pages 202–215. IEEE Computer Society, 2008c. Cited in page 131.
- M. Backes, M. P. Grochulla, C. Hritcu, and M. Maffei. Achieving security despite compromise using zero-knowledge. In *22nd IEEE Computer Security Foundations symposium, CSF 2009*, pages 308–323. IEEE Computer Society, 2009. Cited in page 111.
- M. Backes, M. Maffei, and D. Unruh. Computationally sound verification of source code. In *17th ACM conference on Computer and Communications Security, CCS 2010*, New

- York, 2010. ACM. Cited in page 133.
- E. Bangerter, J. Camenisch, S. Krenn, A.-R. Sadeghi, and T. Schneider. Automatic generation of sound zero-knowledge protocols. Cryptology ePrint Archive, Report 2008/471, 2008. Cited in pages 130 and 131.
- E. Bangerter, J. Camenisch, and S. Krenn. Efficiency limitations for Sigma-protocols for group homomorphisms. In *7th Theory of Cryptography conference, TCC 2010*, volume 5978 of *Lecture Notes in Computer Science*, pages 553–571. Springer, 2010. Cited in page 130.
- G. Barthe, J. Cederquist, and S. Tarento. A machine-checked formalization of the generic model and the random oracle model. In *Automated Reasoning, 2nd International Joint conference, IJCAR 2004*, volume 3097 of *Lecture Notes in Computer Science*, pages 385–399, Berlin, 2004. Springer. Cited in page 133.
- G. Barthe, B. Grégoire, S. Heraud, and S. Zanella Béguelin. Formal certification of ElGamal encryption. A gentle introduction to CertiCrypt. In *5th International workshop on Formal Aspects in Security and Trust, FAST 2008*, volume 5491 of *Lecture Notes in Computer Science*, pages 1–19, Berlin, 2009a. Springer. Cited in pages 14 and 25.
- G. Barthe, B. Grégoire, C. Kunz, and T. Rezk. Certificate translation for optimizing compilers. *ACM Trans. Program. Lang. Syst.*, 31(5):1–45, 2009b. Cited in page 131.
- G. Barthe, B. Grégoire, and S. Zanella Béguelin. Formal certification of code-based cryptographic proofs. In *36th ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages, POPL 2009*, pages 90–101, New York, 2009c. ACM. Cited in pages 2, 25, 50, and 62.
- G. Barthe, M. Daubignard, B. Kapron, and Y. Lakhnech. Computational indistinguishability logic. In *17th ACM conference on Computer and Communications Security, CCS 2010*, New York, 2010a. ACM. Cited in page 134.
- G. Barthe, D. Hedin, S. Zanella Béguelin, B. Gregoire, and S. Heraud. A machine-checked formalization of Sigma-protocols. In *23rd IEEE Computer Security Foundations symposium, CSF 2010*, pages 246–260, Los Alamitos, Calif., 2010b. IEEE Computer Society. Cited in page 125.
- M. Bellare and P. Rogaway. Random oracles are practical: a paradigm for designing efficient protocols. In *1st ACM conference on Computer and Communications Security, CCS 1993*, pages 62–73, New York, 1993. ACM. Cited in pages 33, 67, 69, and 83.
- M. Bellare and P. Rogaway. Optimal asymmetric encryption. In *Advances in Cryptology – EUROCRYPT 1994*, volume 950 of *Lecture Notes in Computer Science*, pages 92–111, Berlin, 1994. Springer. Cited in page 88.
- M. Bellare and P. Rogaway. The exact security of digital signatures – How to sign with RSA and Rabin. In *Advances in Cryptology – EUROCRYPT 1996*, volume 1070 of *Lecture Notes in Computer Science*, pages 399–416, Berlin, 1996. Springer. Cited in pages 33, 67, 83, and 134.
- M. Bellare and P. Rogaway. The security of triple encryption and a framework for code-based game-playing proofs. In *Advances in Cryptology – EUROCRYPT 2006*, volume 4004 of *Lecture Notes in Computer Science*, pages 409–426, Berlin, 2006. Springer. Cited in pages 1, 2, 20, 50, 65, and 88.
- M. Bellare, D. Hofheinz, and E. Kiltz. Subtleties in the definition of IND-CCA: When and how should challenge-decryption be disallowed? Cryptology ePrint Archive, Report 2009/418, 2009. Cited in pages 33, 36, and 37.

- N. Benton. Simple relational correctness proofs for static analyses and program transformations. In *31st ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages, POPL 2004*, pages 14–25, New York, 2004. ACM. Cited in pages 41, 46, and 48.
- Y. Bertot, B. Grégoire, and X. Leroy. A structured approach to proving compiler optimizations based on dataflow analysis. In *Types for Proofs and Programs*, volume 3839 of *Lecture Notes in Computer Science*, pages 66–81, Berlin, 2006. Springer. Cited in page 48.
- B. Blanchet. A computationally sound mechanized prover for security protocols. In *27th IEEE symposium on Security and Privacy, S&P 2006*, pages 140–154. IEEE Computer Society, 2006. Cited in page 85.
- B. Blanchet. A computationally sound mechanized prover for security protocols. *IEEE Trans. Dependable Sec. Comput.*, 5(4):193–207, 2008. Cited in page 133.
- B. Blanchet and D. Pointcheval. Automated security proofs with sequences of games. In *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 537–554, Berlin, 2006. Springer. Cited in pages 85 and 133.
- B. Blanchet, A. D. Jaggard, A. Scedrov, and J.-K. Tsay. Computationally sound mechanized proofs for basic and public-key Kerberos. In *15th ACM conference on Computer and Communications Security, CCS 2008*, pages 87–99, New York, 2008. ACM. Cited in page 134.
- M. Blum, P. Feldman, and S. Micali. Non-interactive zero-knowledge and its applications. In *20th Annual ACM symposium on Theory of computing, STOC 1988*, pages 103–112. ACM, 1988. Cited in page 132.
- D. Boneh. Simplified OAEP for the RSA and Rabin functions. In *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 275–291, Berlin, 2001. Springer. Cited in page 88.
- D. Brown. What hashes make RSA-OAEP secure? Cryptology ePrint Archive, Report 2006/223, 2006. Cited in page 109.
- S. Cavallar, B. Dodson, A. Lenstra, W. Lioen, P. Montgomery, B. Murphy, H. te Riele, K. Aardal, J. Gilchrist, G. Guillerm, P. Leyland, J. Marchand, F. Morain, A. Muffett, C. Putnam, and P. Zimmermann. Factorization of a 512-bit RSA modulus. In *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 1–18, Berlin, 2000. Springer. Cited in page 81.
- J.-S. Coron. On the exact security of Full Domain Hash. In *Advances in Cryptology – CRYPTO 2000*, volume 1880 of *Lecture Notes in Computer Science*, pages 229–235. Springer, 2000. Cited in pages 75 and 83.
- J.-S. Coron. Optimal security proofs for PSS and other signature schemes. In *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 272–287, Berlin, 2002. Springer. Cited in page 83.
- V. Cortier and B. Warinschi. Computationally sound, automated proofs for security protocols. In *Programming Languages and Systems, 14th European symposium on Programming, ESOP 2005*, volume 3444 of *Lecture Notes in Computer Science*, pages 157–171. Springer, 2005. Cited in page 131.
- V. Cortier, S. Kremer, and B. Warinschi. A survey of symbolic methods in computational analysis of cryptographic systems. *J. Autom. Reasoning*, pages 1–35, 2010. Cited in page 133.

- J. Courant, M. Daubignard, C. Ene, P. Lafourcade, and Y. Lakhnech. Towards automated proofs for asymmetric encryption schemes in the random oracle model. In *15th ACM conference on Computer and Communications Security, CCS 2008*, pages 371–380, New York, 2008. ACM. Cited in pages 134 and 135.
- R. Cramer. *Modular Design of Secure yet Practical Cryptographic Protocols*. PhD thesis, CWI and Uni. of Amsterdam, 1996. Cited in pages 111, 114, 122, 124, 127, and 132.
- I. Damgård. On the existence of bit commitment schemes and zero-knowledge proofs. In *Advances in Cryptology – CRYPTO 1989*, volume 435 of *Lecture Notes in Computer Science*, pages 17–27. Springer, 1990. Cited in page 132.
- I. Damgård. Efficient concurrent zero-knowledge in the auxiliary string model. In *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 418–430. Springer, 2000. Cited in page 131.
- I. Damgård. On sigma-protocols. *Lecture Notes on Cryptologic Protocol Theory*, 2010. Cited in page 127.
- I. Damgård and B. Pfitzmann. Sequential iteration of interactive arguments and an efficient zero-knowledge argument for NP. In *Automata, Languages and Programming, 25th International Colloquium, ICALP 1998*, volume 1443 of *Lecture Notes in Computer Science*, pages 772–783. Springer, 1998. Cited in page 114.
- J. den Hartog. Towards mechanized correctness proofs for cryptographic algorithms: Axiomatization of a probabilistic Hoare style logic. *Sci. Comput. Program.*, 74(1-2):52–63, 2008. Cited in page 133.
- U. Feige, A. Fiat, and A. Shamir. Zero-knowledge proofs of identity. *J. Cryptology*, 1(2):77–94, 1988. Cited in page 121.
- A. Fiat and A. Shamir. How to prove yourself: practical solutions to identification and signature problems. In *Advances in Cryptology – CRYPTO 1986*, pages 186–194. Springer, 1987. Cited in page 121.
- E. Fujisaki and T. Okamoto. How to enhance the security of public-key encryption at minimum cost. In *2nd International workshop on Practice and Theory in Public Key Cryptography, PKC 1999*, volume 1560 of *Lecture Notes in Computer Science*, pages 634–634, Berlin, 1999. Springer. Cited in page 135.
- E. Fujisaki, T. Okamoto, D. Pointcheval, and J. Stern. RSA-OAEP is secure under the RSA assumption. *J. Cryptology*, 17(2):81–104, 2004. Cited in pages 88 and 101.
- J. A. Garay, P. MacKenzie, and K. Yang. Strengthening zero-knowledge protocols using signatures. *J. Cryptology*, 19:169–209, 2006. Cited in page 131.
- O. Goldreich. *Foundations of Cryptography: Basic Tools*, volume 1. Cambridge University Press, Cambridge, UK, 2001. Cited in page 28.
- O. Goldreich. Zero-knowledge twenty years after its invention. Technical Report TR02-063, Electronic Colloquium on Computational Complexity, 2002. Cited in pages 111 and 132.
- O. Goldreich and Y. Oren. Definitions and properties of zero-knowledge proof systems. *J. Cryptology*, 7(1):1–32, 1994. Cited in page 131.
- O. Goldreich, S. Goldwasser, and S. Micali. How to construct random functions. *J. ACM*, 33(4):792–807, 1986. Cited in page 58.

- S. Goldwasser and S. Micali. Probabilistic encryption. *J. Comput. Syst. Sci.*, 28(2): 270–299, 1984. Cited in pages 1 and 6.
- S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989. Cited in page 111.
- G. Gonthier. Formal Proof — The Four Colour Theorem. *Notices of the AMS*, 55(11): 1382–1393, 2008. Cited in page 4.
- L. Guillou and J.-J. Quisquater. A practical zero-knowledge protocol fitted to security microprocessor minimizing both transmission and memory. In *Advances in Cryptology – EUROCRYPT 1988*, volume 330 of *Lecture Notes in Computer Science*, pages 123–128. Springer, 1988. Cited in page 121.
- T. Hales. Formal Proof. *Notices of the AMS*, 55(11):1370–1380, 2008. Cited in page 4.
- T. Hales, J. Harrison, S. McLaughlin, T. Nipkow, S. Obua, and R. Zumkeller. A revision of the proof of the Kepler conjecture. *Discrete and Computational Geometry*, 44(1): 1–34, 2010. Cited in page 4.
- S. Halevi. A plausible approach to computer-aided cryptographic proofs. Cryptology ePrint Archive, Report 2005/181, 2005. Cited in pages 1 and 2.
- C. Hall, D. Wagner, J. Kelsey, and B. Schneier. Building PRFs from PRPs. In *Advances in Cryptology – CRYPTO 1998*, volume 1462 of *Lecture Notes in Computer Science*, pages 370–389. Springer, 1998. Cited in page 62.
- J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, 1999. Cited in page 14.
- M. Hofmann. A mixed modal/linear lambda calculus with applications to Bellantoni-Cook safe recursion. In *11th International workshop on Computer Science Logic, CSL 1997*, volume 1414 of *Lecture Notes in Computer Science*, pages 275–294, Berlin, 1998. Springer. Cited in page 134.
- J. Hurd, A. McIver, and C. Morgan. Probabilistic guarded commands mechanized in HOL. *Theor. Comput. Sci.*, 346(1):96–112, 2005. Cited in pages 135 and 136.
- R. Impagliazzo and B. M. Kapron. Logics for reasoning about cryptographic constructions. *J. Comput. Syst. Sci.*, 72(2):286–320, 2006. Cited in page 134.
- R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *21st Annual ACM symposium on Theory of Computing, 1989*, pages 44–61, New York, 1989. ACM. Cited in pages 63 and 65.
- B. Jonsson, W. Yi, and K. G. Larsen. Probabilistic extensions of process algebras. In J. Bergstra, A. Ponse, and S. Smolka, editors, *Handbook of Process Algebra*, pages 685–710. Elsevier, Amsterdam, 2001. Cited in pages 17, 18, and 19.
- J. Katz and N. Wang. Efficiency improvements for signature schemes with tight security reductions. In *10th ACM conference on Computer and Communications Security, CCS 2003*, pages 155–164, New York, 2003. ACM. Cited in pages 33 and 84.
- E. Kiltz and K. Pietrzak. On the security of padding-based encryption schemes — or — why we cannot prove OAEP secure in the standard model. In *Advances in Cryptology – EUROCRYPT 2009*, volume 5479 of *Lecture Notes in Computer Science*, pages 389–406, Berlin, 2009. Springer. Cited in page 109.
- E. Kiltz, A. O’Neill, and A. Smith. Instantiability of RSA-OAEP under chosen-plaintext attack. In *Advances in Cryptology – CRYPTO 2010*, volume 6223 of *Lecture Notes in*



- Computer Science*, pages 295–313, Berlin, 2010. Springer. Cited in page 109.
- G. Klein and T. Nipkow. A machine-checked model for a Java-like language, virtual machine and compiler. *ACM Trans. Program. Lang. Syst.*, 28(4):619–695, 2006. Cited in page 4.
- G. Klein, K. Elphinstone, G. Heiser, J. Andronick, D. Cock, P. Derrin, D. Elkaduwe, K. Engelhardt, R. Kolanski, M. Norrish, T. Sewell, H. Tuch, and S. Winwood. seL4: formal verification of an OS kernel. In *22nd ACM symposium on Operating Systems Principles, SOSP 2009*, pages 207–220. ACM Press, 2009. Cited in page 4.
- T. Kleinjung, K. Aoki, J. Franke, A. Lenstra, E. Thomé, J. Bos, P. Gaudry, A. Kruppa, P. Montgomery, D. Osik, H. te Riele, A. Timofeev, and P. Zimmermann. Factorization of a 768-bit RSA modulus. In *Advances in Cryptology — CRYPTO 2010*, volume 6223 of *Lecture Notes in Computer Science*, pages 333–350. Springer, Berlin, 2010. Cited in page 81.
- D. Kozen. Semantics of probabilistic programs. *J. Comput. Syst. Sci.*, 22(3):328–350, 1981. Cited in page 17.
- P. Laud. Semantics and program analysis of computationally secure information flow. In *Programming Languages and Systems, 10th European symposium on Programming, ESOP 2001*, volume 2028 of *Lecture Notes in Computer Science*, pages 77–91, Berlin, 2001. Springer. Cited in pages 133 and 135.
- A. K. Lenstra and H. W. L. Jr. *The Development of the Number Field Sieve*, volume 1554 of *Lecture Notes in Mathematics*. Springer, Berlin, 1993. Cited in page 81.
- A. K. Lenstra and E. R. Verheul. Selecting cryptographic key sizes. *J. Cryptology*, 14(4): 255–293, 2001. Cited in page 81.
- A. K. Lenstra, E. Tromer, A. Shamir, W. Kortsmit, B. Dodson, J. Hughes, and P. Leyland. Factoring estimates for a 1024-bit RSA modulus. In *Advances in Cryptology – ASIACRYPT 2003*, volume 2894 of *Lecture Notes in Computer Science*, pages 55–74, Berlin, 2003. Springer. Cited in page 82.
- X. Leroy. Formal certification of a compiler back-end, or: programming a compiler with a proof assistant. In *33rd ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages, POPL 2006*, pages 42–54, New York, 2006. ACM. Cited in pages 4, 48, and 131.
- M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM J. Comput.*, 17(2):373–386, 1988. Cited in pages 59 and 62.
- U. Maurer. Unifying zero-knowledge proofs of knowledge. In *Progress in Cryptology – AFRICACRYPT 2009*, volume 5580 of *Lecture Notes in Computer Science*, pages 272–286. Springer, 2009. Cited in pages 122 and 130.
- C. Meadows. Formal methods for cryptographic protocol analysis: emerging issues and trends. *IEEE J. Sel. Areas Commun.*, 21(1):44–54, 2003. Cited in page 133.
- D. Nowak. A framework for game-based security proofs. In *9th International conference on Information and Communications Security, ICICS 2007*, volume 4861 of *Lecture Notes in Computer Science*, pages 319–333, Berlin, 2007. Springer. Cited in page 134.
- T. Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In *Advances in Cryptology – CRYPTO 1992*, volume 740 of *Lecture Notes in Computer Science*, pages 31–53. Springer, 1993. Cited in page 120.

- T. Okamoto and D. Pointcheval. The gap-problems: A new class of problems for the security of cryptographic schemes. In *4th International workshop on Practice and Theory in Public Key Cryptography, PKC 2001*, volume 1992 of *Lecture Notes in Computer Science*, pages 104–118, Berlin, 2001a. Springer. Cited in page 10.
- T. Okamoto and D. Pointcheval. REACT: Rapid Enhanced-Security Asymmetric Cryptosystem Transform. In *Topics in Cryptology – CT-RSA 2001*, volume 2020 of *Lecture Notes in Computer Science*, pages 159–174, Berlin, 2001b. Springer. Cited in page 88.
- L. C. Paulson. The inductive approach to verifying cryptographic protocols. *J. of Comput. Secur.*, 6(1-2):85–128, 1998. Cited in page 4.
- D. Pointcheval. Provable security for public key schemes. In *Advanced Courses on Contemporary Cryptology*, chapter D, pages 133–189. Birkhäuser Basel, 2005. Cited in pages 85, 101, 104, 108, and 109.
- N. Ramsey and A. Pfeffer. Stochastic lambda calculus and monads of probability distributions. In *29th ACM SIGPLAN-SIGACT symposium on Principles of Programming Languages, POPL 2002*, pages 154–165, New York, 2002. ACM. Cited in page 17.
- A. Roy, A. Datta, A. Derek, and J. Mitchell. Inductive proofs of computational secrecy. In *Computer Security – ESORICS 2007, 12th European symposium on Research In Computer Security*, volume 4734 of *Lecture Notes in Computer Science*, pages 219–234, Berlin, 2008. Springer. Cited in page 133.
- RSA Data Security, Inc. PKCS #1 v2.1: RSA encryption standard, June 2002. Cited in page 83.
- A. Sabelfeld and D. Sands. A per model of secure information flow in sequential programs. *Higher-Order and Symbolic Computation*, 14(1):59–91, 2001. Cited in page 18.
- C.-P. Schnorr. Efficient signature generation by smart cards. *J. Cryptology*, 4(3):161–174, 1991. Cited in pages 116 and 120.
- A. Shamir and E. Tromer. On the cost of factoring RSA-1024. *RSA CryptoBytes*, 6:10–19, 2003. Cited in page 82.
- V. Shoup. OAEP reconsidered. In *Advances in Cryptology – CRYPTO 2001*, volume 2139 of *Lecture Notes in Computer Science*, pages 239–259, Berlin, 2001. Springer. Cited in pages 88 and 109.
- V. Shoup. Sequences of games: a tool for taming complexity in security proofs. Cryptology ePrint Archive, Report 2004/332, 2004. Cited in pages 1, 7, 41, and 65.
- C. Sprenger and D. Basin. Cryptographically-sound protocol-model abstractions. In *21st IEEE Computer Security Foundations symposium, CSF 2008*, pages 115–129, Los Alamitos, Calif., 2008. IEEE Computer Society. Cited in page 133.
- J. Stern. Why provable security matters? In *Advances in Cryptology – EUROCRYPT 2003*, volume 2656 of *Lecture Notes in Computer Science*, pages 644–644, Berlin, 2003. Springer. Cited in page 1.
- The Coq development team. The Coq Proof Assistant Reference Manual Version 8.2. Online – <http://coq.inria.fr>, 2009. Cited in pages 2 and 21.
- Y. Zhang. The computational SLR: A logic for reasoning about computational indistinguishability. In *8th International conference on Typed Lambda Calculi and Applications, TLCA 2008*, volume 5608 of *Lecture Notes in Computer Science*, pages 401–415, Berlin, 2009. Springer. Cited in page 134.



# Index

- advantage
  - DDH, 7
  - IND-CCA2, 38
  - IND-CPA, 6
  - PRF, 58
  - PRP, 59
- adversary, 31
- CDH, *see* Computational Diffie-Hellman
- CertiCrypt, 2
  - features, 2–3
  - tactics, 47–50
- challenge set, 114
- claw-free permutation pair, 123
- complete partial order, 16
- Computational Diffie-Hellman, 10
- Coq proof assistant, 21
- CryptoVerif, 85, 133
- DDH, *see* Decision Diffie-Hellman
- dead code elimination, 48
- Decision Diffie-Hellman, 7
- deep embedding, 20
- Diffie-Hellman
  - Computational Diffie-Hellman, 10
  - Decision Diffie-Hellman, 7
  - Diffie-Hellman triple, 7, 122
- eager sampling, *see* lazy sampling
- EF-CMA security, 32–34
- efficient program, *see* PPT
- ElGamal, 7
- encryption
  - asymmetric scheme, 5
  - ElGamal, 7
  - Hashed ElGamal, 9
  - IND-CCA2 security, 35–39
  - IND-CPA security, 6
  - OAEP, 88–89
  - padding-based, 87
  - semantic security, 6
- environment, 25
- exact security, 2
- failure event, 53–56
- Failure Event Lemma, 55, 61, 102
- FDH, *see* Full-Domain Hash
- Feistel network, 88
- Fiat-Shamir protocol, 121
- Fubini’s theorem, 19
- Full-Domain Hash, 67
  - unforgeability, 69, 75
- Fundamental Lemma, 13, 53, 60, 74, 93, 102
- game-based proof, 1
- General Number Field Sieve, 81
- Guillou-Quisquater protocol, 121
- Hashed ElGamal, 9
- Honest-Verifier Zero-Knowledge, 113, 114
- IND-CCA2 security, 35–39, 88
- IND-CCA, *see* IND-CCA2 security
- IND-CPA security, 6, 90
- indistinguishability
  - under chosen-ciphertext attacks, 35
  - under chosen-plaintext attacks, 6
- Katz-Wang signature scheme, 84
- knowledge relation, 113
- lazy sampling, 11, 50, 60, 79
- least upper bound, 16
- lift, 18
- lossless program, 29
- measure monad, 16
- memory, 25
- negligible function, 6
- non-interference, 46
- OAEP, 88–89

- IND-CCA2 security, 101
- IND-CPA security, 90
- observational equivalence, 46–47
- one-time pad, 9, 49, 79, 119
- one-way permutation, *see* trapdoor permutation
- optimistic sampling, 13, 49, 94, 96
- padding-based encryption scheme, 87
- partial-domain one-way permutation, 101
- PFDH, *see* Probabilistic FDH
- PPT, *see* probabilistic polynomial-time
- Probabilistic FDH, 83
- probabilistic polynomial-time, 28–30
- probabilistic relational Hoare logic, 41–46
  - judgment, 42
- Probabilistic Signature Scheme, 83
- probability
  - distribution, 16
  - distribution support, 17, 44–45
  - formal definition, 28
  - measure, *see* measure monad
- proof of knowledge, 111
- proof-assistant, 4, 21
- provable security, 1
- PRP/PRF switching lemma, 57–66
- pseudorandom function, 58
- pseudorandom permutation, 58
- pWHILE, 20–28
  - expressions, 23
  - programs, 24–25
  - semantics, 25–27
  - syntax, 21
  - types, 22
- random oracle model, 11, 91
- range, 17
- RSA, 67, 81, 83, 88, 121
- Schnorr protocol, 120
- semantic equivalence, 42
- semantic security, 6
- $\Sigma$ -protocol, 111–114
  - AND-Combination, 125
  - OR-Combination, 126
  - $\Sigma^\phi$ -protocol, 116
- signature scheme, 32
  - EF-CMA security, 32–34
  - Full-Domain Hash, 67
  - Katz-Wang, 84
  - Probabilistic FDH, 83
- special exponent, 117
- special homomorphism, 117
  - direct product, 121
- sup, *see* least upper bound
- tactic, 5
- termination, 29, 45
- trapdoor permutation, 67
- unforgeability, *see* EF-CMA security
- unit interval, 15
- well-formed adversary, *see* adversary
- Zero-Knowledge proof, 111



## Certification Formelle de Preuves Cryptographiques Basées sur les Séquences de Jeux

**Résumé :** Les séquences de jeux sont une méthodologie établie pour structurer les preuves cryptographiques. De telles preuves peuvent être formalisées rigoureusement en regardant les jeux comme des programmes probabilistes et en utilisant des méthodes de vérification de programmes. Cette thèse décrit CertiCrypt, un outil permettant la construction et vérification automatique de preuves basées sur les jeux. CertiCrypt est implémenté dans l'assistant à la preuve Coq, et repose sur de nombreux domaines, en particulier les probabilités, la complexité, l'algèbre, et la sémantique des langages de programmation. CertiCrypt fournit des outils certifiés pour raisonner sur l'équivalence de programmes probabilistes, en particulier une logique de Hoare relationnelle, une théorie équationnelle pour l'équivalence observationnelle, une bibliothèque de transformations de programme, et des techniques propres aux preuves cryptographiques, permettant de raisonner sur les événements. Nous validons l'outil en formalisant les preuves de sécurité de plusieurs exemples emblématiques, notamment le schéma de chiffrement OAEP et le schéma de signature FDH.

**Mots clés :** Cryptographie, langages de programmation, complexité, vérification de programmes, preuves formelles, assistant à la preuve Coq

## Formal Certification of Game-Based Cryptographic Proofs

**Abstract:** The game-based approach is a popular methodology for structuring cryptographic proofs as sequences of games. Game-based proofs can be rigorously formalized by taking a code-centric view of games as probabilistic programs and relying on programming language techniques to justify proof steps. In this dissertation we present CertiCrypt, a framework that enables the machine-checked construction and verification of game-based cryptographic proofs. CertiCrypt is built upon the general-purpose proof assistant Coq, from which it inherits the ability to provide independently verifiable evidence that proofs are correct, and draws on many areas, including probability and complexity theory, algebra, and semantics of programming languages. The framework provides certified tools to reason about the equivalence of probabilistic programs, including a relational Hoare logic, a theory of observational equivalence, verified program transformations, and ad-hoc programming language techniques of particular interest in cryptographic proofs, such as reasoning about failure events. We validate our framework through the formalization of several significant case studies, including proofs of security of the Optimal Asymmetric Encryption Padding scheme against adaptive chosen-ciphertext attacks, and of existential unforgeability of Full-Domain Hash signatures.

**Keywords:** Cryptography, programming languages, complexity, program verification, formal proofs, Coq proof assistant

